Robust Aerocapture Guidance with High-Fidelity Modeling of Aerothermodynamic Uncertainty

by

Jens Anderson Rataczak

B.S., University of Notre Dame, 2020

M.S., University of Colorado, 2023

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Aerospace Engineering Sciences

2025

Committee Members:

Iain D. Boyd, Co-Chair

Jay W. McMahon, Co-Chair

Hisham K. Ali

Robyn L. Macdonald

Eric M. Queen

Rataczak, Jens Anderson (Ph.D., Aerospace Engineering Sciences)

Robust Aerocapture Guidance with High-Fidelity Modeling of Aerothermodynamic Uncertainty

Thesis directed by Prof. Iain D. Boyd and Prof. Jay W. McMahon

Aerocapture is a mass-efficient orbital insertion method that utilizes atmospheric drag to decelerate a spacecraft, making it an enabling technology for Ice Giant exploration. However, the severe hypersonic environments encountered during aerocapture create significant challenges due to the intrinsic coupling between a spacecraft's trajectory and its aerothermodynamic environment. Traditional spacecraft design approaches treat these subsystems separately, relying on iterative methods that neglect their direct interaction. This dissertation takes a step toward a more integrated treatment of hypersonic systems by developing robust aerocapture guidance frameworks that directly integrate aerothermodynamic uncertainty information obtained using high-fidelity computational fluid dynamics (CFD). First, a detailed uncertainty quantification (UQ) and global sensitivity analysis (GSA) is conducted at key points along a Neptune aerocapture trajectory, accounting for variations in chemical kinetic parameters and the freestream composition. This analysis yields high-fidelity uncertainty bounds for convective and radiative heating distributions, as well as aerodynamic coefficients. Next, the information obtained from the UQ/GSA study is distilled into a reduced-order model using Gaussian process regression (GPR), which enables the efficient propagation of aerothermodynamic uncertainty along an entire trajectory. Finally, two aerocapture guidance algorithms, both based in the principles of sequential convex programming (SCP), are developed. The first is the convex predictor-corrector aerocapture guidance (CPAG) algorithm, which repeatedly solves a deterministic, constrained optimal control problem. CPAG eliminates the need for discrete lateral logic to control out-of-plane error and seamlessly integrates system-level constraints, such as heat load and aerodynamic loading, into its corrective logic. The second takes a stochastic optimal control approach to the aerocapture guidance problem, in which aerothermodynamic, aerodynamic, and atmospheric uncertainty are explicitly modeled within the optimization problem. This ap-

proach simultaneously optimizes a feedback control law with uncertainty in mind. Both guidance algorithms are compared against the state of the art and are shown to robustly guide aerocapture vehicles to their target orbit in the presence of uncertainty. By directly integrating aerothermodynamic modeling with trajectory optimization and guidance, this dissertation helps bridge the gap between traditionally separate subsystems in hypersonic vehicle and mission design.

Dedication

To my parents.

Acknowledgements

First and foremost, I would like to thank my advisors Dr. Iain Boyd and Dr. Jay McMahon for their mentorship, advice, and overall support during my graduate studies. They graciously offered to co-advise me, encouraging me to pursue multiple research topics when I presented them with my dual interests in hypersonics and GN&C. I consider myself incredibly lucky to have had the opportunity to study under them, and I would not be where I am today without their support. Furthermore, I would like to thank Dr. Hisham Ali, Dr. Robyn Macdonald, and Dr. Eric Queen for serving on my committee and providing valuable feedback on my research.

During my time at CU, I was fortunate to be a part of both NGPDL and ORCCA, and everyone that I had the pleasure of interacting with in each group truly made my time so much more enjoyable. I owe special thanks to Dr. Ross Chaudhry, Dr. Tim Aiken, and Dr. Davide Amato for helping further my technical understanding on a variety of topics and improving my research overall. From NGPDL, I would like to thank Dr. Tommy Kava, Dr. Amin Taziny, Dr. Jenny Horing, Dr. Pawel Sawicki, Mitch Wall, Kal Monroe, Nick Carter, Charlie Lipscomb, Cate Leszcz, Matt Stasiukevicius, and Marisa Petrusky for both their helpful technical discussions and their comradery and friendship. Additionally, I would like to thank Dr. Don Kuettel, Dr. Spencer Boone, Dr. Ken Oguri, Dr. Jacopo Villa, Dr. Ken Kuppa, Zach Donovan, Dillon Waxman, Kian Shakerin, Grace Calkins, and Rachel Cueva from ORCCA for their friendship, conversation, and mutually beneficial technical discussions.

Beyond my two research groups, I would like to thank the members of the third floor FSM office space. The presence each of you brought to the office made going to work far more enjoyable,

and I treasure the conversations we have had and memories we have made. In addition to those already listed, Sammy Sheppard, Tomaz Remec, Dr. Nils Wunsch, Jeff Hadley, and Connor Morency all made my experiences at CU better.

During my Ph.D., I was fortunate to be funded through the NASA Space Technology Graduate Research Opportunity (NSTGRO) fellowship. Through this award, I was able to visit several NASA centers and collaborate with some of the brightest minds in aerospace. I would particularly like to thank Jeff Hill, Cole Kazemba, and Dr. Joey Schulz for supporting me from an early stage in my graduate studies. Their initial support kick started my career in EDL, and I owe them a great deal of gratitude. I would also like to thank Dan Matz, Breanna Johnson, Gavin Mendeck, Dr. Som Dutta, Dr. Rohan Deshmukh, Dr. JB Scoggins, Dr. Tom West, and Dr. Brett Cruden for welcoming me during each of my NASA visits and providing thoughtful feedback and help on my research. Getting to collaborate with NASA researchers was a highlight of my graduate studies.

Finally, I would like to thank the people most important to me. To my Mom and Dad: Thank you. Your endless and unconditional support has helped me achieve things my childhood self would be proud of, and I owe you everything. I could not have done it without you. To my siblings, Natalie and Eric: Thank you. You have been the best sister and brother someone could ask for. I am proud to be your brother, and I cherish everything we have experienced together. To my cousins, Nicko, Charlie, and Jack: Thank you. Getting to grow up alongside you has been so rewarding, and your support throughout my life has been instrumental to who I am today. Finally, to my partner, Jenny: Thank you. You listened when I needed someone to talk to, you selflessly celebrated my victories, and you sat with me during my lowest times. You supported me through every step of this journey, and I am so lucky to have gotten to go through it with you by my side.

This research was supported by funding through the NSTGRO fellowship award number 80NSSC22K1186. This work utilized the Blanca condo computing resource at the University of Colorado Boulder. Blanca is jointly funded by computing users and the University of Colorado Boulder. This work also utilized the Alpine high-performance computing resource at the University of Colorado Boulder. Alpine is jointly funded by the University of Colorado Boulder, the University

of Colorado Anschutz, Colorado State University, and the National Science Foundation (award 2201538). Lastly, this work utilized the NASA Advanced Supercomputing (NAS) facilities. NAS is a part of the High End Computing Capability (HECC) project and is funded by NASA.

Contents

Chapter

Tables

Table

Figures

Figure

"In the face of overwhelming odds, I'm left with only one option: I'm gonna have to science the shit out of this."

— Mark Watney

Chapter 1

Introduction

"To confine our attention to terrestrial matters would be to limit the human spirit."
— Stephen Hawking

On July 24th, 1969, the Apollo 11 spacecraft reentered the Earth's atmosphere, returning astronauts Neil Armstrong, Buzz Aldrin, and Mike Collins from humanity's first Lunar landing. The sole goal of the spacecraft was to deposit these three men safely in the Pacific Ocean, decelerating from an initial velocity of 11.03 km/s to rest. To do so, the vehicle had to actively guide itself through the Earth's atmosphere, requiring the precise management of the dissipation of the vehicle's kinetic energy. If too much energy is dissipated, the vehicle fails to reach its intended destination, but insufficient energy dissipation will result in the spacecraft impacting the surface at lethal speeds. According to the conservation of energy, energy cannot be created or destroyed, only transferred or transformed from one form to another. So, what happened to the nearly 338 GJ of initial kinetic energy of the spacecraft as it traversed Earth's atmosphere?

The answer lies in the severe aerothermodynamic phenomena that govern hypersonic flight. As the Apollo Command Module plunged into the atmosphere, the air molecules surrounding the spacecraft were rapidly decelerated and compressed, forming a strong shock wave around the vehicle that transferred the majority of the gas' kinetic energy into thermal energy. Due to the immense energy in the freestream, this resulted in temperatures within the shock layer exceeding the surface of the Sun. To survive this harsh environment, the spacecraft required a robust thermal protection system (TPS), which was designed to gradually burn away as the vehicle maneuvered itself through

the atmosphere using its guidance, navigation, and control (GNC) system. This ablative TPS not only absorbed and dissipated the extreme heat but also influenced the vehicle's aerodynamic profile, necessitating precise GNC adjustments to maintain trajectory accuracy.

This fundamental interplay between a vehicle's trajectory and its aerothermodynamics environment is inherent to hypersonic flight, particularly in applications such as planetary entry, aerocapture, and atmospheric reentry. The trajectory of a hypersonic vehicle dictates the rate at which energy is dissipated, directly influencing the heating environment and aerodynamic forces experienced by the vehicle. Conversely, the aerothermodynamic response — including shock formation, boundary layer behavior, and heat transfer — affects aerodynamic forces, which in turn alter the vehicle's trajectory. This two-way coupling is especially critical for vehicles relying on aerodynamic forces to guide themselves. Historically, this coupling has been treated using an iterative approach in which each subsystem (e.g., TPS, GNC, etc.) is treated individually, and the design of one is used to inform the design of another. There is, at present, a gap in the treatment of this coupling in space-exploration settings, and if the interplay between subsystems can be accounted for accurately, the overall system performance and design process could see significant benefit [1, 2]. The central aim of this dissertation is to take a step in the direction of filling this gap, and, while the application of this research is for aerocapture systems, the fundamental physics of the problem render it relevant to any hypersonic system.

1.1 Aerocapture

One such field that is a prime candidate for a more coupled approach is *aerocapture*, which is an orbital insertion technique that relies on atmospheric drag to reduce a spacecraft's energy. A general aerocapture concept of operations is illustrated in Figure 1.1, where an inbound hyperbolic orbit is converted into a bound, elliptical orbit by performing a single pass through a planet's atmosphere. Typically, two orbital correction burns, labeled as ΔV_1 and ΔV_2 in Figure 1.1, are scheduled to raise the periapsis out of the planet's atmosphere and clean up any errors in the apoapsis radius, respectively.

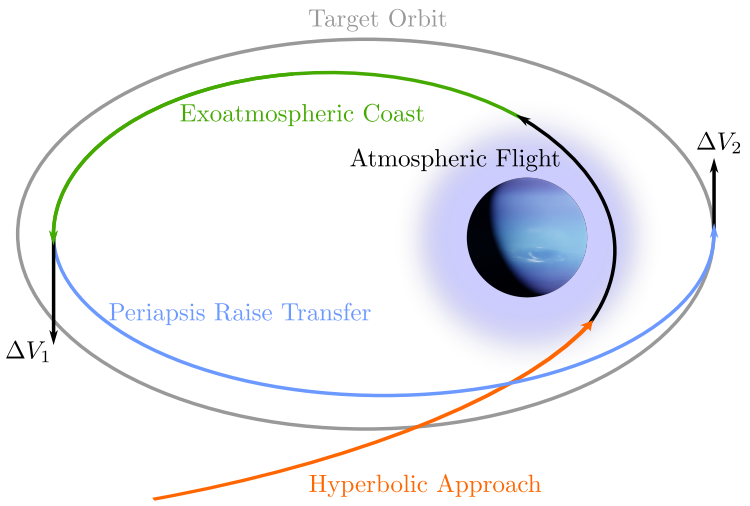

Figure 1.1: Aerocapture concept of operations.

Despite being proposed as early as the 1960s [3, 4, 5], aerocapture, to date, has never been used in a mission. The first proposal to use aerocapture was the Aeroassisted Flight Experiment (AFE) [6], which was developed to gather the required validation data for computational modeling of the nonequilibrium phenomena (discussed in more detail in Chapter 2) that could not be collected in ground-test facilities [7]. Unfortunately, the AFE was canceled due to mass and budget concerns and a refocusing of NASA's priorities to the Space Shuttle program. Aerocapture was proposed for a variety of missions following the AFE, including the Mars Surveyor orbiter [8] and Mars Sample Return [9], however, it was never executed due to factors including perceived risk, cost, and program cancellations. Later, in the early 2000s, a series of system-level analyses were performed for aerocapture concepts at Mars [10], Venus [11], Titan [12], and Neptune [13]. The last of these concluded that successful aerocapture at Neptune would require a substantial development of new technology, including higher lift-to-drag ratio vehicle geometries, deviating substantially from flight-proven heritage technology.

Consequently, when the most recent Planetary Science Decadal Survey [14] identified missions to the outer planets to be of high scientific interest, including a Uranus orbiter and probe as the highest priority flagship mission [15], the participating agencies proposed a mission concept involving a heritage blunt-body aeroshell using a propulsive orbit insertion [16]. Because they are far from

the Sun, a mission to the Ice Giants using standard, fully-propulsive orbit insertion techniques has a transit time ranging from 13 to 15 years. Reducing this transit time would require prohibitively large propulsive ΔV requirements for successful orbit insertion, significantly reducing the allowable payload mass and science capabilities. Aerocapture can alleviate many of these concerns, by enabling faster transit times and increased science payload mass [17, 18, 19, 20, 21], yet its perceived risk prevented it from being proposed as the baseline architecture for Ice Giant exploration. In the last few years, there has been substantial effort directed at demonstrating that the developmental requirements concluded in the original 2004 Neptune aerocapture study [13] are not necessary, and that successful Ice Giant aerocapture can be accomplished using heritage technology [22, 23].

Successful aerocapture requires the integration of numerous technologies. Two subsystems that are particularly relevant for aerocapture, and that will be the focus of this dissertation, are the *aerothermodynamics* and atmospheric *guidance* system of the spacecraft. Owed to the immense entry velocities, substantial aerothermal[1] analysis of the spacecraft is required to design its TPS, and a robust guidance system is needed to control the trajectory during the atmospheric flight phase, such that it can enter as closely to the target orbit as possible. Central to this dissertation is the explicit treatment of the coupling between these two domains, and the following sections provide motivating background for each.

1.1.1 Aerocapture Aerothermodynamic Modeling

Fundamentally, the discipline of aerothermodynamics is concerned with the analysis of the extreme environments produced during hypersonic flight, with the goal of accurately predicting the heating, pressure, and shear-stress profiles on the surface of the spacecraft during its flight through an atmosphere. For aerocapture, the high velocities produce substantial convective heating, and, if the velocity is high enough, radiative heating. Predicting the extent to which each of these sources contributes to the total heating environments requires a detailed understanding of the nonequilibrium thermochemistry in the high-temperature gas surrounding the vehicle. Efforts

[1] "Aerothermodynamic" and "aerothermal" are used interchangeably throughout this dissertation.

supporting this need have been progressing since the 1950s [24], and early success in the space-exploration domain was demonstrated during the Mercury, Gemini, and Apollo eras. A typical aerothermodynamic analysis of a spacecraft consists of experimental techniques, using wind tunnels, shock tunnels, and ballistic ranges, and computational methods, including computational fluid dynamics (CFD) and direct simulation Monte Carlo (DSMC) methods [25].

Each analysis must be tailored to the intended planetary destination, as the aerothermodynamics in an atmosphere like Earth differ substantially from that in the atmosphere of Mars or Neptune. Unlike Earth, the atmosphere of the Ice Giants consists primarily of hydrogen (H_2) and helium (He), with trace amounts of methane (CH_4). In order to determine the magnitude of the convective and radiative heating, it is necessary to predict the state of the gas behind the shock. Behind the shock wave formed around the entry vehicle, the molecules H_2 and CH_4 will dissociate into their constituent atoms H and C, and the short time scales produced in hypersonic flight produce conditions that may or may not be in thermochemical equilibrium. Therefore, it is necessary to determine the finite rate at which reactions occur in H_2/He/CH_4 mixtures.

In the 1970s, the first experimental efforts were directed at characterizing the rates of chemical reactions of He/H_2 mixtures in strongly shocked flows. Leibowitz [26] measured the ionization equilibrium times in a 21% H_2+79% He mixture for shock speeds up to 18 km/s. Shortly thereafter, Livingston and Poon [27] conducted a similar experiment using an 84% H_2+16% He mixture in which the shock speeds reached 40 km/s. In 2002, an experiment by Bogdanoff and Park [28] tested a 100% H_2 mixture for shock speeds up to 30 km/s. Each of these experiments predicted drastically different ionization rates, which Park [29] attributed to post-shock radiation contamination of the data. Despite the speculation over the validity of these rates, the rates measured by Leibowitz [26] and Leibowitz and Kuo [30] remain the standard kinetic model for simulating Ice Giant aerothermodynamics.

On the theoretical side, a series of studies were carried out by Park [24]. He investigated phenomena such as hydrogen Lyman-α radiation [31], and eventually produced a comprehensive theoretical model [29] that attempted to account for the experimental defects observed in the work

of Leibowitz and Livingston and Poon. These were then used to assess the aerothermodynamic environment of a Neptune aerocapture vehicle [32, 33], which concluded that the radiative-heating environments for the proposed Neptune aerocapture scenario [34] were high.

Until recently, very few studies have considered the effects of CH_4 on the aerothermodynamics of an Ice Giant aerocapture scenario. Some computational results have shown the addition of CH_4 can drastically increase the radiative-heating environments [32, 35], while experimental observations [36, 37] have shown a weaker radiation augmentation. This uncertainty is further exacerbated by the inability of all CFD-compatible thermochemical models to accurately predict the radiance measurements obtained experimentally [38]. Instead, Carroll et al. [39, 40] have shown that a detailed treatment of the state-specific kinetics of H_2/He is required to improve the agreements between numerical simulation and experimental measurements. Uncertainties in the aerothermal predictions are not desired because they can lead to unnecessarily heavy TPS, which can have a ripple effect throughout the spacecraft and may limit the payload mass allocation. The discrepancy between computational and experimental results, combined with the effect of aerothermodynamic uncertainty on an aerocapture vehicle design, are the driving motivations for the detailed uncertainty quantification (UQ) effort carried out in Chapter 3.

The role of uncertainty quantification in spacecraft aerothermodynamic design has become more prevalent over the last several years [41, 42, 43]. Modeling of aerothermodynamics phenomena requires hundreds of (potentially uncertain) input parameters for chemical kinetic rates, mixture and species transport properties, exchange of energy between translational and internal energy modes, and gas-surface interactions. Sophisticated global sensitivity analysis (GSA) techniques are shown to handle these high-dimensional problems efficiently [44, 42, 45, 46, 47], and the same techniques used in these works are employed in this dissertation. Apportionment of uncertainty to individual parameters using GSA provides a quantitative basis for parameters that should be investigated further which can lead to increased confidence in TPS and vehicle design. This dissertation also seeks to extend the benefits of uncertainty quantification beyond single trajectory points to an entire trajectory, techniques for which are further explored in Chapter 4.

1.1.2 Aerocapture Guidance

In a general aerocapture setting, the role of the onboard guidance system is to determine the required aerodynamic lift and/or drag force vector orientation, so that the spacecraft can exit the atmosphere and enter its target orbit without expending considerable amounts of propellant to perform orbital-correction maneuvers. Guidance commands are driven by the flight control law, which, by necessity, is closed-loop because the vehicle will likely deviate from the desired trajectory due to external disturbances and parametric uncertainties. The performance of an aerocapture guidance algorithm is typically evaluated statistically through repeated Monte Carlo simulations, which incorporate random variations in atmospheric conditions, initial vehicle states, and numerous other uncertain parameters. A "robust" aerocapture guidance algorithm is, thus, one that can properly respond to the uncertainty present in the environment and adapt its commands accordingly.

Despite never being flow, aerocapture guidance algorithm development has received substantial attention over the years. Early iterations of aerocapture guidance were derived from the Apollo entry guidance law, which uses a linear feedback control law obtained using the calculus of variations about a pre-determined reference trajectory [48]. This form of guidance, often referred to as a terminal point controller (TPC), is sensitive to the design of the reference trajectory and requires the tuning of the feedback gains. Despite the lightweight nature of TPC guidance, the gain-tuning process is problem dependent and can limit its generality and robustness.

An alternative to reference-tracking guidance are *predictor-corrector* algorithms. Predictor-corrector algorithms *predict*, either analytically or numerically, the remainder of the trajectory using the current control policy, which is then used to calculate the required *corrections* to the control, typically such that some performance index is minimized [49]. Due to the increased onboard computational requirements, predictor-correctors typically apply significant simplifications, including reduced sets of equations of motion and univariate, parametric control profiles. Analytic predictor-correctors (APC) utilize analytic solutions to the equations of motion within the prediction phase, which are limited to specific classes of flight, such as an equilibrium-glide condition [50] or ballis-

tic motion [51]. While APC benefits from the reduction in computation burden, the assumptions made to obtain the analytic solutions degrade the accuracy of the predictor, inherently limiting its robustness. The Hybrid Predictor-Corrector Aerocapture Scheme (HYPAS) combines an APC with reference-tracking logic, and was the guidance algorithm considered in NASA's extensive 2004 aerocapture systems-analysis campaign. Because it combines predictor-corrector logic and reference-tracking logic, HYPAS bypasses the need to design a reference trajectory in advance. However, its prediction accuracy suffers when the true motion of the vehicle is not well approximated by an equilibrium glide.

On the other hand, *numerical* predictor-correctors (NPC) make predictions by numerically integrating the spacecraft equations of motion onboard [52]. Although historically deemed impractical due to high computational demands, advancements in space-rated flight computer performance have made NPC guidance feasible, as demonstrated during the Artemis 1 skip entry. [53]. An extension of the Orion NPC, called PredGuid+A [54], seeks a constant bank-angle such that the ΔV required to correct the orbit is minimized. The fully numerical predictor corrector guidance (FNPAG) [55] relaxes the assumption of a constant bank angle, using a bang-bang bank-angle profile that follows from the solution to a deterministic optimal control problem. FNPAG has been the subject of numerous aerocapture studies [56, 57, 58, 59] and is considered to be the state of the art.

Because of the onboard computational requirements, existing NPC algorithms typically solve an *unconstrained* optimization problem, rendering the inclusion of system-level constraints – such as aerodynamic heating and loading – infeasible. A method for augmenting the FNPAG control profile to compensate for path constraints is presented in [52], however, doing so negates the optimality of form originally derived in [55]. With the increase in onboard computational capabilities, though, comes the ability to solve more complex optimization problems. This has been adopted heavily in the powered-descent community, in which convex-optimization based algorithms have been shown to be capable of running in real time onboard a spacecraft [60]. Convex-optimization based guidance has seen great success for a variety of applications including, but not limited to, powered descent guidance [61, 62, 63], atmospheric entry [64, 65, 66], low-thrust trajectory design [67], aerocapture

trajectory optimization [68, 69, 70, 71], and stochastic feedback control policy design [72, 73]. The inability to impose mission-critical constraints in an NPC setting, combined with the recent advancement in numerical optimization techniques, are the fundamental motivators for the advanced NPC guidance algorithm developed in Chapter 6.

Regardless of the underlying architecture, an aerocapture guidance algorithm has to be robust to uncertainty. Generally, the sources of environmental[2] uncertainty can be placed into one of two categories: aerodynamic and aerothermodynamic. Aerodynamic uncertainty stems from the lack of knowledge in a planet's atmospheric properties and the aerodynamic properties of the spacecraft itself, among others. Aerothermodynamic uncertainty, on the other hand, is produced by the uncertainty in the predictions of the convective- and radiative-heating environments along the aerocapture trajectory. As highlighted at the beginning of this chapter, these two disciplines are inherently coupled, yet in current entry and aerocapture guidance algorithm performance analyses, aerodynamic uncertainty is the only environmental uncertainty considered. In [74], atmospheric density uncertainty is *explicitly* included in the optimization of an aerocapture feedback control policy, highlighting the benefit of a stochastic treatment of the guidance problem. The lack of aerothermodynamic uncertainty inclusion in aerocapture guidance design, and the demonstrated success of stochastic-optimal-control based guidance techniques form the basis for the stochastic aerocapture guidance algorithm developed in Chapter 7.

1.2 Dissertation Overview

This dissertation takes a step toward a more direct integration of aerothermodynamic modeling and guidance algorithm design for hypersonic systems. Aerocapture is chosen as the central problem due to its unique challenges at the intersection of these disciplines. Despite its significance, an explicit, high-fidelity treatment of the coupling between aerothermodynamics and guidance remains largely unexplored in the hypersonics and space exploration research community. This work

[2] The *environment* classifier is used because uncertainty is also present in other aspects of an aerocapture mission, such as arrival state uncertainty and onboard navigation but are not relevant to the current discussion.

addresses key challenges on both fronts, with efforts focused on high-fidelity aerothermodynamic uncertainty quantification, the propagation of these uncertainties along an entire trajectory, and the development of sophisticated, robust aerocapture guidance algorithms that explicitly account for aerothermodynamic variability. By advancing the understanding of the interplay between aerothermodynamic uncertainties and guided hypersonic flight, this dissertation improves the state of the art in aerocapture modeling and control, contributing to the development of more reliable and efficient planetary entry systems.

1.2.1 Organization

This dissertation consists of eight chapters. Chapter 2, introduces the high-fidelity numerical tools used to model the aerothermodynamic environment of an aerocapture vehicle. It first explains key hypersonic flow phenomena and their impact on vehicle design, then details the computational tools used, including CFD and radiative-transport models. Chapter 2 also covers methods for assessing uncertainty in these models, identifying key contributing factors through uncertainty quantification (UQ) and global sensitivity analysis (GSA). Finally, it introduces Gaussian process regression (GPR) as a reduced-order modeling technique to efficiently estimate uncertainty along an entire aerocapture trajectory.

The primary aerothermodynamic-focused results are contained in Chapters 3 and 4. Chapter 3 applies the high-fidelity computational tools and uncertainty quantification (UQ) techniques to analyze the aerothermodynamic environment of an Ice Giant aerocapture trajectory. Baseline profiles for convective- and radiative-heating distributions and the axial-force coefficient are established, followed by UQ and GSA to identify key parameters that most influence predictions of each aerothermal quantity of interest. The effects of including $H_2{}^+$, a species typically omitted from Ice Giant aerothermodynamic analysis, are studied in an additional UQ/GSA analysis. In Chapter 4, a Gaussian-process-regression-based framework is developed to approximate the uncertainty intervals predicted in Chapter 3 along an entire aerocapture trajectory. The predictive capabilities of the Gaussian-process based models are enhanced through the use of invariant kernel design, which

enables physical constraints to be encoded within the surrogate structure.

Chapter 5 provides the preliminary material required for modeling aerocapture trajectories and the fundamental requirements for active guidance systems. The mathematical foundations of the guidance algorithms developed in the dissertation are reviewed, focusing on convex optimization and sequential convex programming (SCP) for both deterministic and stochastic trajectory optimization. In Chapter 6, the Convex Predictor-Corrector Aerocapture Guidance (CPAG) algorithm is presented, which improves on current NPC algorithms by using convex optimization in the correction phase. This allows for more complex control profiles and directly enforces aerothermodynamic constraints. An extensive Monte Carlo campaign is conducted, comparing the performance of CPAG to FNPAG. Then, a quantitative analysis of an energy-based apoapsis targeting expression capable of accounting for non-spherical gravitational perturbations is given, highlighting its benefit over conventional apoapsis expressions. Chapter 7 takes a stochastic optimal control approach to the aerocapture guidance problem, in which uncertainty in the atmospheric density, vehicle aerodynamics, and aerothermodynamic environment are explicitly considered in the design of a feedback control policy. In support of this explicit uncertainty treatment, the theoretical underpinnings of the effect of multiple, interacting, correlated noise sources are derived. Lastly, a summary, conclusion, and recommended future work are provided in Chapter 8.

1.3 List of Publications

The following publications have been produced by the author throughout this dissertation.

Peer-Reviewed Journal Articles

(J1) **J. A. Rataczak**, I. D. Boyd, J. W. McMahon, "Uncertainty Quantification and Sensitivity Analysis of Ice Giant Aerocapture Aerothermodynamics" *Journal of Thermophysics and Heat Transfer* (published 2025).

(J2) **J. A. Rataczak**, J. W. McMahon, I. D. Boyd, "Convex Predictor-Corrector Aerocapture Guidance," *Journal of Guidance, Control, and Dynamics* (published 2025).

(J3) **J. A. Rataczak**, J. W. McMahon, I. D. Boyd, "Stochastic Aerocapture Guidance," *Journal of Guidance, Control, and Dynamics* (in preparation).

(J4) **J. A. Rataczak**, I. D. Boyd, J. W. McMahon "Probabilistic Surrogate Modeling for Hypersonic Aerothermodynamics using Gaussian Process Regression," *Journal of Spacecraft and Rockets* (in preparation).

(J5) **J. A. Rataczak**, J. W. McMahon, I. D. Boyd, "Investigation of Surface-Catalycity Effects on Hypersonic Glide Vehicle Trajectory Optimization", *Journal of Spacecraft and Rockets* (published 2024).

(J6) **J. A. Rataczak**, D. Amato, J. W. McMahon, "Density Estimation for Entry-Guidance Problems using Deep Learning," *Journal of Guidance, Control, and Dynamics* (published 2025).

Conference Papers

(C1) **J. A. Rataczak**, I. D. Boyd, J. W. McMahon, "Surrogate Models for Hypersonic Aerothermodynamics and Aerodynamics using Gaussian Process Regression", AIAA SciTech Forum, Orlando, FL, January 2024.

(C2) **J. A. Rataczak**, J. W. McMahon, I. D. Boyd, "Predictor-Corrector Aerocapture Guidance using Convex Programming", AAS Guidance, Navigation & Control Conference, Breckenridge, CO, February 2024.

(C3) **J. A. Rataczak**, I. D. Boyd, J. W. McMahon, "Parametric Sensitivity Analysis and Uncertainty Quantification of Ice Giant Aerocapture Aerothermodynamics", AIAA Aviation Forum, Las Vegas, NV, August 2024.

(C4) **J. A. Rataczak**, I. D. Boyd, J. W. McMahon, "Stochastic Aerocapture Trajectory Optimization Under Aerothermodynamic Uncertainty," AIAA SciTech Forum, Orlando, FL, January 2025.

(C5) **J. A. Rataczak**, J. W. McMahon, I. D. Boyd, “Predictor-Corrector Aerocapture Guidance using Convex Programming,” AIAA SciTech Forum, Orlando, FL, January 2025.

(C6) **J. A. Rataczak**, D. Amato, J. W. McMahon, "Density Estimation for Entry-Guidance Problems using Deep Learning," AIAA SciTech Forum, Orlando, FL, January 2024.

(C7) **J. A. Rataczak**, J. W. McMahon, I. D. Boyd, “Reachability Analysis of a Hypersonic Glide Vehicle using Particle Swarm Optimization”, AIAA SciTech Forum, National Harbor, MD, January 2023.

(C8) K. Monroe, **J. A. Rataczak**, I. D. Boyd, "Hypersonic Flight Regimes for Electron Transpiration Cooling," AIAA Aviation Forum, Las Vegas, NV, July 2025 (accepted).

Chapter 2

Preliminaries - Aerothermodynamic Modeling and Uncertainty Quantification

"Any dogma, whatever it may be, is bound to be wrong at some point because it is a model of reality, not reality itself."
— Isaac Asimov (Foundation, 1951)

2.1 Introduction

This chapter serves two primary purposes. First, it seeks to provide the reader with a background of the high-fidelity numerical tools used to model the aerothermodynamic environments of an aerocapture vehicle. To set the stage, the chapter begins with a general overview of the aerothermodynamic phenomena in hypersonic flows and their importance to aerocapture vehicle design. In particular, a generalized description of the aerothermodynamic phenomena present in hypersonic flows is presented, and their relevance to the design of aerocapture vehicles is highlighted. With this context established, the chapter then describes the two main computational aerothermodynamic tools employed in this dissertation: a computational fluid dynamics (CFD) tool and a radiative-transport tool. The CFD tool section covers the mathematical models for chemical reactions, thermodynamics, transport properties, and boundary conditions, followed by an explanation of the radiative-transport tool.

Second, this chapter introduces the mathematical principles used to answer key questions about each of the aforementioned models' predictions. These questions are specifically: 1) What is the uncertainty in each model's predictions? 2) Which input parameters are contributing most to that uncertainty? 3) How can the cumulative effects of that uncertainty along a trajectory be

approximated? The first two of these questions have to do with the concepts of uncertainty quantification (UQ) and global sensitivity analysis (GSA), respectively. It is therefore necessary to present the theory behind polynomial chaos expansion (PCE) for UQ and GSA, including sparse point-collocation techniques and the Sobol' decomposition for sensitivity indices. The third question deals with the topic of reduced-order modeling, as it is computationally intractable to perform a detailed UQ analysis at every point along an aerocapture trajectory. Specifically, this research uses Gaussian process regression (GPR), which is described in the final part of this chapter.

2.2 Physical Characteristics of Hypersonic Aerothermodynamics

This section explores the physical phenomena in hypersonic flows and their relevance to aerocapture vehicle design and analysis. As noted in Chapter 1, the motion of an aerocapture vehicle and its aerothermodynamic environment are governed by the conservation of mass, momentum, and energy. When moving at hypersonic speeds, the amount of kinetic energy in the surrounding fluid is extraordinarily high. Upon crossing the strong shock waves inherent to hypersonic flight, the gas particles slow down, and the majority of their kinetic energy is converted into internal energy. This surge in internal energy kick-starts many (if not all) of the defining characteristics of high Mach-number flight, some of which are illustrated in Figure 2.1.

Consider, for example, a molecule of hydrogen (H_2). When this molecule crosses the shock wave in Figure 2.1, its velocity is slowed dramatically, and its internal energy begins to increase. But what are these internal energy modes, and how does the amount of energy in them increase? Generally, the constituent particles of a gas interact by colliding with each other, and these collisions transfer energy between the particles and internal energy modes. From statistical mechanics, the following five internal energy modes may be defined.

(1) Translational: the translational mode describes the kinetic energy of a particle and is related to the random molecular motion of the gas particles. From classical mechanics, it is expressed as $\varepsilon_{tr} = \frac{1}{2}mv^2$. All particles (atoms and molecules) have translational energy.

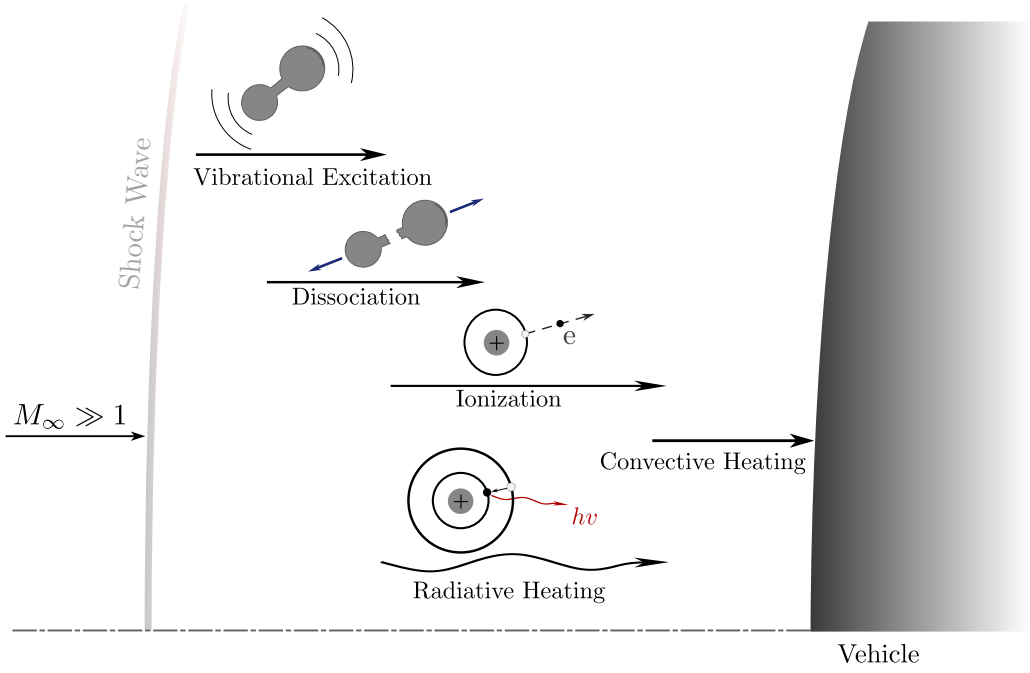

Figure 2.1: Typical processes occurring behind a hypersonic shock wave.

(2) Rotational: the rotational mode describes the energy related to the rotation of a molecule and is typically described using a rigid rotor model [75]. Generally, diatomic molecules (e.g. H_2) have two rotational degrees of freedom, while triatomic or larger molecules (e.g CH_4) have three. Only molecules have rotational energy.

(3) Vibrational: the vibrational energy describes the oscillations between atoms in molecules due to the change in the length of the chemical bonds and is typically described using a harmonic oscillator model [75]. Only molecules have vibrational energy.

(4) Electronic: the electronic energy mode describes the amount of energy contained in a particle's electronic states. Bound electrons in atoms or molecules occupy discrete (quantized) energy levels. This internal energy mode is typically only activated at high temperatures (>10,000 K), and quantifying the populations of the excited electronic states is necessary for predicting the radiative-heating environment. All atoms and molecules have electronic energy.

(5) Free-Electron: electrons that are no longer bound to an atom also have energy attributed to their translational motion. Free electrons are produced through *ionization* reactions which typically require large amounts of energy.

Returning now to Figure 2.1, a more apt description of the processes occurring behind the strong shock wave can be given. As the hydrogen molecule crosses the shock wave, its translational energy is transferred into the internal energy modes via collisions with other particles. At standard conditions, the rotational mode is fully excited and in equilibrium with the translational mode. At high temperatures, the next of these internal energy modes to typically *activate* is the vibrational mode, which is referred to as "vibrational excitation." When a particle is vibrationally excited, the amount of energy contained within the molecule begins to approach the strength of the chemical bonds, and at some point, the molecule will collide with another particle and *dissociate* into its constituent atoms. Dissociation tends to occur preferentially for molecules in vibrationally excited states, as their internal energy is more adjacent to the dissociation enthalpy of the molecule.

As the gas continues to dissociate, more and more particles will be produced (e.g. the dissociation of a hydrogen molecule produces two hydrogen atoms) which increases the rate at which the particles collide with one another. If the total amount of energy in the gas is sufficiently high, these collisions will eventually begin to populate the excited electronic states of the particles. Similar to dissociation, the removal of an electron from an atom, or *ionization*, preferentially occurs from electronically excited configurations because the high-lying electronic states are more energetically adjacent to the ionization potential. The higher-energy, bound electron thus requires less additional energy to increase its energy beyond the ionization potential. Excited bound electrons reside in higher orbitals (i.e. further from the nucleus) which are less stable than lower-lying electron states, so they eventually relax into lower-energy electronic configurations. The conservation of total energy consequently releases energy from the particle in the form of a photon. Emitted photons can traverse the shock layer and be absorbed by the vehicle – a phenomenon known as *radiative* heating. While not typically a concern for lower velocities, the high entry velocities of Ice Giant aerocapture

can produce substantial radiative heating. Radiative emission/absorption can also come in other forms (i.e. bound-bound, bound-free, free-free), but they are beyond the scope of this discussion.

Moving now to the right-hand side of Figure 2.1, the high-temperature gas encounters a relatively cool surface, creating a steep temperature gradient. This form of energy transfer is known as *convective* heating – a local phenomenon produced by the direct contact of the vehicle surface with the high-temperature gas. Convective heating can also be augmented by the release of energy produced by chemical bonds being formed as the dissociated/ionized gas recombines through the boundary layer. A more rigorous characterization of both the convective and radiative heating for Ice Giant aerocapture vehicles is found in Chapter 3.

In general, it is helpful to think about the various energy modes as "buckets" into which some portion of the total energy can be stored. As more modes activate, more buckets become available, and quantifying how much and the rate at which each fills is key to modeling the aerothermodynamic environment of aerocapture vehicles. For example, the amount of hydrogen that has dissociated from H_2 into H is a crucial quantity to know if the radiative-heating environment is to be predicted confidently, as high-temperature H_2 contributes strongly to the total radiation. To develop this quantitative understanding, it is essential to examine the principles of thermochemical equilibrium and nonequilibrium, which govern the behavior of high-temperature gas flows in hypersonic regimes.

2.2.1 Thermochemical Equilibrium

In the context of hypersonic applications, the term "equilibrium" generally signifies that a process or quantity is not changing with time, or any changes that do occur happen in an infinitesimal amount of time. From the previous section, changes in a given process are governed by particle collisions, so an equilibrium condition implies that a large number of collisions for a given process occur before any other flow properties change. For hypersonic vehicles, it is necessary to distinguish *thermal* equilibrium from *chemical* equilibrium. The former deals with the distribution of energy among the various energy modes described in the previous section, while the latter deals with the chemical composition of the gas. While these concepts are distinct, they are not independent.

Instead, they are inherently coupled, as the distribution of energy will affect the types of chemical reactions that can occur (e.g. dissociation occurs preferentially from vibrationally excited states).

2.2.1.1 Thermal Equilibrium

From quantum mechanics, the energy levels of each energy mode are "quantized", meaning that they exist in discrete, specific values, not in a continuous range. The translational and rotational modes, however, are typically treated in a classical sense, as the energy spacing between the possible energy levels for each is small. The other internal energy modes have relatively larger energy spacing and are therefore treated discretely. For each mode, the distribution of energy within each of the states is given by the Boltzmann distribution

$$\frac{n_i}{n} = \frac{g_i \exp\left(-\frac{\epsilon_i}{k_B T}\right)}{\sum_j g_j \exp\left(-\frac{\epsilon_j}{k_B T}\right)} \tag{2.1}$$

where n_i is the number density of particles in the i-th state, $n = \sum_i n_i$ is the total number density, ϵ_i is the energy level of the i-th state, $k_B = 1.38 \times 10^{-23}$ J/K is the Boltzmann constant, and g_i is the *degeneracy* which accounts for scenarios in which multiple states occupy the same energy level. The final variable in the Boltzmann distribution is the temperature T. This brings to light an important, physical interpretation of what *temperature* means – it quantifies the amount of energy in a given energy state at *equilibrium*. Consequently, *thermal* equilibrium signifies that the distribution of energy across all modes is in equilibrium such that $T_{\text{trans}} = T_{\text{rot}} = T_{\text{vib}} = T_{\text{elec}} = T_{e^-}$.

The denominator of the Boltzmann distribution is called the *partition function*

$$Q(T) = \sum_i g_i \exp\left(\frac{\epsilon_i}{k_B T}\right) \quad . \tag{2.2}$$

Physically, the partition function is the number of available quantum states where each term is proportional to the number of particles in a given energy group. Statistical mechanics assumes that the energy modes may be treated separately, each with their own partition function, so the total partition function is expressed as the product of the individual partition functions for all contributing energy modes.

While the Boltzmann distribution is particularly useful for quantifying the energy distribution within a given energy mode, it does not account for *ionization*. Ionization cannot be neglected in this research because, as will be highlighted in Chapter 3, ionization of atomic hydrogen plays a key role in radiative- and convective-heating environments for Ice Giant aerocapture. Specifically, the Boltzmann distribution only deals with the distribution of excited states within a single ionization state (e.g. H or H^+, not H and H^+), which is an insufficient description in the context of radiative-heating predictions. The Saha-Boltzmann equation extends the Boltzmann distribution to account for ionization and relates the number density of two ionization states at equilibrium

$$\frac{n_i}{n_+} = \frac{g_i n_e}{2Q_+}\left(\frac{h^2}{2\pi m_e k_B T_e}\right)^{3/2} \exp\left[\frac{-hc\left(\epsilon_i - \epsilon_{\text{ionize}}\right)}{k_B T_e}\right] \quad . \tag{2.3}$$

Here n_+ is the ion number density, n_e is the electron number density, n_i is the number density of the i-th excited state, Q_+ is the ion partition function, $m_e = 9.1091 \times 10^{-31}$ kg is the electron mass, $h = 6.6256 \times 10^{-34}$ J·s is Planck's constant, $c = 2.9979 \times 10^8$ m/s is the speed of slight, and ϵ_{ionize} is the ionization energy. Because ionization occurs preferentially from electronically excited configurations, it is often more accurate to assume that the excited states of a neutral species (e.g. H) are in equilibrium with the ions (e.g. H^+) [76]. Figure 2.2 compares the state populations predicted by Equations 2.1 and 2.3 for atomic hydrogen, where the electronic energy of level i is given as $\epsilon_i = \epsilon_{\text{ionize}}\left(1 - \frac{1}{i^2}\right)$. Here, $\epsilon_{\text{ionize}} = 13.6$ eV is the ionization potential of atomic hydrogen. The temperature is assumed to be $T = 10{,}000$ K, the total number density is taken to be $n = 10^{23}$ $1/\text{m}^3$, and the electron and ion number densities are both taken to be $n_+ = n_e = 10^{22}$ $1/\text{m}^3$. For these conditions, the Saha-Boltzmann line is below the Boltzmann line at all energy levels.

The location of the Saha-Boltzmann line relative to the Boltzmann line provides insight into the amount of chemical equilibrium at the local conditions. If there are fewer ions than at equilibrium, the gas is net ionizing, and the Saha-Boltzmann line will be below the Boltzmann line, as the excited states are underpopulated. On the other hand, if there are more ions than at equilibrium, the gas is net recombining, and the Saha-Boltzmann line will be above the Boltzmann line, as the excited states are overpopulated. At chemical equilibrium, Equations 2.1 and 2.3 are

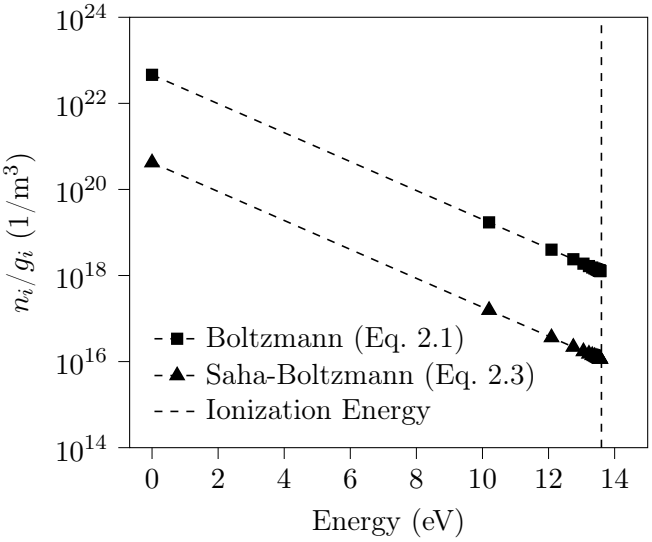

Figure 2.2: Comparison of Boltzmann and Saha-Boltzmann state-population distributions for atomic hydrogen.

equivalent and the Boltzmann and Saha-Boltzmann lines lie on top of each other.

When dealing with a single chemical species, the conditions for thermal equilibrium relate to the distribution of internal energy of that species. These conditions may easily be extended to multi-species mixtures (e.g. H and H^+ or H_2, H, He, etc.). Specifically, thermal equilibrium in a multi-species mixture implies that the internal energy modes are also in equilibrium between species, such that the internal energy distribution for each is described by the same temperature.

2.2.1.2 Chemical Equilibrium

Instead of dealing with the distribution of internal energy, *chemical* equilibrium describes the population distribution of individual chemical species within a multi-species mixture. As such, *thermochemical* equilibrium implies both thermal *and* chemical equilibrium. The chemical composition of a multi-species mixture is governed by the types of chemical reactions that can occur at the local thermodynamic conditions. For instance, H is formed from a variety of reactions, such as the dissociation of H_2 and the recombination of H^+ into H. A general chemical reaction may be

written as

$$\text{Reactants} \leftrightarrow \text{Products} \tag{2.4}$$

where the reactants and products are any of the species considered in the mixture. The condition for chemical equilibrium is that the rates of the forward and reverse directions of Equation 2.4 are equal and opposite such that the net rate is zero. The forwards and backwards rates are typically expressed using rate coefficients, and the ratio of the two rate coefficients is known as the *equilibrium constant*. More information on how the rate coefficients and equilibrium constants are calculated are provided in Section 2.3. An important detail regarding chemical equilibrium is that the chemical composition is directly linked to the local thermodynamic state (e.g. temperature and pressure). Therefore, if a chemical equilibrium condition is assumed, it signifies that when temperature, pressure, and other flow properties change, the chemical composition is changed instantaneously to the composition defined by the local thermodynamic state. Figure 2.3 shows the equilibrium composition of an H_2/He mixture for a range of temperatures and a pressure of 1 atm.

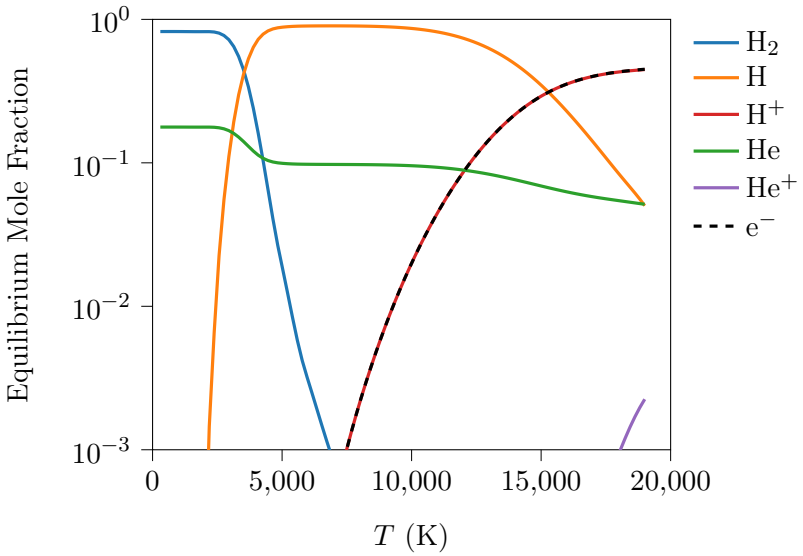

Figure 2.3: Equilibrium composition of H_2/He mixture at $P = 1$ atm.

At these conditions, molecular hydrogen begins to dissociate around $T = 3000$ K and is completely dissociated once the temperature exceeds 10,000 K. Atomic hydrogen begins to ionize

around $T = 6000$ K, and H^+ is the dominant species for $T > 15,000$ K. Note that helium ions do not start to appear until very high temperatures, as its ionization potential (24.6 eV) is much larger than that of atomic hydrogen (13.6 eV). Despite the small amount of He^+, the neutral helium mole fraction decreases with temperature because the number of hydrogen atoms and ions is increasing. The electron mole fraction is essentially equal to that of H^+ because H^+ is the dominant ionized species. Understanding the conditions at which certain species persist can help reduce the number of chemical species tracked in a numerical simulation, as those that are not expected to form can be omitted from the thermochemical model.

2.2.2 Thermochemical Nonequilibrium

The previous discussion on thermochemical equilibrium assumes that any changes in processes occur instantaneously or, equivalently, sufficient time has passed for collisions between particles to produce equilibrium conditions. One can define the equilibrium condition for a given process using a characteristic time τ_x

$$\tau_x \ll \tau_f \tag{2.5}$$

where τ_f is the flow characteristic time which is inversely proportional to the freestream velocity. This definition acknowledges that the processes occurring behind a hypersonic shock wave take a finite amount of time to proceed but assumes that the time scales are sufficiently disparate that they may be considered independently. However, in most hypersonic applications, including aerocapture, the flow characteristic time approaches the process characteristic time, and the finite-rate nature of the process must be accounted for explicitly. When these time scales become comparable, the fundamental assumption of equilibrium breaks down, and *nonequilibrium* effects must be considered.

Recall that thermal equilibrium allows for the energy distribution of all chemical species to be described by a single temperature. Thermal nonequilibrium thus arises when the internal energy distributions of all species have not had sufficient time to relax. During this time, multiple temperatures are required to describe the flow. A common example of a multi-temperature model is the two-temperature model [77]. In this model, the translational and rotational modes are assumed

to be in equilibrium such that $T_{\text{trans}} = T_{\text{rot}} = T_{tr}$, and the vibrational, electronic, and free-electron modes are assumed to be in equilibrium such that $T_{\text{vib}} = T_{\text{elec}} = T_{\text{e}^-} = T_{ve}$. This particular grouping of energy modes is chosen because the translational and rotational modes respond rapidly to changes in the local conditions, while the other internal energy modes are relatively slower to relax. More complex multi-temperature models may be used if the conditions under investigation are suspected to violate the two-temperature model's assumptions [78]. Additionally, each species could be assigned its own T_{ve} and T_{tr}, however, this is usually not done in practice, as tracking species-specific temperatures would drastically increase the computational complexity.

Similarly, chemical nonequilibrium arises when the forward and reverse rates of chemical reactions within the flowfield are not balanced. Depending on the local thermodynamic conditions, this results in a net production or depletion of certain species which consequently affects the internal energy distribution of the mixture. Because of their effect on the flowfield properties, capturing the combined effects of thermal and chemical nonequilibrium is essential for high-fidelity predictions of critical quantities such as convective and radiative heating. Achieving this level of accuracy typically requires advanced computational models, which are the focus of the next two sections.

2.3 Computational Fluid Dynamics

The numerical aerothermodynamic analysis in this dissertation uses the CFD code LeMANS [79]. LeMANS is a finite-volume, multidimensional, parallelized CFD code designed to model the thermochemical nonequilibrium phenomena relevant to hypersonic flows. It has been thoroughly benchmarked against existing state-of-the-art hypersonic CFD solvers such as DPLR, LAURA, and US3D. For the original documentation on LeMANS, the reader is referred to Ref. [79], and further implementation details may be found in Ref. [80]. An overview of LeMANS is provided in this section, along with models for thermal and chemical nonequilibrium in hypersonic flowfields. Additional detail is provided for the models in LeMANS that are considered by the uncertainty quantification and sensitivity analysis performed in Chapter 3.

2.3.1 Navier-Stokes Equations

LeMANS solves the compressible, reacting Navier-Stokes equations that account for an arbitrary number of species ns and thermal and chemical nonequilibrium. The flow is modeled under the continuum assumption. Translational and rotational modes of all species are assumed to be in equilibrium and described by a single temperature T_{tr}, as the rotational mode only requires a few collisions to equilibrate with the translational mode. The vibrational, electronic, and free electron modes are assumed to be in equilibrium and described by a single temperature T_{ve}.

With these assumptions in mind, the Navier-Stokes equations are stated as

$$\frac{\partial \boldsymbol{Q}}{\partial t} + \nabla \cdot (\boldsymbol{F} - \boldsymbol{F}_v) = \boldsymbol{S}_{cv} \tag{2.6}$$

where $\boldsymbol{Q}$ is a vector of conserved variables, $\boldsymbol{F}$ and $\boldsymbol{F}_v$ are the inviscid and viscous fluxes, respectively, and $\boldsymbol{S}_{cv}$ is a vector of the source terms. The vector of conserved variables is explicitly stated as

$$\boldsymbol{Q} = \begin{Bmatrix} \rho_1 \\ \vdots \\ \rho_{ns} \\ \rho u \\ \rho v \\ \rho w \\ E \\ E_{ve} \end{Bmatrix} \quad . \tag{2.7}$$

Here, ρ_s is the density of species s, ρ is the bulk density, u, v, and w are the bulk velocity components, and E and E_v are the total and vibrational-electronic-electron energy per unit volume of mixture,

respectively. The vector of source terms is

$$\boldsymbol{S}_{cv} = \begin{Bmatrix} \dot{w}_1 \\ \vdots \\ \dot{w}_{ns} \\ 0 \\ 0 \\ 0 \\ 0 \\ \dot{w}_{ve} \end{Bmatrix} \tag{2.8}$$

where $\dot{w}_s$ is the net production of species s, and $\dot{w}_{ve}$ is the vibrational-electronic-electron energy source term. The inviscid and viscous flux vector components in the x-direction are given by

$$\boldsymbol{F} = \begin{Bmatrix} \rho_1 u \\ \vdots \\ \rho_{ns} u \\ \rho u^2 + p \\ \rho u v \\ \rho u w \\ (E+p)\, u \\ E_{ve} u \end{Bmatrix} \qquad \boldsymbol{F}_v = \begin{Bmatrix} -J_{x,1} \\ \vdots \\ -J_{x,ns} \\ \tau_{xx} \\ \tau_{xy} \\ \tau_{xz} \\ \tau_{xx} u + \tau_{xy} v + \tau_{xz} w - (q_{tr,x} + q_{ve,x}) - \sum (J_{x,s} h_s) \\ -q_{ve,x} - \sum (J_{x,s} e_{ve,s}) \quad . \end{Bmatrix} \tag{2.9}$$

Here, p is the pressure, τ_{ij} are the viscous stress components, and $q_{tr,i}$ and $q_{ve,i}$ are the translational-rotational and vibrational-electronic-electron heat fluxes in the i-th direction, respectively. Furthermore, h_s is the specific enthalpy of species s, and $J_{i,s}$ is diffusion flux in the i-direction for species s. Flux vectors for the y and z directions are similar. Spatial integration is performed over the inviscid and viscous fluxes using the finite-volume method, and the time integration is performed implicitly using the Backward Euler scheme. Temporal integration is first-order accurate, and spatial integration is second-order accurate [79].

2.3.2 Thermal and Chemical Nonequilibrium

Modeling of thermal nonequilibrium in LeMANS is accomplished using Park's two-temperature model [77]. Consequently, the energy levels of each mode are assumed to follow a Boltzmann distribution at the corresponding temperature. Exchange of energy between these modes is captured through the source term for the vibrational-electronic-electron equation in Equation 2.8

$$\dot{w}_{ve} = \dot{w}_{epg} + \dot{w}_{c-v} + \dot{w}_{h-e} + \dot{w}_{e-i} + \dot{w}_{t-v} \quad . \tag{2.10}$$

Here, $\dot{w}_{epg}$ models the work done on the electrons by the electronic field induced by the electron pressure gradient. The next term, $\dot{w}_{c-v}$ is the vibrational energy-electronic-electron energy added or removed by chemical reactions. A non-preferential dissociation model is assumed in this work, in which molecules are created or destroyed at the average vibrational energy of the mixture. Energy transfer between heavy particles and electrons is modeled using $\dot{w}_{h-e}$, and energy removed from electron-impact ionization reactions is modeled using $\dot{w}_{e-i}$. Details for each of these source terms are found in [79].

An explicit description of the vibrational-relaxation source term is given because the aerothermodynamic uncertainty quantification study performed in this research assumes the Millikan-White coefficients are uncertain. Relaxation between the translational-rotational modes and vibrational-electronic-electron modes is modeled using the Landau-Teller equation [81]

$$\dot{w}_{t-v} = \sum_s \rho_s \frac{e^*_{v,s} - e_{v,s}}{\tau_s} \quad . \tag{2.11}$$

Here, $e^*_{v,s}$ and $e_{v,s}$ are the vibrational energy per unit mass of species s evaluated at T_{tr} and T_{ve}, respectively, and τ_s is the vibrational relaxation time of species s. Vibrational relaxation times are expressed as

$$\tau_s = < \tau_s > + \tau_{ps} \tag{2.12}$$

where τ_{ps} is Park's high-temperature correction [79], and the molar averaged Landau-Teller relaxation time is

$$< \tau_s > = \frac{\sum_r X_r}{\sum_r X_r / \tau_{sr}} \quad . \tag{2.13}$$

Here, X_r is the mole fraction of species s, and τ_{sr} is given by the Millikan-White expression [82]

$$\tau_{sr} = \frac{101325}{p} \exp\left[a_{sr}\left(T^{-1/3} - b_{sr}\right) - 18.42\right] \tag{2.14}$$

where a_{sr} and b_{sr} are interaction-dependent constants usually fit from experimental data. The factor of 101325 converts the pressure p into atmospheres. Park's high-temperature correction term τ_{ps} is introduced to account for the fact that the baseline Millikan-White curve fits underpredict the relaxation times at high temperatures [24].

Chemical nonequilibrium in LeMANS is modeled using a finite-rate chemistry model. A generic reaction may be specified as

$$\sum \alpha_s \left[S\right] \leftrightarrow \sum \beta_s \left[S\right] \tag{2.15}$$

where $[S]$ represents one of the chemical species, and α and β are the stoichiometric coefficients. The source term in Equation 2.8 for the production of species s is given by

$$\dot{w}_s = M_s \sum_k \dot{w}_{sk} \tag{2.16}$$

where M_s is the molar mass of species s and $\dot{w}_{sk}$ is the production of species s by reaction k

$$\dot{w}_{sk} = (\beta_{sk} - \alpha_{sk}) \left[10^3 k_{fk} \prod_j \left(10^{-3} \frac{\rho_j}{M_j}\right)^{\alpha_{jk}} - 10^3 k_{bk} \prod_j \left(10^{-3} \frac{\rho_j}{M_j}\right)^{\beta_{jk}}\right] \tag{2.17}$$

In Equation 2.17, k_{fk} and k_{bk} are the forward and backward reaction-rate coefficients, respectively. LeMANS models forward reaction rates coefficients using a modified-Arrhenius form

$$k_f = A_f T_f^{\eta_f} \exp\left(-D_f/T_f\right) \quad \left[\mathrm{cm}^3 \cdot \mathrm{mol}^{-1} \cdot \mathrm{s}^{-1}\right] \tag{2.18}$$

where A_f, η_f, and D_f are reaction-specific constants, and T_f is the controlling temperature of the forward reaction. Backward reaction rate coefficients are obtained using the principle of detailed balance in which the forward and backward rate coefficients are related through the equilibrium constant K_{eq} evaluated at the backward controlling temperature T_b

$$k_b(T_b) = \frac{k_f(T_b)}{K_{eq}(T_b)} \quad . \tag{2.19}$$

For the hydrogen-helium-methane mixtures studied in this work, a Gibbs free energy approach is used to calculate the equilibrium constant for all reactions. The factors of 10^3 and 10^{-3} in Equation 2.17 are used to convert the rate-coefficient units, which are typically specified in $\mathrm{cm}^3{\cdot}\mathrm{mol}^{-1}{\cdot}\mathrm{s}^{-1}$, into SI units.

2.3.3 Transport Processes and Properties

The shock waves produced during hypersonic flight produce large gradients of mass, momentum, and energy in the flowfield. In general, transport processes serve to smooth out these gradients through the transfer of either mass, momentum, or energy throughout the mixture. Specifically, these processes are accomplished by diffusion, viscosity, and heat flux, respectively. This section provides explicit descriptions for the models used to describe each of these phenomena, as the uncertainty quantification study performed Chapter 3 considers uncertainty in the parameters underpinning the calculation of each transport coefficient.

Diffusion in LeMANS is modeled using a modified version of Fick's law

$$J_{s\neq e} = -\rho D_s \nabla Y_s + Y_s \sum_{r\neq e} \rho D_r \nabla Y_r \tag{2.20}$$

where D_s is the species diffusion coefficient, and $Y_s = \rho_s/\rho$ is the species mass fraction. Previous work in air [83] has shown that this model well approximates the full multicomponent diffusion solution given by the Stefan-Maxwell equations. Because the mass of an electron is significantly less than the heavy particles, the diffusion flux of electrons is treated using an ambipolar diffusion model

$$J_e = M_e \sum_{s\neq e} \frac{J_s C_s}{M_s} \tag{2.21}$$

where M_e is the electron molecular mass, and C_s is the charge of species s (i.e. $+1$ for ions and 0 for neutral species). The ambipolar assumption implies that the electrons move with the ions and is used to guarantee charge neutrality of the flowfield.

Momentum fluxes and shear stresses are modeled in LeMANS assuming a Newtonian fluid

and Stokes' hypothesis

$$\tau_{ij} = \mu\left(\frac{\partial u_j}{\partial x_i} + \frac{\partial u_i}{\partial x_j}\right) + \delta_{ij}\lambda\nabla\boldsymbol{u} \tag{2.22}$$

$$\lambda = -\frac{2}{3}\mu \tag{2.23}$$

where μ is the mixture coefficient of viscosity and δ_{ij} is the Kronecker delta.

Heat fluxes for both the translational-rotational and vibrational-electronic-electron modes are computed using Fourier's law

$$q_{tr} = -\kappa_{tr}\nabla T_{tr} \tag{2.24}$$

$$q_{ve} = -\kappa_{ve}\nabla T_{ve} \tag{2.25}$$

where κ_{tr} and κ_{ve} are the mixture thermal conductivity for the translational-rotational and vibrational-electronic-electron modes, respectively.

Each of the aforementioned transport processes requires knowledge of specific mixture transport properties. These are the diffusion coefficient, viscosity, thermal conductivity for diffusion, momentum flux, and heat flux, respectively. A common model for the mixture transport properties is Wilke's semi-empirical mixing rule [84]. However, it has been shown that the Wilke-based model underpredicts the heat transfer when temperatures in the shock layer exceed 10,000 K [83], which occurs frequently for the Ice Giant aerocapture environments studied in this work. Instead, the transport properties are computed using Gupta's mixing rule [85], which is considered to be more accurate than the Wilke model because it accounts for collision integrals in the multicomponent nature. For the two-temperature model assumed by LeMANS, the collision integrals are evaluated at an appropriate controlling temperature. For heavy-heavy collisions (e.g. H-H), the controlling temperature is taken to be T_{tr}, and for collisions involving electrons, the controlling temperature is taken to be T_{ve}.

Using Gupta's mixing rule, the mixture viscosity is calculated by

$$\mu = \sum_{s\neq e} \frac{m_s\gamma_s}{\sum\limits_{r\neq e}\gamma_r\Delta^{(2)}_{s,r}(T_{tr}) + \gamma_e\Delta^{(2)}_{s,e}(T_{ve})} + \frac{m_e\gamma_e}{\sum\limits_{r}\gamma_r\Delta^{(2)}_{e,r}(T_{ve})} \tag{2.26}$$

where the collision terms $\Delta_{s,r}$ are evaluated at the controlling temperatures. The molar concentration of each species, γ_s, is calculated by

$$\gamma_s = \frac{\rho_s}{\rho M_s} \quad . \tag{2.27}$$

The mass of each species, m_s, is given by

$$m_s = \frac{M_s}{N_{\text{av}}} \tag{2.28}$$

where $N_{\text{av}} = 6.022 \times 10^{-23}$ mol^{-1} is Avogadro's number.

The mixture thermal conductivity for the translational mode is calculated by

$$\kappa_t = \frac{15}{4} k_B \sum_{s \neq e} \frac{\gamma_s}{\sum\limits_{r \neq e} a_{s,r} \gamma_r \Delta^{(2)}_{s,r}(T_{tr}) + 3.54 \gamma_e \Delta^{(2)}_{s,e}(T_{ve})} \tag{2.29}$$

where

$$a_{s,r} = 1 + \frac{\left[1 - (m_s/m_r)\right]\left[0.45 - 2.54\,(m_s/m_r)\right]}{\left[1 + (m_s/m_r)\right]^2} \quad . \tag{2.30}$$

The thermal conductivity for the rotational mode is given by

$$\kappa_r = k_B \sum_{s=\text{molecules}} \frac{\gamma_s}{\sum\limits_{r \neq e} \gamma_r \Delta^{(1)}_{s,r}(T_{tr}) + \gamma_e \Delta^{(1)}_{s,e}(T_{ve})} \tag{2.31}$$

where it is emphasized that the summation is over molecules. The mixture thermal conductivity for the translational-rotational mode is then the sum of Equations 2.29 and 2.31. The thermal conductivity for the vibrational-electronic mode is calculated by

$$\kappa_{vel} = k_B \frac{C_{v_{ve}}}{R_u} \sum_{s=\text{molecules}} \frac{\gamma_s}{\sum\limits_{r \neq e} \gamma_r \Delta^{(1)}_{s,r}(T_{tr}) + \gamma_e \Delta^{(1)}_{s,e}(T_{ve})} \tag{2.32}$$

where $C_{v_{ve}}$ is the specific heat at constant volume for the vibrational-electronic mode, and $R_u = 8.314$ J/mol/K is the universal gas constant. The thermal conductivity for electrons is

$$\kappa_e = \frac{15}{4} k_B \frac{\gamma_e}{\sum\limits_{r} 1.45 \gamma_r \Delta^{(2)}_{e,r}(T_{ve})} \quad . \tag{2.33}$$

The collision terms, $\Delta^{(1)}_{s,r}$ and $\Delta^{(2)}_{s,r}$, in the previous equations are given by

$$\Delta^{(1)}_{s,r}(T) \;=\; \frac{8}{3}\left[\frac{2 M_s M_r}{\pi R_u T\,(M_s + M_r)}\right]^{1/2} 10^{-20} \pi \Omega^{(1,1)}_{s,r}(T) \tag{2.34}$$

$$\Delta^{(2)}_{s,r}(T) = \frac{16}{5}\left[\frac{2M_sM_r}{\pi R_u T\left(M_s+M_r\right)}\right]^{1/2} 10^{-20}\pi\Omega^{(2,2)}_{s,r}(T) \tag{2.35}$$

where the constant 10^{-20} converts Angstroms, the standard unit for collision integrals, to square meters. In Equations 2.34 and 2.35, the collision integral terms $\pi\Omega^{(l,j)}_{s,r}$ must be calculated according to the type of collision (i.e. neutral-neutral, neutral-electron, neutral-ion, etc.). For collisions involving neutral species, the collision integrals are expressed using a curve fit of the form

$$\pi\Omega^{(l,j)}_{s,r} = DT^{A(\ln T)^2+B\ln T+C} \tag{2.36}$$

where A, B, C, and D are interaction-dependent constants that are obtained from experiments and/or high-fidelity simulations. For interactions including hydrogen and helium, values for these curve-fit coefficients are taken from Palmer et al. [86], and the data from Park et al. [87] are used for all other neutral-neutral and neutral-ion interactions. Interactions between two charged species are modeled using Coulumbic theory as described by Wright et al. [88].

Finally, the binary diffusion coefficients between heavy particles and electrons are calculated by

$$D_{s,r} = \frac{k_B T_{tr}}{p\Delta^{(1)}_{s,r}(T_{tr})} \tag{2.37}$$

$$D_{e,r} = \frac{k_B T_{ve}}{p\Delta^{(1)}_{e,r}(T_{ve})} \quad . \tag{2.38}$$

The species diffusion coefficient used in the conservation equations is given by

$$D_s = \frac{\gamma_t^2 M_s\left(1-M_s\gamma_s\right)}{\sum\limits_{r\neq s}\left(\gamma_s/D_{sr}\right)} \tag{2.39}$$

where

$$\gamma_t = \sum_s \gamma_s \quad . \tag{2.40}$$

2.3.4 Boundary Conditions

The surface boundary condition requires specification of the thermodynamic state and the chemical composition. For the thermodynamic state, this dissertation assumes a radiative-equilibrium

boundary condition in which the surface temperature T_w is related to the wall-incident heat flux q_{wall} via

$$q_{\text{wall}} = \varepsilon \sigma_{SB} T_w^4 \quad . \tag{2.41}$$

Here, ε is the surface emissivity, and $\sigma_{SB} = 5.670 \times 10^{-8}$ W/m^2/K^4 is the Stefan-Boltzmann constant. Furthermore, the chemical composition boundary condition assumes the surface is fully catalytic to ions and electrons. Catalytic recombination occurs when dissociated/ionized species recombine at the surface. This process releases energy and increases the surface heating, which is a more conservative assumption than a non-catalytic surface. Accurate surface-heating predictions require that the near-surface gradients be sufficiently resolved. A quantitative description of the boundary-layer resolution is made by defining the nondimensional lengthscale

$$y^+ = \frac{u_\tau \Delta y \rho}{\mu} \tag{2.42}$$

where $u_\tau = \sqrt{\tau_w / \rho}$ is the friction velocity, and Δy is the height of the first cell normal to the surface. Generally, a value of $y^+ < 1$ is desired, which ensures that the first computational cell lies within the viscous sublayer of the boundary layer.

2.4 Radiative Transport Modeling

While LeMANS is capable of providing high-fidelity predictions of the convective heating and aerodynamic coefficients of an aerocapture vehicle, it cannot calculate the radiative-heating environment on its own. Radiative-heating calculations require knowledge of excited-state population distributions for each chemical species, and modeling each excited state as a separate species in LeMANS is computationally intractable for vehicle-scale simulations. Instead, the calculation of the state-population distributions and resulting radiative heating is performed by the NASA Ames tool NEQAIR. As an input, NEQAIR takes the number densities and temperatures predicted by LeMANS along a given line of sight. This information is then used to calculate the populations of excited species and spectral radiance produced at wavelength regions of interest using a line-by-line radiation solver. The net wall-directed radiative heating is obtained by integrating the spectral ra-

diance using a tangent-slab approximation, which assumes that the properties along a line of sight extend infinitely parallel to the line. Using a tangent-slab approximation reduces the necessary amount of calculations and is a reasonable assumption for predicting the radiative heating along the forebody of a blunt entry vehicle [89].

NEQAIR offers multiple methods for computing excited state populations. Under strong thermochemical nonequilibrium, assuming a Boltzmann distribution for all species can lead to inaccurate radiation predictions, as excited states may not have had sufficient time to relax. A more accurate approach is the quasi-steady-state (QSS) model, which determines electronic state populations by balancing excitation and quenching rates to achieve a net-zero time rate of change for each state. While incorporating the finite rates of excitation and quenching improves the fidelity of state-population calculations, it increases the computational complexity and requires reliable excitation and quenching rate data at relevant conditions.

For the H_2/He system, a combination of Boltzmann and QSS models [90] are used to compute species state populations. Helium state populations are calculated with an assumed Boltzmann distribution, while the hydrogen molecule and atom are calculated using QSS methods. As described in [90], helium is an inert substituent, so assuming a Boltzmann state distribution for He does not affect the radiative heating calculation. Atomic carbon is modeled using a QSS assumption, but the remaining carbon-based species state populations are modeled with a Boltzmann assumption. For additional details on the hydrogen QSS models, the reader is referred to Ref. [90].

2.5 Uncertainty Quantification and Sensitivity Analysis Methodology

This section details the methods utilized to perform UQ and GSA. Specifically, it describes the mathematical theory necessary to answer the first two questions described in the introduction regarding the predictions of the computational models. These are: 1) What is the uncertainty in each model's predictions? 2) Which input parameters are contributing most to that uncertainty? A polynomial chaos expansion UQ technique is used to answer the first question, and Sobol' indices are used to answer the second.

Before getting into the theory behind each of these techniques, it is helpful to define some terminology that is central to UQ and GSA. Generally, there are two types of uncertainty: aleatory and epistemic [91]. Aleatory uncertainty comes from the random variation in physical data used as an input to the computational model, and it is often referred to as "irreducible" because its variability cannot be controlled. Typically, aleatory uncertainty is modeled using a probability distribution function (e.g. Gaussian). Epistemic uncertainty comes from the lack of knowledge with regard to a specific model parameter and is typically represented using intervals. Epistemic uncertainty is also referred to as "reducible" uncertainty, as it may be decreased by collecting more data. This research treats all uncertainty as epistemic, as additional data for each source could reduce the corresponding uncertainty interval.

In the context of sensitivity analysis, a distinction must be made between *local* and *global* sensitivity analysis. Both are methods used to assess how variations in input parameters affect a model's output. Local sensitivity analysis examines how small changes in a single parameter affect the model while keeping others fixed, often using derivatives. It is computationally efficient and useful for identifying trends but is generally valid only within a narrow range and does not account for parameter interactions. Global sensitivity analysis, on the other hand, evaluates the influence of input variations across the entire parameter space, using statistical methods like variance decomposition (e.g., Sobol' indices) or Monte Carlo sampling. This approach accounts for nonlinear effects and parameter interactions, providing a more comprehensive understanding of system behavior. While computationally more expensive, it is particularly useful for complex, nonlinear models, such as those used to study aerocapture systems, where parameter interactions and uncertainties play a significant role. As such, unless stated otherwise, any usage of the term "sensitivity analysis" throughout this dissertation implies a *global* sensitivity analysis.

2.5.1 Polynomial Chaos Expansion

Recently, the use of polynomial chaos expansion (PCE) for large-scale UQ and global sensitivity analysis has become quite popular in the aerospace field [92, 93, 94, 95, 42, 44, 96, 46, 47].

PCE is a surrogate modeling technique based on a spectral representation of the uncertainty. A surrogate-based approach is attractive for UQ problems in which the deterministic model is expensive to evaluate, as statistical information can be computed directly from the surrogate. Hundreds of thousands of samples from the uncertain parameter space can be efficiently propagated through the fit PCE, bypassing the need to evaluate the deterministic model a statistically-significant number of times. Here, the "deterministic model" refers to a single LeMANS or NEQAIR simulation.

The PCE surrogate is used to approximate the functional relationship between random parameters and a given quantity of interest (QoI) by decomposing the QoI response function R into a linear combination of deterministic and stochastic components

$$R = \sum_{i=0}^{Q} \alpha_i \Psi_i(\boldsymbol{\xi}) \quad . \tag{2.43}$$

Here, α_i and Ψ_i are the deterministic component and random variable basis function, respectively, of the i-th mode of the response R, and $\boldsymbol{\xi}$ is a vector of normalized random (uncertain) parameters. The choice of the basis functions Ψ_i depends on the distribution of each random parameter and may be determined using the Ashey key [97]. Because this work treats all sources of uncertainty as epistemic, Legendre polynomials are used as the polynomial basis functions, as they have a bounded domain over the interval $[-1, 1]$. By definition, Equation 2.43 is an infinite series ($Q = \infty$), however, in practice, it is typical to truncate the sum over a discrete set of output modes. The number of terms required to form a complete expansion in Equation 2.43 is a function of number of the uncertain parameters N_ξ, and the order of the polynomial expansion p

$$N_t = Q + 1 = \frac{(N_\xi + p)!}{N_\xi! p!} \quad . \tag{2.44}$$

The form of Equation 2.44 highlights the "curse of dimensionality" commonly associated with PCE, as the number of terms required grows increasingly large for problems with many uncertain parameters. As shown in Figure 2.4, the number of terms for a total-order expansion quickly becomes computationally prohibitive for high-dimensional problems, as is the case in the present work.

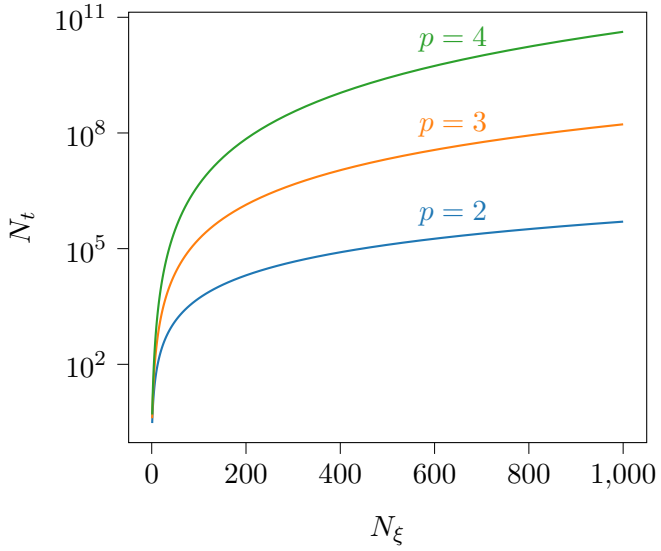

Figure 2.4: Number of terms required for total-order PCE.

2.5.2 Sparse Point-Collocation PCE

Many methods exist for obtaining the expansion coefficients α_i in Equation 2.43. Generally, they may be categorized into one of two approaches: intrusive and non-intrusive. Intrusive methods are generally more complex to implement, as they may require significant modification to the underlying computational model. On the other hand, non-intrusive polynomial chaos (NIPC) methods only require the functional inputs and corresponding outputs to construct a surrogate model capable of approximating the complex computational model (i.e. they treat the computational model as a black box). Due to its success in many aerothermodynamic UQ studies [92, 93, 94, 95, 42, 44], this work uses a non-intrusive point-collocation technique to obtain the unknown basis coefficients α_i by solving the following linear system given N_s evaluations of the computational model

$$\begin{pmatrix} R(\boldsymbol{\xi}_1) \\ R(\boldsymbol{\xi}_2) \\ \vdots \\ R(\boldsymbol{\xi}_{N_s}) \end{pmatrix} = \begin{pmatrix} \Psi_0(\boldsymbol{\xi}_1) & \Psi_1(\boldsymbol{\xi}_1) & \cdots & \Psi_Q(\boldsymbol{\xi}_1) \\ \Psi_0(\boldsymbol{\xi}_2) & \Psi_1(\boldsymbol{\xi}_2) & \cdots & \Psi_Q(\boldsymbol{\xi}_2) \\ \vdots & \vdots & \ddots & \vdots \\ \Psi_0(\boldsymbol{\xi}_{N_s}) & \Psi_1(\boldsymbol{\xi}_{N_s}) & \cdots & \Psi_Q(\boldsymbol{\xi}_{N_s}) \end{pmatrix} \begin{pmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_Q \end{pmatrix} . \tag{2.45}$$

Here, $R(\boldsymbol{\xi}_i)$ is the deterministic model output for random sample $\boldsymbol{\xi}_i$, where $i = 1, \ldots, N_s$. Note that, to obtain an analytical solution to Equation 2.45, N_t evaluations of the deterministic model are required. As highlighted in the previous section, the minimum number of terms can become substantially large for high-dimensional problems, such that obtaining an analytical solution for the expansion coefficients is computationally prohibitive. If fewer than N_t samples are available, the linear system is underdetermined, and an infinite number of solutions exist.

However, in high-dimensional parametric UQ problems, it is unlikely that all parameters contribute substantially to the total uncertainty in the model response. In this case, many of the coefficients in the polynomial expansion will be equal or close to zero, and the solution vector $\boldsymbol{\alpha}$ will be sparse. For underdetermined systems ($N_s < N_t$), techniques such as L^1-regularization [98] may be used to encourage the solution vector to be sparse. Examples of the successful use of this approach for aerothermodynamic UQ problems may be found in Refs. [93] and [44]. However, as motivated by West and Johnston [42], this work utilizes least-angle regression (LAR), as it does not require the solution to an optimization problem, reduces computational complexity, and has been shown to give similar results to optimization-based techniques like L^1-regularization. Introduced by Efron et al. [99], LAR is an iterative regression algorithm for solving underdetermined linear systems in which the regressors are moved from a *candidate* set to an *active* set based on their correlation with the solution residual. For more details on the LAR algorithm implementation in an NIPC-PCE setting, the reader is referred to Ref. [42].

2.5.3 Propagation of Uncertainty through PCE

Once the expansion coefficients α_i have been computed using the N_s evaluations of the computational model, the PCE surrogate may be used to efficiently obtain statistical information about the select QoI. To do so, a large number (usually on the order of 10^5–10^6) of random samples $\boldsymbol{\xi}_j$ are propagated through the surrogate to obtain the corresponding approximate function response $\hat{R}(\boldsymbol{\xi})$

$$\hat{R}_j(\boldsymbol{\xi}_j) = \sum_{i=0}^{Q} \alpha_i \Psi_i \left(\boldsymbol{\xi}_j\right) \quad . \tag{2.46}$$

The set of function-response values obtained is statistically significant enough to then compute quantities such as 95% confidence bounds on the QoI. In addition to percentiles, the use of a PCE surrogate allows for the first two statistical moments, the mean and variance, to be computed analytically from the expansion coefficients α_i

$$\text{mean}\left(\hat{R}\right) = \hat{\mu} = \alpha_0 \tag{2.47}$$

$$\text{var}\left(\hat{R}\right) = \hat{D}_{\text{tot}} = \sum_{i=1}^{Q} \alpha_i^2 \mathbb{E}\left[\Psi_i^2(\boldsymbol{\xi})\right] \quad . \tag{2.48}$$

For the Legendre polynomials used in this work, the variance reduces to

$$\hat{D}_{\text{tot}} = \sum_{i=1}^{Q} \alpha_i^2 \quad . \tag{2.49}$$

2.5.4 Sobol' Decomposition

The use of a PCE-based UQ technique lends itself naturally to a GSA methodology that uses Sobol' indices to allocate portions of the total QoI variance to individual parameters. Sobol' indices are derived using a variance-based sensitivity analysis technique that leverages a Sobol' decomposition of the total variance D_{tot}

$$D_{\text{tot}} = \sum_{i=1}^{N_\xi} D_i + \sum_{1 \le i < j \le N_\xi}^{i=N_\xi-1} D_{i,j} + \sum_{1 \le i < j < k \le N_\xi}^{i=N_\xi-2} D_{i,j,k} + \ldots + D_{1,2,\ldots,N_\xi} \tag{2.50}$$

which is a sum over partial variances due to variations in the N_ξ unique input parameters. Partial variances with a single subscript, D_i, account for variance that is solely due to variations in the i-th parameter and are commonly referred to as "first-order" variances. Partial variances that have more than one subscript account for variance due to interactions between two or more uncertain parameters and are typically referred to as "higher-order" variances. Applying the Sobol' decomposition to the variance obtained from the PCE surrogate yields

$$\hat{D}_{\text{tot}} = \sum_{i=1}^{N_\xi} \hat{D}_i + \sum_{1 \le i < j \le N_\xi}^{i=N_\xi-1} \hat{D}_{i,j} + \sum_{1 \le i < j < k \le N_\xi}^{i=N_\xi-2} \hat{D}_{i,j,k} + \ldots + \hat{D}_{1,2,\ldots,N_\xi} \tag{2.51}$$

where the partial variances are expressed as a function of the expansion coefficients

$$\hat{D}_{i_1,\ldots,i_m} = \sum_{\beta \in \{i_1,\ldots,i_m\}} \alpha_\beta^2 \mathbb{E}\left[\Psi_\beta^2(\boldsymbol{\xi})\right], \quad 1 \le i_1 < \ldots < i_m \le N_\xi \quad . \tag{2.52}$$

An m-th order Sobol' index is then explicitly defined as

$$S_{i_1,\ldots,i_m} = \frac{\hat{D}_{i_1,\ldots,i_m}}{\hat{D}_{\text{tot}}} \tag{2.53}$$

which gives the relative contribution of the uncertain parameters $(i_1, \ldots, i_s)$ to the total variance and satisfies the equation

$$1 = \sum_{i=1}^{N_\xi} S_i + \sum_{1\le i<j\le N_\xi}^{i=N_\xi-1} S_{i,j} + \sum_{1\le i<j<k\le N_\xi}^{i=N_\xi-2} S_{i,j,k} + \ldots + S_{1,2,\ldots,N_\xi} \quad . \tag{2.54}$$

First-order Sobol' indices, or those with a single subscript, are interpreted as the percentage of the total variance that may be attributed to variations in a single uncertain parameter. Similarly, higher-order Sobol' indices represent the amount of uncertainty that is allocated to interactions between uncertain parameters. It is also possible to define a total Sobol' index, which is the sum of all Sobol' indices that include a given parameter i

$$S_{T_i} = \sum_{\mathcal{L}_i} \frac{\hat{D}_{i_1,\ldots,i_m}}{\hat{D}_{\text{tot}}}, \quad \mathcal{L}_i = \{(i_1, \ldots, i_m) : \exists\ k, 1 \le k \le m, i_k = i\} \quad . \tag{2.55}$$

For example, if $N_\xi = 3$, the total Sobol' index for the first uncertain parameter would be

$$S_{T_1} = S_1 + S_{1,2} + S_{1,3} + S_{1,2,3} \quad . \tag{2.56}$$

2.5.5 Monitoring Convergence

Predictions from the PCE surrogate are a direct function of the expansion coefficients. Because the sparse approximation used to obtain the expansion coefficients implicitly assumes $N_s < N_t$, it is necessary to ensure that expansion coefficients have converged. Instead of tracking changes in the expansion coefficients themselves, a common technique for assessing the convergence of the PCE is to use the change in Sobol' indices from one set of sample points to the next. In general, there are far more expansion coefficients than Sobol' indices, so using Sobol' indices to track convergence requires fewer values to be tracked. In particular, the total Sobol' index is well-suited for monitoring the PCE convergence, as it captures the sensitivity contributions from individual parameter variation as well as mixed contributions caused by parameter interactions. This research computes both

the maximum change in total Sobol' index, and the mean change in total Sobol' index across all parameters as the expansion sample size grows. Specifically, for the j-th total Sobol' index and sample size $n > 1$, these are expressed as

$$\Delta S_{T,\max,n} = \max\left\{|S_{T_j,n} - S_{T_j,n-1}| : j = 1,\ldots,N_\xi\right\} \tag{2.57}$$

$$\Delta S_{T,\text{mean},n} = \text{mean}\left\{|S_{T_j,n} - S_{T_j,n-1}| : j = 1,\ldots,N_\xi\right\} \quad . \tag{2.58}$$

Equations 2.57 and 2.58 provide simple metrics for the change across all Sobol' indices as more samples are used to create the polynomial expansion, and a maximum change of 1–2% over several consecutive iterations is considered an acceptable tolerance for convergence.

2.5.6 Implementation Details

Here, it is worth mentioning how the previously described UQ/GSA theory is applied in practice. Generally, one may assume that the PCE surrogate will be inaccurate for very small values of N_s. Therefore, instead of fitting the PCE and checking the total Sobol' index convergence after each evaluation of the computational model, batches of simulations of size N_b are run and the PCE is fit after each batch has completed. In this work, a batch size of $N_b = 100$ is used. For each batch, N_b samples of the uncertain parameters are generated from a joint-marginal distribution using a space-filling Latin-hypercube sampling algorithm, where each marginal distribution corresponds to one of the input parameters. Once the input space has been sampled, the UQ and GSA is performed as follows (summarized visually in Figure 2.5):

(1) Evaluate LeMANS and NEQAIR N_b times using the generated samples.

(2) Compute N_d values of the desired QoIs from each LeMANS and NEQAIR solution.

(3) Solve Equation 2.45 for the expansion coefficients using the LAR algorithm.

(4) Compute total Sobol' indices using Equation 2.55.

(5) Propagate uncertainty through the fit PCE surrogate and compute statistics of response.

(6) Evaluate convergence metrics using Equations 2.57 and 2.58.

(7) Generate an additional N_b samples and repeat steps 1-6 if convergence criteria is not met.

The sampling of the input space, construction of the orthogonal basis polynomials, LAR algorithm, and calculation of the Sobol' indices are all implemented using the Matlab toolbox `UQLab` [100].

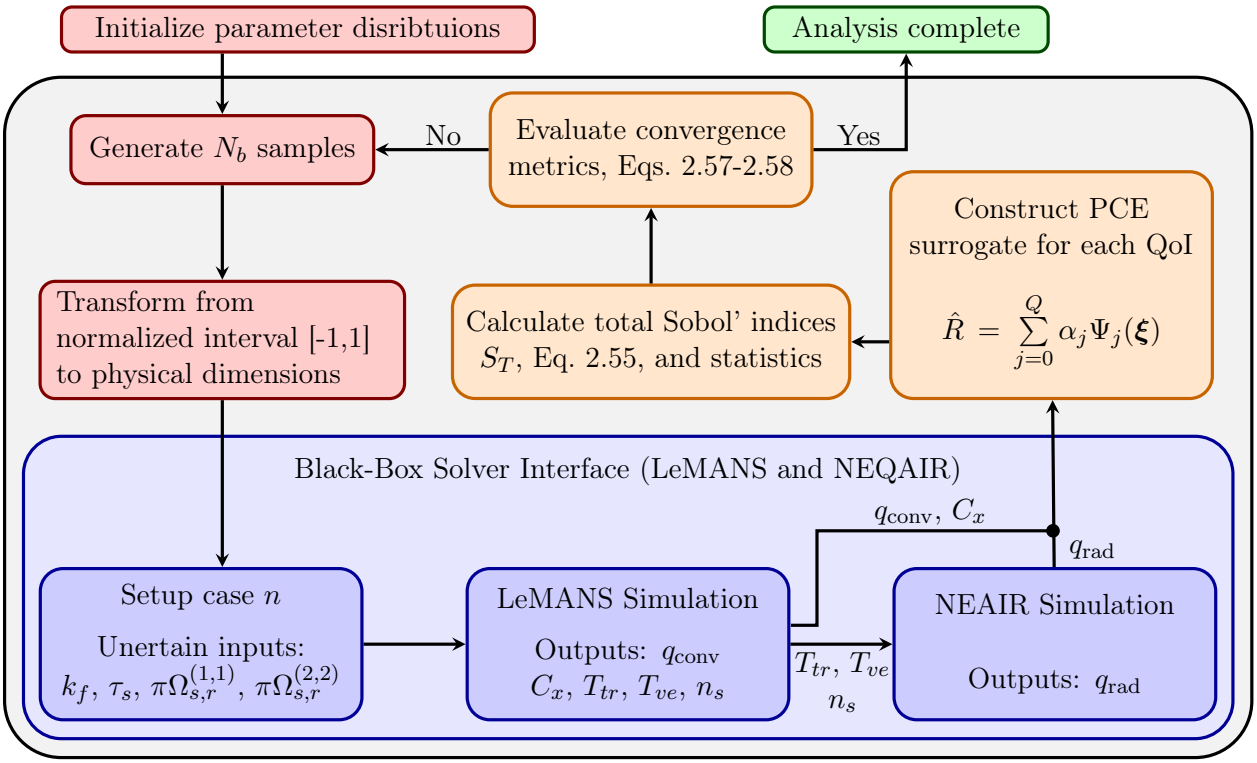

Figure 2.5: Flowchart of UQ/GSA framework used in this research.

2.6 Reduced-Order Modeling using Gaussian Process Regression

The third and final question related to the predictions from LeMANS and NEQAIR is: how can the cumulative effects of that uncertainty along a trajectory be approximated? Performing a detailed UQ study at each point along an aerocapture trajectory is computationally intractable, but capturing the cumulative effects of the aerothermodynamic uncertainty along the entire trajectory is necessary for robust thermal protection system (TPS) design. As such, it is desirable to develop

a reduced-order modeling framework that is capable of incorporating aerothermal uncertainties into its predictions. This section describes the fundamental concepts of Gaussian process regression (GPR), which is a powerful surrogate-modeling technique for uncertain systems.

2.6.1 Gaussian Processes

In a general supervised learning problem, the goal is to learn some unknown mapping $f : \mathbb{R}^D \mapsto \mathbb{R}$ from empirical input-output data. Typically, the user only has access to noisy outputs/measurements such that for an input (feature) $\boldsymbol{x} \in \mathbb{R}^D$, the available output data (target) $y \in \mathbb{R}$ is of the form

$$y(\boldsymbol{x}) = f(\boldsymbol{x}) + \epsilon \tag{2.59}$$

where $\epsilon \sim \mathcal{N}(0, \sigma_n^2)$ is modeled as independent and identically distributed (i.i.d.) Gaussian noise with zero mean and standard deviation σ_n. Numerous techniques exist to approximate the function f, but most are limited by assuming that the underlying function takes a parametric form defined by a finite set of parameters. GPR is considered a non-parametric regression method, or more intuitively, a method where the underlying functional form has an infinite number of parameters, as it takes a Bayesian approach to the general regression problem in Equation 2.59. A Gaussian process (GP) is a generalization of the Gaussian distribution and may be thought of as a distribution over all functions [101], where for a series of inputs in the set $\boldsymbol{X} = \left\{\boldsymbol{x}_i \in \mathbb{R}^D : \ i = 1, \ldots, N\right\}$, the corresponding outputs in the set $\boldsymbol{Y} = \{y_i \in \mathbb{R} : \ i = 1, \ldots, N\}$ are assumed to be jointly Gaussian distributed. Similar to a standard multivariate Gaussian distribution, a GP is completely defined by its mean function $\boldsymbol{\mu}(\boldsymbol{x}) : \mathbb{R}^D \mapsto \mathbb{R}$ and covariance matrix $K \in \mathbb{S}_+^{N \times N}$ where the i,j^{th} element of the covariance matrix is defined by a kernel function $k(\cdot, \cdot)$ such that $K_{ij} = k(\boldsymbol{x}_i, \boldsymbol{x}_j)$. A function represented by a GP is denoted by

$$f(\boldsymbol{x}) \sim \mathcal{GP}\left(\mu(\boldsymbol{x}), k(\boldsymbol{x}, \boldsymbol{x}')\right) \qquad . \tag{2.60}$$

2.6.2 Kernel Functions

The form of the covariance function $k(\cdot,\cdot)$ must produce valid covariance matrices, meaning that for any set of inputs, the mapping $k(\cdot,\cdot)$ produces a symmetric, positive definite matrix. Numerous functions exist for the choice of k, and the reader is referred to Ref. [102] for a more in-depth analysis of valid kernel functions. In this dissertation, the two primary kernel functions used are the squared exponential (SE) kernel

$$k(\boldsymbol{x},\boldsymbol{x}')_{\mathrm{SE}} = \sigma^2 \exp\left(-\frac{1}{2}r^2\right) \tag{2.61}$$

and the Matérn 5/2 kernel

$$k(\boldsymbol{x},\boldsymbol{x}')_{5/2} = \sigma^2\left(1+\sqrt{5}r+\frac{5}{3}r^2\right)\exp\left(-\sqrt{5}r\right) \tag{2.62}$$

where

$$r = \sqrt{\sum_{d=1}^{D}\frac{(x_d - x'_d)^2}{l_d^2}} \quad . \tag{2.63}$$

Here, σ^2 is the variance, and l_d is the length scale associated with dimension d. Note that it is not strictly necessary to assign a length scale for each input dimension d. One could just as easily use a single length scale for all input dimensions, however, this generally restricts the regressive capabilities of the GP. Particularly for scenarios in which the response function behaves differently in each input dimension, it is more robust to assign each input dimension its own length scale. This technique is known as automatic relevance determination (ARD), and, unless specified otherwise, is used for all GP models constructed in this dissertation.

The set of parameters $\boldsymbol{\theta} = \{\sigma,\ l_d :\ d = 1,\ldots,D\}$ is often referred to as the hyperparameter set for the GP, and it defines the structure of the covariance matrix K. Proper selection of the hyperparameters is central to GPR, as the covariance matrix K not only determines the shape of the GP distribution, but ultimately defines the characteristics of the functions sampled from the GP. Before discussing how to determine the best (optimal) values for the GP hyperparameters, it is useful to demonstrate the qualitative effects of the variance and length scale parameters. Generally, the variance σ^2 controls how far away from the mean a sampled function lies. A larger variance will

produce a "wider" distribution of functions, while a smaller variance produces sampled functions that are more clustered near the mean. On the other hand, the length scale l controls how quickly the correlation between data points decays with distance. Physically, a smaller length scale will produce sampled functions that are more "wiggly," while a larger length scale will produce more smoothly varying functions. To help illustrate these concepts, functions sampled from a GP prior using an SE kernel with varying hyperparameter values are visualized in Figure 2.6, where the shaded region denotes the 3σ uncertainty of the prior distribution.

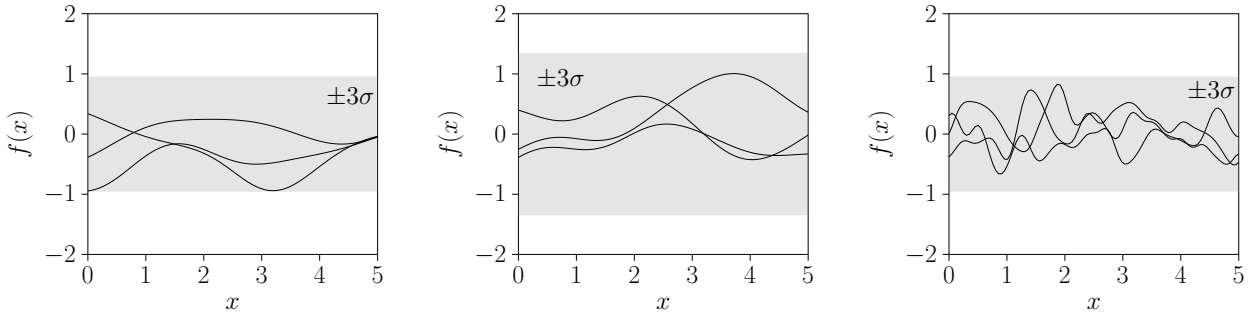

(a) $\sigma^2 = 0.1$, $l = 1.0$ (b) $\sigma^2 = 0.2$, $l = 1.0$ (c) $\sigma^2 = 0.1$, $l = 0.5$

Figure 2.6: Functions sampled from GP priors using squared-exponential kernel with various variances, σ^2, and length scales, l.

Selecting a kernel function, like the SE, is not a "once-size-fits-all" approach to constructing a GP surrogate, meaning that some kernel functions will perform better than others for a given regression problem. As an example, the SE kernel function works extremely well for a broad class of sufficiently smooth functions, but it tends to struggle if there are discontinuities in the first few derivatives of the underlying function f. It is for this reason that some GP models fit in this dissertation use an SE kernel, while others use the Matérn 5/2 kernel. For example, aerodynamic coefficients generally vary smoothly across the relevant input dimensions (density, velocity, angle of attach), while surface-heating profiles tend to have sharp peaks near the shoulder of the vehicle.

2.6.3 Posterior Predictive Distribution

In practice, one does not typically sample from the GP prior, as the sampled functions are in no way related to some underlying set of data/measurements. Instead, it is desirable to "infer" functions that are most likely to be represented by the available data, thereby allowing the modeler to predict the value of $f(\boldsymbol{x})$ at locations where data is not available. Inference on f is performed by conditioning the prior distribution in Equation 2.60 on training data, and the resulting conditional distribution is referred to as the posterior distribution. For a set of N training points $\boldsymbol{X} = \left\{\boldsymbol{x}_i \in \mathbb{R}^D : \ i = 1, \ldots, N\right\}$ and corresponding noisy outputs $\boldsymbol{Y} = \{y_i \in \mathbb{R} : \ i = 1, \ldots, N\}$, the GP prior placed on the function f means that the observations $\boldsymbol{Y}$ are, by definition, jointly Gaussian with the corresponding outputs $\boldsymbol{f}(\boldsymbol{x}) \in \mathbb{R}^N$. Without loss of generality, assume that the GP prior has zero mean. Then, at any set of test points $\boldsymbol{X}^*$

$$\begin{bmatrix} \boldsymbol{Y} \\ \boldsymbol{f}(\boldsymbol{X}^*) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \ \begin{bmatrix} K(\boldsymbol{X}, \boldsymbol{X}) + \sigma_n^2 I_N & K(\boldsymbol{X}, \boldsymbol{X}^*) \\ K(\boldsymbol{X}^*, \boldsymbol{X}) & K(\boldsymbol{X}^*, \boldsymbol{X}^*) \end{bmatrix} \right) \tag{2.64}$$

where the notation $K(\cdot, \cdot)$ is introduced to represent the matrix whose i, j^{th} element is given by $K_{ij} = k(\cdot, \cdot)$. Using the standard results for conditioning multivariate Gaussian distributions, the posterior predictive mapping $\hat{f}$ is also represented by a GP of the form

$$\hat{f}(\boldsymbol{X}^*) \sim p(f(\boldsymbol{X}^*)|\boldsymbol{X}^*, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{\theta}) = \mathcal{GP}(\hat{\mu}, \hat{k}) \tag{2.65}$$

where the posterior predictive mean $\hat{\mu}$ and covariance matrix $\hat{K}$ are given by

$$\hat{\mu}(\boldsymbol{X}^*) = K(\boldsymbol{X}^*, \boldsymbol{X}) \left[K(\boldsymbol{X}, \boldsymbol{X}) + \sigma_n^2 I\right]^{-1} \boldsymbol{Y} \tag{2.66}$$

$$\hat{K}(\boldsymbol{X}^*, \boldsymbol{X}^*) = K(\boldsymbol{X}^*, \boldsymbol{X}^*) - K(\boldsymbol{X}^*, \boldsymbol{X}) \left[K(\boldsymbol{X}, \boldsymbol{X}) + \sigma_n^2 I\right]^{-1} K(\boldsymbol{X}, \boldsymbol{X}^*) \quad . \tag{2.67}$$

Note that the predictive mean in Equation 2.66 is a linear combination of the noisy observations $\boldsymbol{Y}$, which is referred to as a *linear predictor* [101]. On the other hand, the predictive covariance in Equation 2.67 is only a function of the observation and test inputs, which is a known property of Gaussian distributions and is reflected in the form of the kernel function. The predictive covariance

is also implicitly a function of the kernel hyperparameters, which are not known *a priori*, but are rather learned during the training process.

The regression, or "training", aspect of GPR involves determination of the hyperparameters $\boldsymbol{\theta}$, and optionally the observation noise σ_n, of the GP (the observation noise is sometimes referred to as the "likelihood variance" in the GPR literature [101]). This is typically achieved through numerical optimization, in which the negative of the log marginal likelihood is minimized to determine the optimal set of hyperparameters $\boldsymbol{\theta}^*$

$$\boldsymbol{\theta}^* = \underset{\boldsymbol{\theta},\sigma_n}{\arg\min} \ \ln p(\boldsymbol{Y}|\boldsymbol{X},\boldsymbol{\theta},\sigma_n) \tag{2.68}$$

$$\ln p(\boldsymbol{Y}|\boldsymbol{X},\boldsymbol{\theta},\sigma_n) = -\frac{1}{2}\boldsymbol{Y}^T\left(K+\sigma_n^2 I\right)^{-1}\boldsymbol{Y} - \frac{1}{2}\ln\left|K+\sigma_n^2 I\right| - \frac{N}{2}\ln 2\pi \quad . \tag{2.69}$$

Here, the shorthand notation $K = K(\boldsymbol{X},\boldsymbol{X};\boldsymbol{\theta})$ is introduced for clarity. The minimization of Equation 2.69 may be achieved by a variety of numerical optimization algorithms, such as the Adam [103] optimizer.

2.6.4 Simple GPR Example

To help demonstrate the capabilities of GPR, a simple one-dimensional regression problem is used. Consider the scenario where the true function is $f(x) = \sin(x)$, however, only noisy measurements of the form

$$y(x) = \sin(x) + \epsilon \tag{2.70}$$

are available. The noise variance is assumed to be $\sigma_n^2 = 5.0 \times 10^{-2}$, and 15 noisy observations are available that are randomly sampled from the uniform distribution $\mathcal{U}(1,4)$ as shown in Figure 2.7.

A GP prior with a zero mean and squared-exponential covariance function is placed on the function $f(x)$, and the Adam optimizer [103] is used to optimize the kernel variance, kernel length scale, and the likelihood variance. The optimized hyperparameters are summarized in Table 2.1. The kernel parameters offer valuable insights into the regularity of the GP model function and the underlying data. The kernel variance parameter, σ^2, quantifies the uncertainty in the random

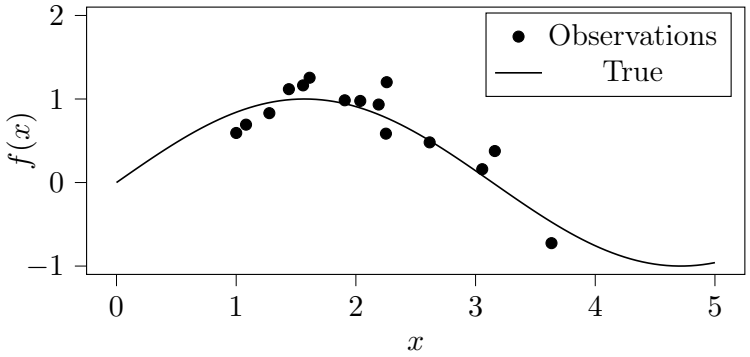

Figure 2.7: Noisy measurements and true function for simple GPR problem.

function predictions and directly reflects the spread of the data. This differs from observation noise, σ_n^2, which represents measurement accuracy, whereas variance captures the uncertainty of the underlying function given the noisy observations. The length-scale parameter, l, expressed in units of the input dimension, indicate the distance over which significant variations in the output variable occur. As seen in Table 2.1, the optimized length scale for this problem is on the order $\mathcal{O}(1)$. Intuitively, this means that the underlying function $f(x)$ should not vary significantly over input ranges smaller than this. For this simple problem, this implication is easily confirmed by inspection of Figure 2.7.

Table 2.1: Optimized hyperparameters for toy GPR problem.

Hyperparameter	Kernel variance σ^2	Kernel length scale l	Observation noise
Value	0.86	1.36	4.34×10^{-2}

Figure 2.8 illustrates the posterior mean and 3σ uncertainty, as well as three random samples drawn from the posterior distribution computed. In practice, the random samples are obtained by sampling a multivariate-normal distribution with the mean defined by Equation 2.66 and the covariance matrix defined by Equation 2.67. In general, the posterior mean shows excellent agreement with the true function $f(x)$, and the 3σ confidence interval contains the true function everywhere. Note that the posterior distribution can be sampled at any input location, including outside (i.e. extrapolating) the bounds of the training data. However, when sampled far enough away from the

training data, the uncertainty in the model's predictions grow, as the covariance function uses the distance between the training data and the point at which the posterior is being sampled. This characteristic is reflected in Figure 2.8, where the variance in the model's predictions is generally low near the training data and grows larger away from the training data set. Such behavior is a well-known characteristic of GPR and allows for the "out-of-distribution" (i.e. extrapolating instead of interpolating) uncertainties to be estimated. This feature is useful when the underlying function is not known a priori, as the increased posterior variance directly informs the user where additional training data is needed.

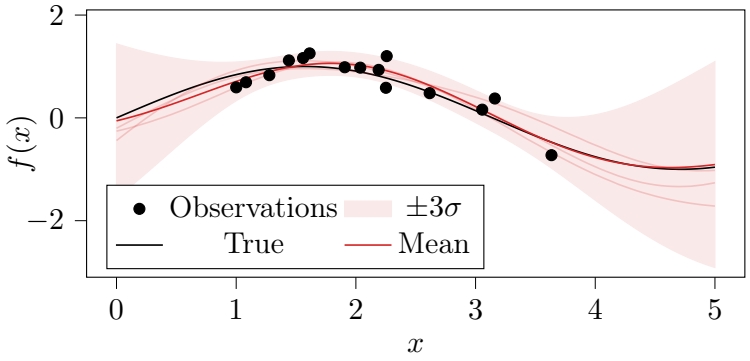

Figure 2.8: Simple 1D Gaussian process regression posterior distribution.

The material presented in this section covers the most basic usage of GPR for the development of probabilistic surrogate models. For most of the surrogates developed in this dissertation, however, more advanced features are required. More complex frameworks enable the use of multi-output, multi-fidelity, and constrained GPR, which are particularly useful for hypersonic applications because physical information about the underlying function can be used to improve the model's predictive capabilities. For example, it is well known that heat flux increases with freestream velocity, so it is beneficial to encode this constraint within the structure of the GP itself. The mathematical background and construction of GPs with these advanced features is discussed in detail in Chapter 4. All Gaussian process models in this work are constructed using the toolbox GPflow [104], a python toolbox built on top of TensorFlow 2 [105] for efficient and highly flexible GP design.

2.7 Summary

This chapter introduced the fundamental concepts of thermochemical nonequilibrium and the computational tools used to model it. First, the physical processes occurring behind strong shock waves were discussed, emphasizing their relevance to aerocapture vehicles. The distinction between thermochemical equilibrium and nonequilibrium was then established, along with key terminology for describing hypersonic flow states. Next, two high-fidelity numerical tools for modeling thermochemical nonequilibrium in aerocapture aerothermodynamics were presented. The first, the CFD code LeMANS, predicts aerodynamic forces and convective heating, while the second, NEQAIR, computes radiative heating.

Following the description of the computational models, the theory behind uncertainty quantification (UQ) and sensitivity analysis techniques used in this research was introduced. This included the mathematical foundations of sparse point-collocated polynomial chaos expansion for estimating uncertainty in a given quantity of interest, as well as Sobol' indices for sensitivity analysis. Finally, the fundamental principles of Gaussian process regression (GPR), a probabilistic reduced-order modeling technique, were explained, and an introductory example usage of GPR was given.

Chapter 3

Aerothermodynamic Analysis and Uncertainty Quantification

"When this baby hits 88 miles per hour, you're gonna see some serious stuff."
— Doc, Back to the Future

3.1 Introduction

At this point in the dissertation, the path to answering questions related to high-fidelity aerothermodynamic modeling of an Ice Giant aerocapture vehicle has been paved. In Chapter 1, the pitfalls of the typical, low-fidelity, aerothermodynamic models used in aerocapture trajectory simulations were emphasized, particularly their inability to account for aerothermodynamic uncertainty. Then, Chapter 2 laid the computational groundwork for improving upon these low-fidelity models by using high-fidelity CFD, UQ, and surrogate modeling. Now comes the time to apply the high-fidelity computational tools and UQ techniques described in Sections 2.3 and 2.5, respectively.

In this Chapter, a high-fidelity computational analysis of the aerothermodynamic environments of an Ice Giant aerocapture scenario is performed. First, a review of existing high-fidelity studies, both experimental and computational, on the aerothermodynamic environments in Ice Giant atmospheres is performed. Then, the specific thermochemical models used to simulate an H_2/He/CH_4 atmosphere in this research are described. Next, a reference aerocapture trajectory is generated, along which four key points are selected for analysis. The aerothermodynamic analysis is then conducted in two parts. In the first part, the nominal values for all input parameters are used to generate baseline profiles for three aerothermodynamic quantities of interest (QoIs). These

QoIs are namely the convective- and radiative-heating distributions, and the axial-force coefficient. Next, the UQ/GSA framework presented in Chapter 2.5 is used to quantify the uncertainty in each of the QoIs and identify the parameters contributing most to the uncertainty. Finally, the effects of including $H_2{}^+$ in the thermochemical model, a species typically omitted from Ice Giant aerocapture aerothermodynamic analyses, on the resulting QoIs are presented.

3.2 Background

Predicting the convective- and radiative-heating environments of an aerocapture vehicle requires highly sophisticated analysis tools, in which high-fidelity models for chemical kinetic rates, mixture and species transport properties, exchange of energy between translational and internal energy modes, and gas-surface interactions are used. The Ice Giants' atmospheres are primarily composed of hydrogen (H_2) and helium (He), and there have been several studies designed to obtain two-temperature compatible models for the chemical reaction rates in H_2/He mixtures. These include the works of Leibowitz et al. [26, 30], Park [29], and Boyd [106]. Owed to the large entry velocities, modeling the high-temperature, post-shock environments of an Ice Giant aerocapture scenario requires an accurate description of the kinetics of molecular hydrogen dissociation and atomic hydrogen ionization. However, each of the aforementioned models predict vastly different hydrogen ionization rates, and previous work from Hansson et al. [107] has shown that none are able to produce post-shock species distributions that agree well with experimental results. Recently developed state-to-state chemical kinetic models from Colonna et al. [108, 109, 110, 111] have demonstrated that the inclusion of otherwise omitted minority species, such as $H_3{}^+$ and H^-, can drastically alter the post-shock species distribution and improve agreement with experimental data. On the other hand, a series of similar state-to-state based studies from Carroll et al. [112, 40] have found that slight modifications to existing state-specific rate data produce an improved agreement between computational and experimental results without needing to model $H_3{}^+$ or H^-. Unfortunately, these state-specific models are not compatible with existing two-temperature solvers that would be used to model vehicle-scale geometries. However, Carroll et al. [112] determined that increasing the

existing ionization rates of Leibowitz et al. [30] by a factor of 5 produced reasonable agreement between two-temperature and state-to-state radiance predictions. Regardless, any uncertainty in the level of dissociation and ionization within the shock layer can produce substantial variation in the aerothermodynamic environment, particularly the radiative heating.

In a past experimental shock-tube campaign, Cruden et al. [113] sought to characterize the radiative heating contribution to the aerothermodynamic environment of H_2/He mixtures. They observed significant levels of nonequilibrium and radiation for shock speeds in excess of 25 km/s; however, for shock speeds below 25 km/s, little radiation was observed. In addition to H_2 and He, the Ice Giants also contain trace amounts of methane (CH_4). While typically omitted from Ice Giant aerothermodynamic analyses, there is evidence that the inclusion of CH_4 can substantially alter the radiative-heating environment. A recent computational study from Coelho et al. [35] demonstrated that introducing a concentration of 1.5% CH_4 by mole resulted in an aerothermodynamic environment in which 50% of the heating was attributed to radiation. However, the hydrogen ionization rates used in [35] are reported to be two orders of magnitude lower than the actual values presented by Leibowitz et al. [30]. As observed by Carroll et al. [112], the Leibowitz rates already underpredict the radiance produced at conditions relevant to Ice Giant aerocapture, so the slower ionization modeled by Coelho et al. [35] is likely underpredicting the proportion of radiation produced by hydrogen. Furthermore, Coelho et al. concluded that considering CH_4 in the freestream produces significant radiation for shock speeds as low as 19 km/s. Recent experimental efforts from Steer et al. [114] observed that mixtures of 0.845/0.15/0.005 and 0.8/0.15/0.05 of H_2/He/CH_4 (by mole) produced radiation signatures that were not present in pure H_2/He mixtures for shock speeds up to 18.9 km/s. Cruden and Tibere-Inglesse [37] also found that the addition of 0.4% by mole CH_4 produced radiative features that were not present in H_2/He-only mixtures for shock speeds up to 26 km/s, however, they were smaller than those observed by Steer et al. [114] (owed to the increased CH_4 and He concentrations in [114]) and were far outweighed by hydrogen-based features. The actual amount of CH_4 in the atmospheres of Neptune and Uranus is highly uncertain, such that the amount of CH_4 in the troposphere can range from 1-5% by mole [115]. While the majority of

an Ice Giant aerocapture trajectory remains within the stratosphere (where less CH_4 is expected), the potentially large sensitivity of radiative heating to the amount of freestream CH_4, and the high level of uncertainty in the amount of freestream CH_4 motivates further quantification of the aerothermodynamic uncertainty for H_2/He/CH_4 mixtures.

To date, there has only been one major effort to quantify the aerothermodynamic uncertainty in H_2/He mixtures at conditions representative of an Ice Giant aerocapture mission. Palmer et al. [86] quantified the uncertainty in convective and radiative stagnation-point heating rates for a 1-meter aeroshell entering the atmospheres of Saturn and Uranus at velocities of 28.2 km/s and 22.3 km/s, respectively. Their study performed a large Monte Carlo analysis consisting of thousands of CFD simulations in which chemical reaction-rate coefficients, transport properties, and vibrational relaxation times were assumed to be uncertain. An extensive review of the transport properties in H_2He mixtures was performed in [86] which was used to inform the uncertainty intervals on collision-integral parameters. On the other hand, no review of the chemical reaction-rate coefficients was made. Rather, the nominal rate coefficients were taken from Leibowitz et al. [26, 30], and a Gaussian distribution with a 3σ uncertainty of an order of magnitude was imposed. This gives no consideration to other models [29, 106] for the same reactions and can result in overly-conservative aeroheating uncertainty estimates. Palmer et al. [86] concluded that, for catalytic surfaces, the convective heating 2σ-uncertainty was only a few percent, but as high as 18% for non-catalytic surfaces. For the highest-velocity Saturn entry case, significant radiation was observed that nominally accounted for ∼10-20% of the total heating. Additionally, the radiative-heating uncertainty was significantly larger with 2σ-uncertainties of up to 280%. No radiative heating results were presented for the Uranus scenario due to the lower entry velocity. Furthermore, Palmer et al. [86] did not model CH_4 as a freestream species and only considered the stagnation-point heating. In reality, an aerocapture mission at the Ice Giants could have entry velocities in excess of 30 km/s [13, 116], and the surface heating at the vehicle's shoulder can be as high as at the stagnation-point. More recently, a study from Erb et al. [117] assessed the impact of modeling practices like coupled radiation and ablation on surface heating uncertainties, but their work was directed at Jovian entry

and not Ice Giant aerocapture. Due to the immense entry velocity for Jovian entry (47.4 km/s), the aerothermodynamic environment is dominated by radiation, so the same conclusions of Erb et al. [117] cannot necessarily be made for an Ice Giant aerocapture environment.

In light of these existing works and the increased attention toward Ice Giant aerocapture, there is a need for an updated and more comprehensive quantification of the aerothermodynamic uncertainty in $H_2/He/CH_4$ mixtures. Specifically, the current work builds upon previous analysis [45] and performs a sensitivity analysis and uncertainty quantification (UQ) study of the aerothermodynamic environment of a Neptune aerocapture scenario, where the imposed uncertainty intervals are informed by existing data, when possible. Furthermore, consideration is given to the uncertainty in the heating *distributions*, not just the stagnation point, and the freestream composition is taken to be uncertain in order to reflect the lack of knowledge in the Ice Giant atmosphere composition. A parametric analysis, considering uncertainty in reaction-rate coefficients, collision integral fitting parameters, and vibrational-relaxation time fitting parameters is performed to determine the uncertainty in the axial-force coefficient, and convective- and radiative-heating distributions. Finally, a global sensitivity analysis is performed in which Sobol' indices are used to identify the processes contributing most to the uncertainty in the selected aerothermodynamic quantities of interest. Because an Ice Giant aerocapture mission has yet to be flown, it is necessary to characterize the aerothermodynamic environment across a broad range of densities and velocities. As such, an aerocapture mission to Neptune is chosen over Uranus because of the nominally larger entry velocity[1] . Neptune and Uranus are believed to have similar atmospheres, with Uranus having slightly less methane than Neptune. However, the results in this Chapter demonstrate a minor level of sensitivity to the freestream methane concentration, so the differences in atmospheric compositions between Uranus and Neptune will not produce drastically different results. Therefore, the findings in this work are also applicable to a Uranus aerocapture scenario.

[1] Some Uranus aerocapture concepts have similarly high entry velocities [116]

3.3 Computational Modeling Approach and Sources of Uncertainty

This section provides an overview of the computational models used to simulate the aerothermodynamic environment of an Ice Giant aerocapture mission. Section 3.3.1 describes the setup of the computational fluid dynamics and radiative transport codes, and Section 3.3.2 discusses the primary uncertainty sources considered.

3.3.1 Computational Fluid Dynamics and Radiation Models

The gas is assumed to be in thermochemical nonequilibrium where the translational and rotational modes are described by a single temperature T_{tr}, and the vibrational, electronic, and free electron energy modes are described by a single temperature T_v. To model an atmosphere composed of H_2, He, and CH_4, a 12-species thermochemical model is used that considers: H_2, CH_4, CH_3, CH_2, CH, H, C_2, C, H^+, C^+, He, and e^-. Only neutral helium is considered due to its large ionization potential, as He^+ is not likely to form for $V_\infty < 30$ km/s. Other minor species (e.g. C_2H, C_2H_2, CH^+, etc.) are not considered because they are either not expected to form in the gas phase or do not contribute significantly to the aerothermodynamic quantities of interest [32]. The reaction data set for the 12-species model includes 106 chemical reactions. Dissociation of H_2 is modeled using rate data from Leibowitz [26]. Baseline atomic hydrogen ionization rate data is taken from Leibowitz and Kuo [30], but all hydrogen-ionization reactions are increased by a factor of 5 as recommended by Carroll et al. [112]. Finally, the rate data of Lino da Silva et al. [118] is used to model all carbon-based reactions. Forward reaction rates are modeled using the modified-Arrhenius form in Equation 2.18. The specifics of all 106 chemical reactions employed in this work are summarized in Appendix A.1.

As described in Section 2.3.3, diffusion in LeMANS is modeled using a modified version of Fick's law where the diffusion coefficients are computed using collision integrals. Previous work in air [83] has shown that this model well approximates the full multicomponent diffusion solution given by the Stefan-Maxwell equations. However, Erb et al. [117] found that adopting the full Stefan-

Maxwell approach was necessary due to the substantial difference in molecular weight between hydrogen and the large amount of heavier carbon-based species blown into the boundary layer via ablation. This work does not model ablation, so the boundary layer is predominantly composed of hydrogen, as the only source of carbon is the trace amount in the freestream. Therefore, the simpler modified Fick's law with species transport properties obtained from collision integrals is used. Future work that considers coupled CFD-material response should consider the more rigorous Stefan-Maxwell approach. Interactions between neutral-neutral and neutral-ion species are modeled using the approach of Gupta et al. [85], whereas, interactions between two charged species are modeled using Coulumbic theory as described by Wright et al. [88]. Collision integral fits for interactions including hydrogen and helium are taken from Palmer et al. [86], and the data from Park et al. [87] are used for all other neutral-neutral and neutral-ion interactions. The Gupta parameterization of the collision integral is given in Equation 2.36. For the 12 species considered in this work, there are 69 interactions involved corresponding to 138 total collision integrals (i.e. the $\Omega^{(1,1)}$ and $\Omega^{(2,2)}$ collision integrals). Physically, the $\Omega^{(1,1)}$ and $\Omega^{(2,2)}$ collision integrals represent how interactions between particles in the gas affect the transport properties of the mixture (i.e. viscosity, diffusivity, and thermal conductivity).

LeMANS models the relaxation of the vibrational and translational modes using the Landau-Teller model, where the vibrational-relaxation times are computed using the Millikan-White formulation [82] with Park's high-temperature correction [24]. The expression for the Millikan-White vibrational-relaxation time curve fit is given in Equation 2.14, and the values of a and b used in the Millikan-White cure fits for interactions involving H_2 and He are summarized in Table 3.1. The vibrational relaxation of CH_x species is not explicitly modeled. Because the freestream concentration of CH_4 is small, and because CH_4 tends to dissociate quickly, neglecting the vibrational relaxation of CH_x will not substantially affect the overall energy balance within the flow.

Radiative heating calculations are performed using the NASA line-by-line radiation code NEQAIR [119], which requires the temperatures and species number densities along a specific line of sight. For computational simplicity, the temperatures and densities calculated using LeMANS

Table 3.1: Millikan-White fitting parameters for interactions involving H_2 and He [86].

Interaction	a	b
H_2-H_2	65.110	6.8210×10^{-3}
H_2-He	69.971	4.6820×10^{-3}
H_2-H	9.6730	7.2500×10^{-2}

are used directly as the NEQAIR inputs. As such, the radiative heating is determined in an uncoupled fashion in which the flowfield computed by LeMANS is provided to NEQAIR, but the radiative heating source term within LeMANS is assumed to be zero. The use of a coupled LeMANS-NEQAIR simulation could be addressed in future work. NEQAIR computes the radiation spectra in a wavelength region of interest and provides the integrated radiance at specific points in space, and the wall-incident radiative heat flux is obtained using a tangent-slab approximation. As described in Chapter 2, NEQAIR provides numerous options for calculating the species excited state populations. A QSS model is used for species that have QSS rate data available, while all other species are assumed to follow a Boltzmann distribution. Helium state populations are assumed to follow a Boltzmann distribution, while the hydrogen molecule and atom populations are calculated using QSS methods. As described in [90], helium is an inert substituent, so assuming a Boltzmann state distribution for He does not affect the radiative heating calculation. Atomic carbon is modeled using a QSS assumption, and the remaining carbon-based species state populations are modeled with a Boltzmann assumption. Uncertainty in the NEQAIR databases and QSS models are not considered and are left as future work. The version of NEQAIR used in this work also implements a preliminary H free-free contribution to the total radiation calculation.

3.3.2 Sources of Uncertainty

This work considers uncertainty in each of the physical models described in Section 3.3.1. In addition to uncertain chemical reaction-rate coefficients, transport properties, and vibrational-relaxation times, uncertainty in the freestream CH_4 mass fraction is considered. There is significant

speculation over the impact of CH_4 on the resulting flowfield [35, 114], and the actual quantity of CH_4 present in the Ice Giants' atmospheres is not well known [115]. Therefore, this work seeks to include the uncertainty in the freestream CH_4 quantity to more comprehensively assess the uncertainty in the Ice Giant aerothermodynamic environment. This study treats all uncertainty as epistemic, and Table 3.2 summarizes the uncertain parameters varied in this study as well as their uncertainty multiplier intervals. The choice of the uncertainty intervals is intended to reflect that knowledge (or lack thereof) in each parameter. Intervals for the reaction-rate coefficients are informed from available rate data [26, 30, 120, 87, 29, 118], and intervals for the collision integrals and vibrational-relaxation times are taken from Palmer et al. [86]. The current data in NeptuneGRAM is believed to serve as an upper limit for the trace CH_4 levels in Neptune's atmosphere [115]. As such, the uncertainty interval for the freestream methane concentration uses the value taken from NeptuneGRAM as the upper limit and 25% of that value as the lower limit.

Table 3.2: Uncertain parameters for H_2/He/CH_4 flows.

Category	Model	Parameter Varied	Uncertainty Interval, U_i
Rate Coefficients	$k_f = A_f T_f^{\eta_f} \exp(-D_f/T_f)$	A_f	Table A.1
Collision Integrals	$\Omega^{(1,1)}, \Omega^{(2,2)} = Df(T)$	D	$\pm$10% for neutral-neutral $\pm$50% for all others [86]
V-T Relaxation Time	Millikan-White	Slope	$\pm$10%
Freestream CH_4	NeptuneGRAM	Mass fraction	$[0.25, 1.0]$

In total, 248 sources of uncertainty are considered. These include 106 forward-reaction-rate leading coefficients, 138 collision-integral leading coefficients, 3 vibrational-relaxation parameters, and the freestream methane concentration.

3.4 Selection of Freestream Conditions

This section provides an overview of the selection of freestream conditions to be analyzed using the UQ methodology presented in the previous section. A representative trajectory for an aerocapture mission at Neptune is presented in Section 3.4.1 along with the four conditions at which

the UQ analysis is run. The nominal aerothermal environments for each condition are discussed in Section 3.4.3.

3.4.1 Neptune Aerocapture Trajectory

The Neptune aerocapture entry state is derived from an existing Neptune aerocapture systems analysis [13] which gives an initial altitude and inertial velocity of $h_0 = 1000$ km and $V_0 = 29.0$ km/s, respectively. Figure 3.1 displays a representative aerocapture trajectory obtained by guiding the entry vehicle to a target orbit with an apoapsis altitude of 100,000 km and target periapsis altitude of 3,000 km using the fully numerical predictor-corrector aerocapture guidance (FNPAG) algorithm [55]. In the trajectory simulation, the vehicle is assumed to have a constant lift-to-drag ratio of $L/D = 0.3$, and the inertial entry flight path angle is taken to be $\gamma_0 = -11.5$ degrees.

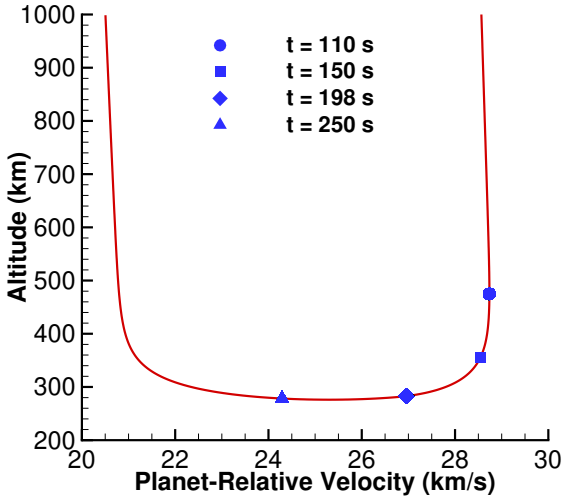

Figure 3.1: Trajectory points selected for sensitivity analysis and UQ.

Four locations along the trajectory are chosen for detailed UQ analysis. The first occurs at $t = 198$ s after entry interface and corresponds to the maximum convective heating condition, as the maximum heating typically informs the TPS material selection. The second and third points occur at $t = 110$ s and $t = 150$ s after entry interface, respectively. These two points are chosen to be before the maximum heating condition where the heat flux is still expected to be high and

the modeling of nonequilibrium phenomena is more important due to the decreased freestream density. The final point occurs at $t = 250$ s after entry interface. This point occurs after the maximum convective heating condition and is chosen to investigate the effects of lower-speed flight on the aerothermodynamic environment. Table 3.3 summarizes the freestream conditions for all four trajectory points, where the freestream density and species concentrations are obtained from NeptuneGRAM [121]. The nominal freestream CH_4 mass fraction is taken to be the midpoint of the uncertainty interval defined in Table 3.2, and the freestream H_2 and He mass fractions are modified holding the ratio of H_2/He from NeptuneGRAM constant.

Table 3.3: Freestream conditions at which detailed UQ analysis is performed.

Trajectory Point	ρ (kg/m^3)	V (km/s)	H_2/He/CH_4 Mass Fraction (%)
$t = 110$ s	3.63×10^{-7}	28.7	67.3/31.4/1.30
$t = 150$ s	3.40×10^{-6}	28.5	65.3/30.9/3.80
$t = 198$ s	1.42×10^{-5}	27.0	64.5/30.7/4.80
$t = 250$ s	1.55×10^{-5}	24.3	64.4/30.7/4.90

3.4.2 Numerical Setup

The entry vehicle is modeled as an axisymmetric, 70-degree sphere-cone with a base diameter of 5 meters and a nose radius of 1.25 meters, which is consistent with the most recent Ice Giant aerocapture studies [22, 122]. While the aerocapture trajectory is produced assuming lifting flight, quantification of the uncertainty in the full three-dimensional aerothermal quantities of interest is left as future work. A grid convergence study is performed to determine the proper resolution of the computational domain. Figure 3.2 shows the difference in the convective- and radiative-heating distributions obtained on "coarse", "medium", and "fine" grids, which have wall-tangential and wall-normal resolutions of 81×161, 161×321, and 161×641, respectively. The results obtained on the medium and fine grids deviate by less than 1% of each other for all QoIs. Therefore, all further simulations are performed on the medium-resolution grid. Mesh adaptation is performed to improve the solution quality by aligning the outer boundary of the computational domain with

the bow shock, such that 90% of the grid normal to the surface lies within the shock layer. Wall-normal grid spacing is selected to ensure $y^+ < 1$ in order to properly resolve the boundary layer. A representative computational domain is shown in Figure 3.3, along with contours of the vibrational-electronic-free-electron temperature. For all simulations, the vehicle surface is assumed to be fully catalytic and in radiative equilibrium with a surface emissivity of 0.9. Inclusion of a finite-rate surface chemistry model within the UQ analysis is left as future work. Finally, a laminar flow assumption is made. While modeling turbulent flow would increase the convective heating, the proper selection of a turbulence model and quantification of the uncertainty in its parameters is beyond the scope of this work.

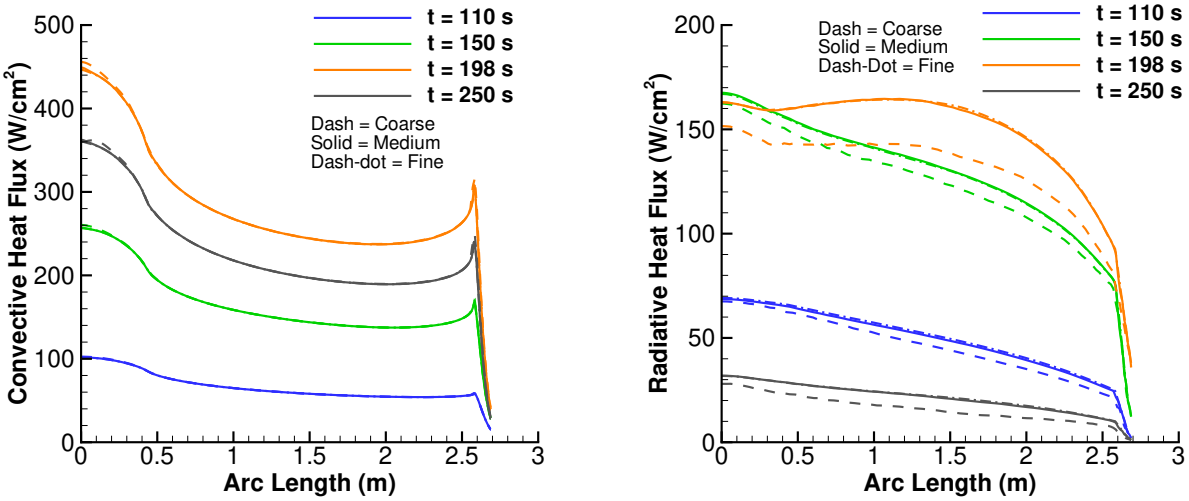

(a) Convective heating.

(b) Radiative heating.

Figure 3.2: Grid convergence study for convective (a) and radiative (b) heating.

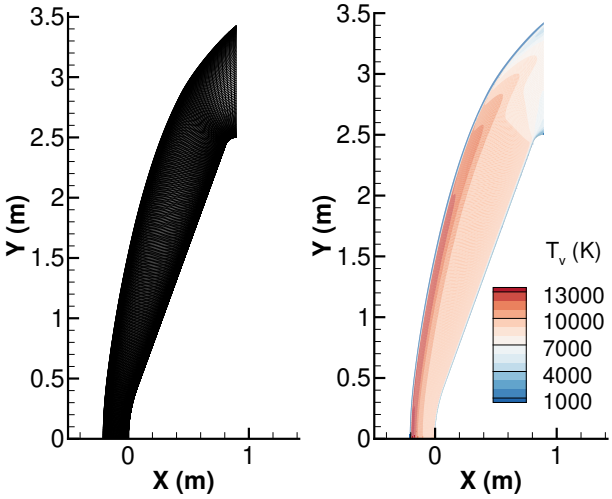

Figure 3.3: Representative computational grid and corresponding contours of vibrational-electronic-electron temperature.

3.4.3 Nominal Solutions

Before performing the detailed UQ analysis, the aerothermal environments for each trajectory point are obtained using the baseline/nominal models. At each point, the convective- and radiative-heating distributions are obtained along the vehicle surface, and the axial-force coefficient is computed using the resulting pressure and shear-force profiles.

Figure 3.4 shows the stagnation-line profiles of the temperatures and pressures for each of the trajectory points. The $t = 150$ s, 198 s, and 250 s trajectory points all reach thermal equilibrium, while the $t = 110$ s point does not. A peak translational temperature just below 25,000 K is observed, and the three cases that reach thermal equilibrium display and equilibrium temperature of approximately 10,000 K. Slower relaxation is observed for the $t = 110$ s and 150 s cases due to the lower freestream density. Furthermore, as the freestream density decreases, the shock standoff distance increases. A peak stagnation-point pressure of approximately 10 kPa is predicted for the $t = 198$ s condition. Convective- and radiative-heating distributions at each of the trajectory points are shown in Figure 3.5. Generally, the convective heating decreases as the freestream density and velocity decrease, as expected. For all but the $t = 250$ s trajectory point, the radiative heating

is a large percentage of the total heating along the surface, suggesting that radiative heating is of significant importance for freestream velocities in excess of 25 km/s. Despite a lower relative convective-heating distribution, the radiative heating at the mid-altitude trajectory point ($t = 150$ s) is similar to that at the maximum convective heating trajectory point ($t = 198$ s) with peak values of 168 W/cm^2 and 163 W/cm^2, respectively. However, the maximum-heating condition is the only trajectory point where the radiative heating does not decrease monotonically along the vehicle surface due to the more severe post-shock conditions.

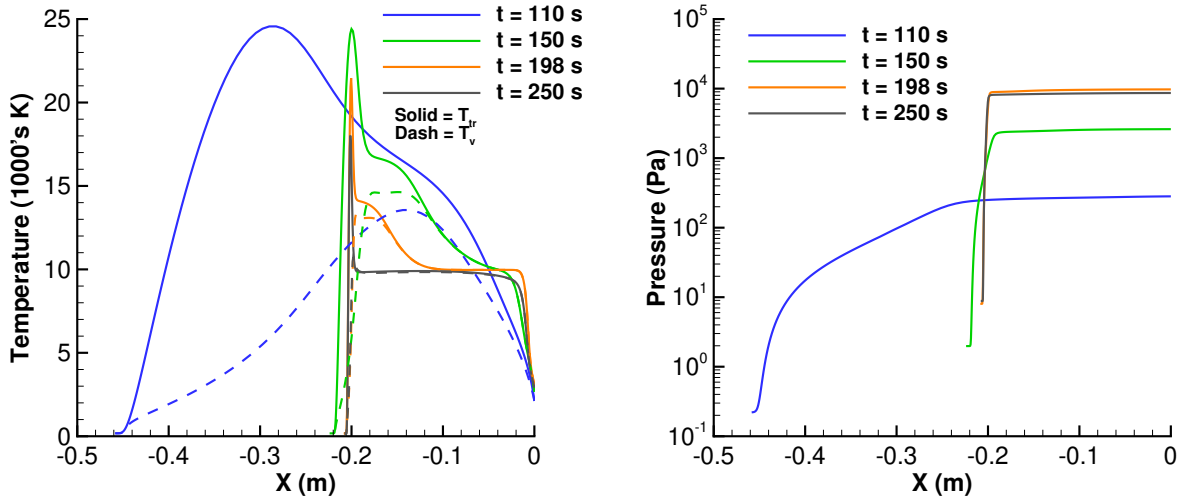

(a) Temperature.

(b) Pressure.

Figure 3.4: Nominal stagnation-line temperature (a) and pressure (b) profiles for all four trajectory points.

Figure 3.6 shows the nominal spectral and integrated radiance profiles along the stagnation line of sight for each trajectory point. For all four trajectory points, the majority of the radiative heating is accumulated between wavelengths of 100-200 nm in which the primary radiative features are the H_2 Lyman band (100-160 nm) and the H Lyman-α line (121.57 nm). The spectral radiance profiles for the maximum heating ($t = 198$s) condition and the mid-altitude condition ($t = 150$ s) show similar spectral characteristics, with higher overall spectral radiance values observed at the maximum heating condition. However, H_2 dissociates more slowly at higher altitudes, so the

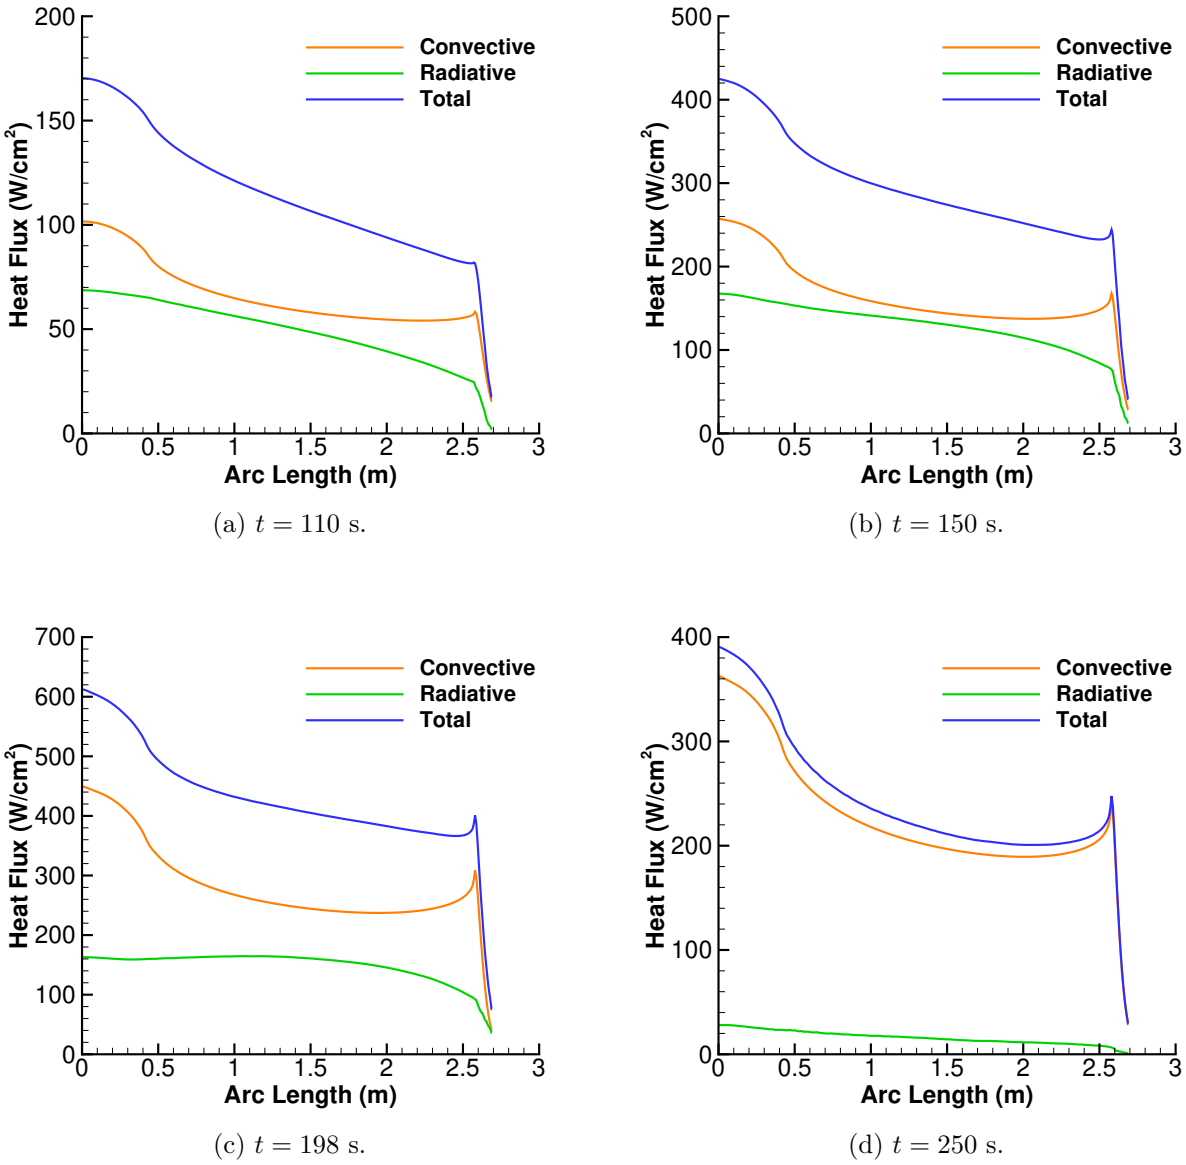

(a) $t = 110$ s.

(b) $t = 150$ s.

(c) $t = 198$ s.

(d) $t = 250$ s.

Figure 3.5: Nominal heating environments for the four trajectory points used in detailed analysis.

amount of radiance produced from the H_2-Lyman band is higher for the mid-altitude case than the max-heating case due. The three most prominent atomic lines observed are the H Lyman-α line (121.57 nm), H Balmer-α line (656.28 nm), and the H Paschen-α line (1875 nm). There are also a series of prominent carbon emission lines in the spectral range 146-250 nm. These features contribute some to the total radiance, indicating that the presence of CH_4 in the freestream does increase the radiative heating, albeit slightly. As the freestream density decreases, a lower overall contribution to radiance from atomic lines is observed, as there is both a decrease in H_2 dissociation and reduced freestream CH_4 concentration.

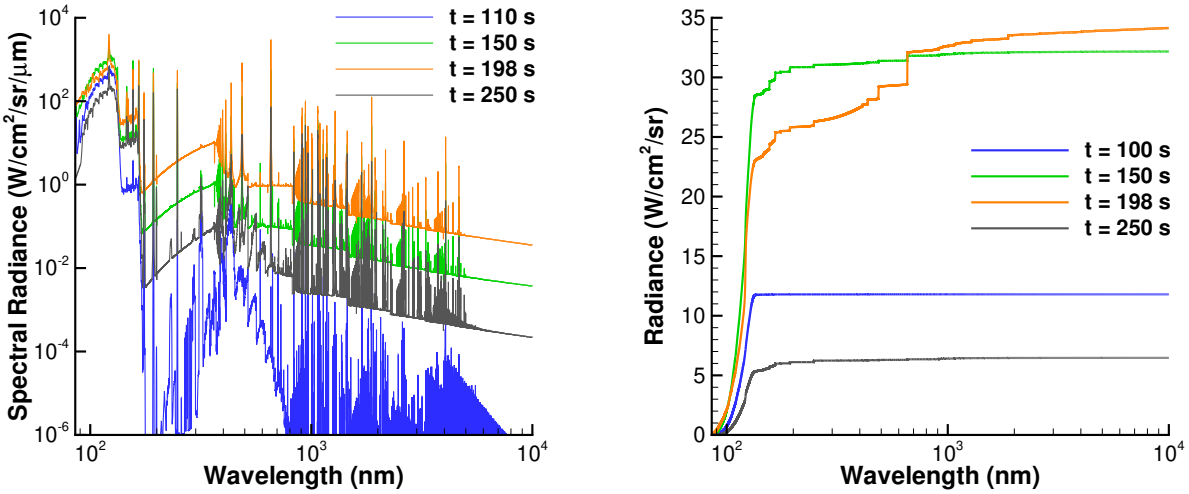

(a) Spectral radiance.

(b) Radiance.

Figure 3.6: Nominal spectral radiance (a) and integrated radiance (b) along the stagnation line of sight.

Overall, the radiative environment in these flight regimes is dominated by hydrogen features, which is consistent with previous work [37]. This is emphasized in Figure 3.7, which shows the fractional contributions of the species dominating the radiative-heating distribution. At higher altitudes (t = 110 s), nearly all of the radiative heating is from H_2. As the vehicle descends into the atmosphere, a larger contribution from H is observed due to the increasing levels of dissociation and ionization within the shock layer. However, as the freestream velocity decreases below 25 km/s,

the prominent radiative species again becomes H_2. For altitudes below 400 km, there is a minor contribution from carbon-based species due to the higher freestream CH_4 concentration. Specifically, atomic carbon is found to contribute ∼5-8% of the total stagnation-point radiative heating.

Finally, the nominal axial-force coefficients for each trajectory point are computed using the pressure and shear distributions along the vehicle surface and are summarized in Table 3.4. The axial force coefficient is relatively invariant across the four trajectory points which is expected for high hypersonic Mach numbers.

Table 3.4: Nominal values of axial-force coefficient.

Trajectory Point	C_x
$t = 110$ s	1.62
$t = 150$ s	1.64
$t = 198$ s	1.64
$t = 250$ s	1.62

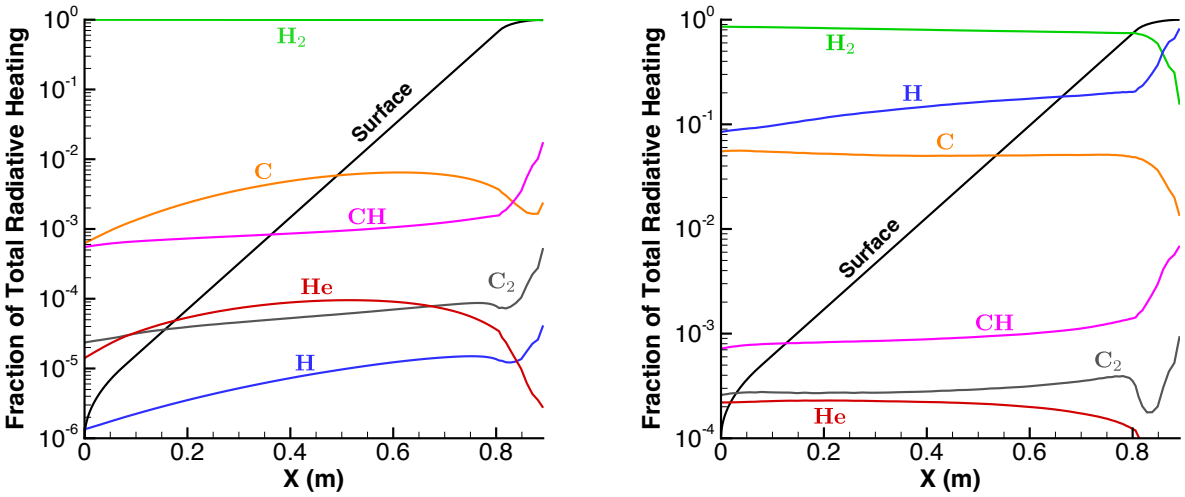

(a) $t = 110$ s.

(b) $t = 150$ s.

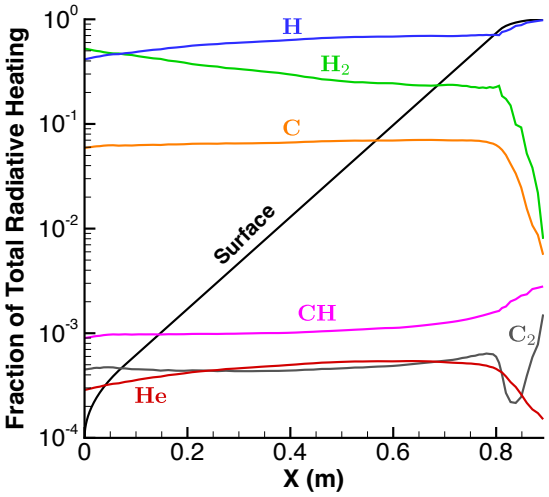

(c) $t = 198$ s.

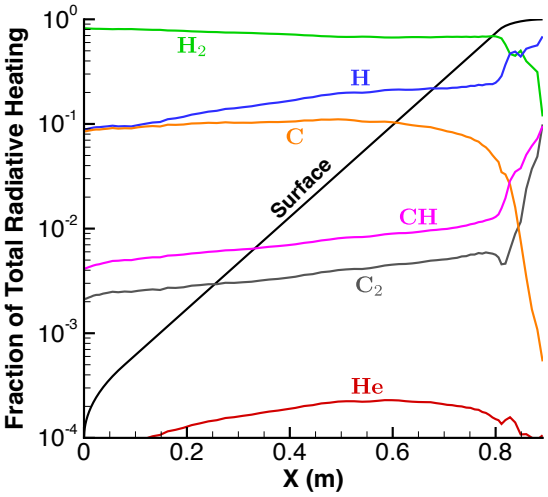

(d) $t = 250$ s.

Figure 3.7: Fractional species contribution to radiative-heating distribution at each trajectory point.

3.5 Uncertainty in Aerothermodynamic Quantities of Interest

This section presents a detailed uncertainty quantification and sensitivity analysis of the convective heating, radiative heating, and axial-force coefficient at the four trajectory points along the representative Neptune aerocapture trajectory. At each condition, the deterministic models are evaluated a total of 200 times. Next, the NIPC-PCE methodology described in Section 2.5.2 is used to determine the overall uncertainty in each QoI, and total Sobol' indices are computed to assess the global sensitivity of each QoI to the uncertain model parameters. For each QoI, the convergence of the PCE is monitored by tracking the total Sobol' indices as described in Section 2.5.5. Results for the convective and radiative heating at the four trajectory points are presented in Sections 3.5.1–3.5.4, and results for the axial-force coefficient at all four conditions are presented in Section 3.5.5.

3.5.1 Maximum Heating Condition Heating ($t = 198$ s)

Figure 3.8 shows the total Sobol' indices and epistemic uncertainty intervals for the convective heating along the vehicle surface at the maximum heating condition. Note that Figure 3.8a only shows the Sobol' indices corresponding to parameters that contribute $\geq 1\%$ to the convective-heating uncertainty. The uncertainty in the convective heating remains similar across the majority of the surface with a maximum of 10.7% above and -5.6% below the nominal prediction. At the stagnation point, the convective heating is nominally 449 W/cm^2 with a 95% confidence interval of 432–467 W/cm^2. At the shoulder (arc length = 2.6 m), the convective heating is nominally 308 W/cm^2 with a 95% confidence interval of 299–321 W/cm^2.

At the stagnation point, the variability in the convective heating is primarily attributed to the uncertainty in the H_2-H $\Omega^{(1,1)}$ collision integral and hydrogen electron-impact ionization (EII) reaction-rate coefficient, which account for 59.2% and 26.6% of the uncertainty, respectively. Because the surface is modeled as fully catalytic to ions, the species concentration gradients near the wall are primarily H_2 and H, explaining the dependence of the convective heating on the H_2-H

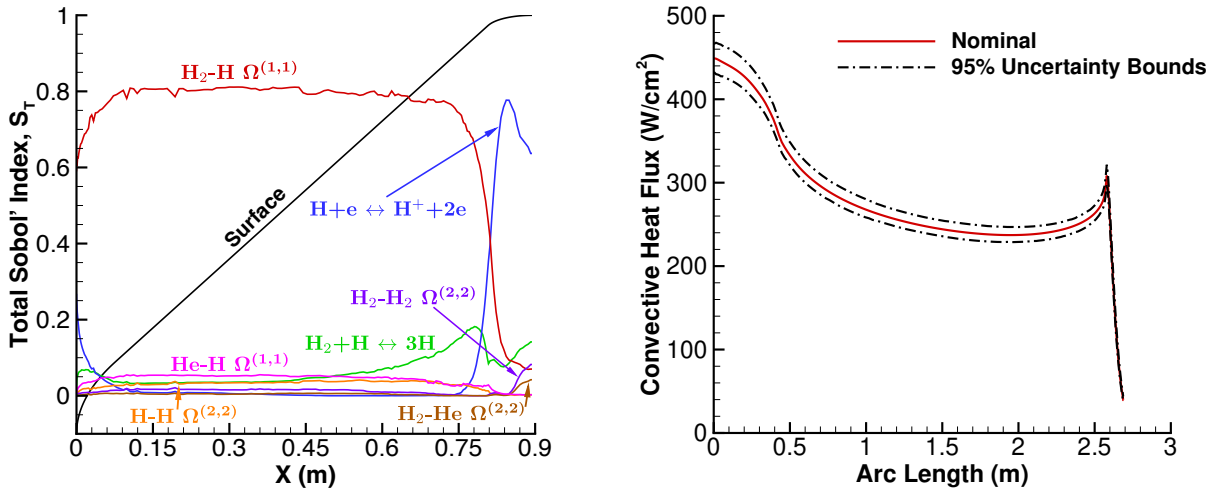

(a) Sobol' indices.

(b) Uncertainty bounds.

Figure 3.8: Total Sobol' indices (a) and epistemic uncertainty bounds (b) for the convective-heating distribution at the maximum heating trajectory point.

$\Omega^{(1,1)}$ collision integral. In addition, there is a significant level of ionization along the stagnation line. The hydrogen EII reaction strongly affects the boundary-layer edge temperature which, in turn, affects the convective heating (and transport properties throughout the boundary layer).

Moving downstream from the stagnation point, the convective-heating uncertainty becomes driven almost exclusively by collision integrals. For most of the frustum ($0 < x < 0.8$ m), the variability in the H_2-H $\Omega^{(1,1)}$ collision integral accounts for approximately 80% of the uncertainty in the convective heating. The sensitivity of the convective heating to the hydrogen EII reaction-rate coefficient nearly vanishes, as most ionization occurs in the stagnation region. Approaching the shoulder, the uncertainty in the rate coefficient for $H_2 + H \leftrightarrow 3\,H$ starts to contribute more to the convective-heating uncertainty, and it peaks around $x = 0.78$ m with a total contribution of 18.2%. The gas moving along streamlines in the boundary layer at $x = 0.78$ m has already spent a significant amount of time reacting within the boundary layer, so the hydrogen dissociation reaction will affect the species concentration distribution, and therefore, convective heating more in this region than others.

At the shoulder, the uncertainty in the convective heating becomes dominated by the hydrogen EII rate coefficient. Because the flow is cooling near the shoulder, the chemical processes in this region are dominated by recombination phenomena. To more clearly assess the effects of recombination versus ionization on the resulting heat transfer, the chemical production terms for the electron-impact ionization reaction

$$\mathrm{H} + \mathrm{e}^- \leftrightarrow \mathrm{H}^+ + 2\mathrm{e}^- \tag{3.1}$$

are computed. Hydrogen atoms are "lost" from the forward direction of Equation 3.1 and are formed from the reverse direction of Equation 3.1. As such, the net production of hydrogen atoms due to Equation 3.1 is computed as

$$r_{\mathrm{net}} = r_{\mathrm{form}} + r_{\mathrm{loss}} = n_{\mathrm{H}^+} n_{\mathrm{e}^-} n_{\mathrm{e}^-} k_b(T_b) - n_{\mathrm{H}} n_{\mathrm{e}^-} k_f(T_f) \tag{3.2}$$

where n_i is the number density of species i, and k_f and k_b are the forward and backward rate coefficients evaluated at the forward and backward controlling temperatures T_f and T_b, respectively.

Figure 3.9 shows the chemical source term for the EII reaction computed using Equation 3.2 and the nominal values for all CFD parameters. There is a significant loss of hydrogen atoms (i.e. ionization) in the stagnation region directly behind the shock. This is expected, as the gas in this region is being compressed at high temperatures. There is also significant recombination closer to the surface in the stagnation region. The amount of recombination of hydrogen ions into neutral hydrogen atoms remains high throughout much of the flowfield, and the size of the recombination zone increases approaching the shoulder. Right at the shoulder, there is a spike in the rate at which the reverse of Equation 3.1 proceeds. This local increase in recombination affects the temperature throughout the boundary layer and therefore surface heating, confirming that the large Sobol' index for the EII reaction rate at the shoulder is due to the reverse of Equation 3.1.

Figure 3.10 shows the total Sobol' indices and epistemic uncertainty intervals for the radiative heating along the vehicle surface at the maximum-heating condition. The uncertainty in the radiative heating is high, with a maximum of 68.8% above and -64.0% below the nominal prediction.

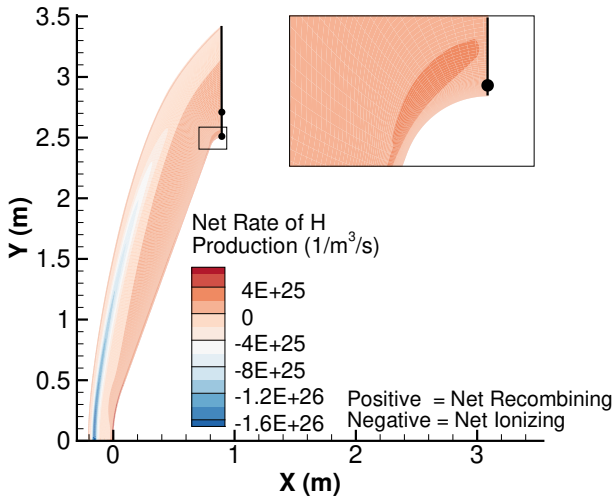

Figure 3.9: Net rate of production of hydrogen atoms from the reaction $\mathrm{H} + \mathrm{e}^- \leftrightarrow \mathrm{H}^+ + 2\,\mathrm{e}^-$ at the maximum heating condition using nominal aerothermal models.

These uncertainties are greater than the radiative heating margin used by Morgan et al. [123] to size the TPS for a Uranus aerocapture system but less than the 2σ uncertainties reported by Palmer et al. [86]. Additionally, the magnitude of the radiative heating along the entire surface is high, including the shoulder. At the stagnation point, the radiative heating is nominally 163 W/cm^2 with a 95% confidence interval of 118–254 W/cm^2. At the shoulder, the peak radiative heating is nominally 93.9 W/cm^2 with a 95% confidence interval of 78.0–132 W/cm^2.

Along the majority of the surface, the uncertainty in the radiative heating is dominated by the variability in the hydrogen dissociation reaction-rate coefficient $\mathrm{H}_2 + \mathrm{H} \leftrightarrow 3\mathrm{H}$. This is because any undissociated molecular hydrogen that persists into the high-temperature shock layer contributes significantly to the overall radiative heating. Additionally, this reaction is the primary source of atomic hydrogen, which is the dominant source of radiation at this flight condition (see Figure 3.7c). A similar finding was reported by Palmer et al. [86].

Moving downstream from the stagnation point, two interesting changes occur. First, an increase in the contribution to the radiative-heating uncertainty from the freestream CH_4 concentration is observed. As shown in Figure 3.7c, the dominant source of radiation from carbon-based

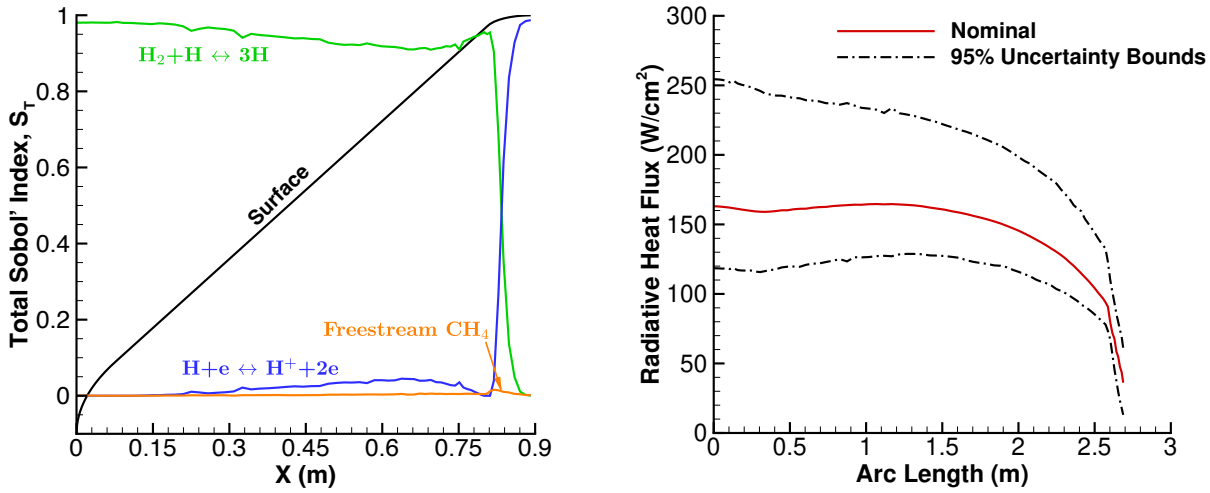

(a) Sobol' indices.

(b) Uncertainty bounds.

Figure 3.10: Total Sobol' indices (a) and epistemic uncertainty bounds (b) for the radiative-heating distribution at the maximum heating trajectory point.

species at this condition is atomic C. Additionally, ionization of atomic carbon is observed to peak downstream of the stagnation point due to the relatively slower carbon EII reaction rate. As such, any increase in the freestream CH_4 concentration results in more carbon ions, and therefore electrons, along the frustum, explaining the increase in the Sobol' index for the freestream CH_4 approaching the shoulder. Overall, the contribution of freestream CH_4 uncertainty to the variation in the radiative heating is small, with a maximum value of 2.0% at $x = 0.81$ m.

Second, at the shoulder, 99% of the uncertainty in the radiative heating is attributed to uncertainty in the hydrogen EII rate coefficient. Figure 3.11 compares the electron temperature and electron number density along the line of sight at the shoulder (shown as the black line in Figure 3.9) and the resulting spectral and integrated radiance for three cases: the nominal case, the case with the highest hydrogen EII rate (labeled "max"), and the case with the lowest hydrogen EII rate (labeled "min"). Electron temperature and electron number density are key parameters in the radiance calculation, so comparing them can provide insight into the source of the radiative heating differences. Generally, as the EII rate *decreases*, the level of ionization along the shoulder line of

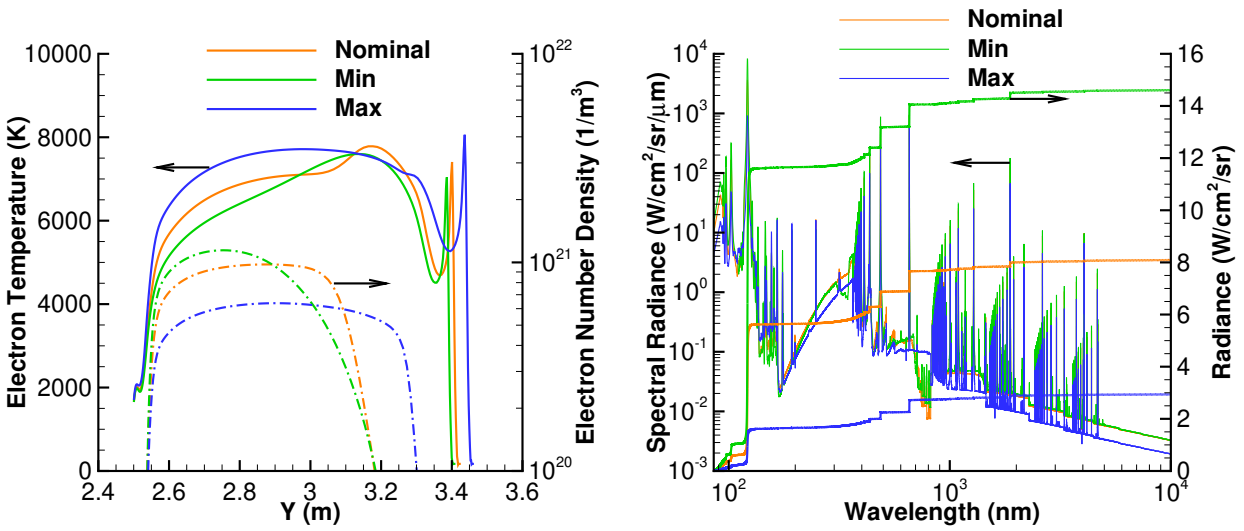

(a) T_v and n_e.

(b) Radiance.

Figure 3.11: Electron temperature and electron number density profiles (a) and radiance (b) along the shoulder line of sight for the nominal, minimum, and maximum hydrogen EII rates.

sight *increases*. This is because the shoulder is a recombination driven region, so a slower ionization rate also means a slower recombination rate. Sufficient ionization occurs in the stagnation region for all three cases, so the electron number density at the shoulder is significantly influenced by the EII reaction recombination rate. Additionally, as the EII reaction rate decreases, the radiance at the shoulder increases.

In particular, there are significant differences in the radiance from the H Lyman-α line at 121.57 nm. Emission from the H Lyman-α line corresponds to an electron transitioning from H($n = 2$) to H($n = 1$), so the decreased EII rate causes the population of the second energy level ($n = 2$) of atomic hydrogen to increase. This is confirmed in Figure 3.12a which shows the atomic hydrogen electronic state population computed by NEQAIR along the shoulder line of sight for the nominal and lowest hydrogen EII rates. The three colored lines denote the first three electronic levels, and the grayed lines denote all higher electronic states. On average, the population of H($n = 2$) is higher when the EII rate is decreased, which increases the contribution of the H Lyman-α line to the total radiative heating. Additionally, Figure 3.12b compares the atomic

hydrogen state populations for the minimum hydrogen EII rate at two locations (shown as the black circles in Figure 3.9) along the shoulder line of sight computed assuming nonequilibrium state population distributions and state populations determined assuming a Saha-Boltzmann distribution. For both locations along the shoulder line of sight, the nonequilibrium H($n = 2$) state population is the furthest below the corresponding Saha-Boltzmann value, while other states are much closer to (or even on top of) equilibrium with the free electrons and continuum species. Because H($n = 2$) is the furthest from Saha-Boltzmann equilibrium, its population is most likely to be affected by any changes in the electron number density. Due to the increased electron number density profile for the "min" case in Figure 3.11b, there should be a larger net flux of electrons to the H($n = 2$) state through both ground state excitation and recombination. The recombination event is more likely near the surface, where the H($n = 2$) state is significantly depleted, which is consistent with the conclusion that the shoulder is a recombination-driven region.

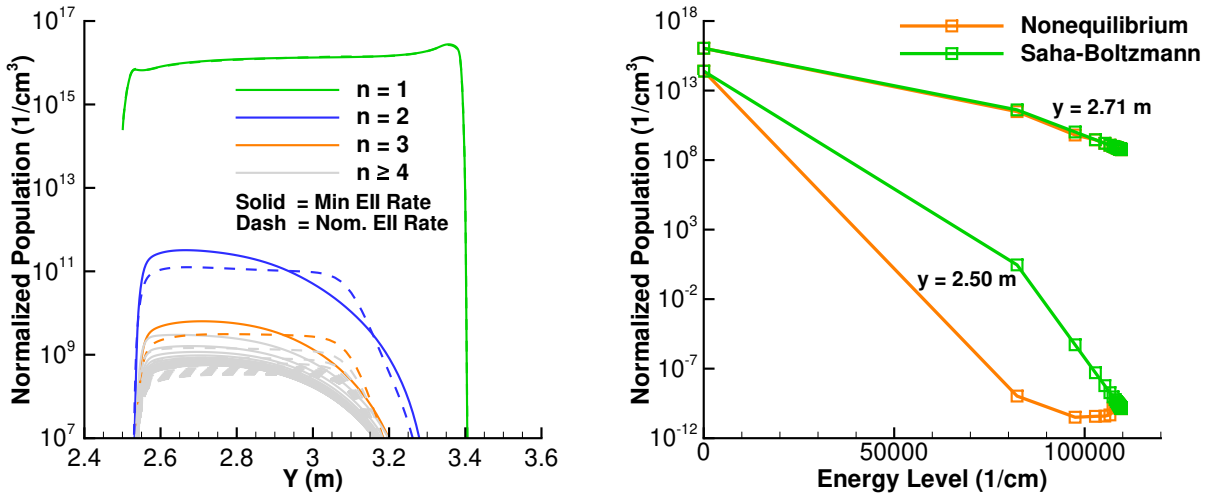

(a) State populations.

(b) Comparison to Saha-Boltzmann.

Figure 3.12: Electronic state populations of H along the shoulder line of sight (a) and comparison of state populations computed using nonequilibrium and Saha-Boltzmann models (b).

3.5.2 Mid-Altitude Condition Heating ($t = 150$ s)

Figure 3.13 shows the total Sobol' indices and epistemic uncertainty intervals for the convective heating along the vehicle surface at the mid-altitude condition. The uncertainty in the convective heating is slightly larger than that computed for the maximum heating condition, with a maximum of 7.5% above and -4.4% below the nominal prediction. At the stagnation point, the convective heating is nominally 257 W/cm^2 with a 95% confidence interval of 246–273 W/cm^2. At the shoulder, the convective heating is nominally 168 W/cm^2 with a 95% confidence interval of 162–176 W/cm^2.

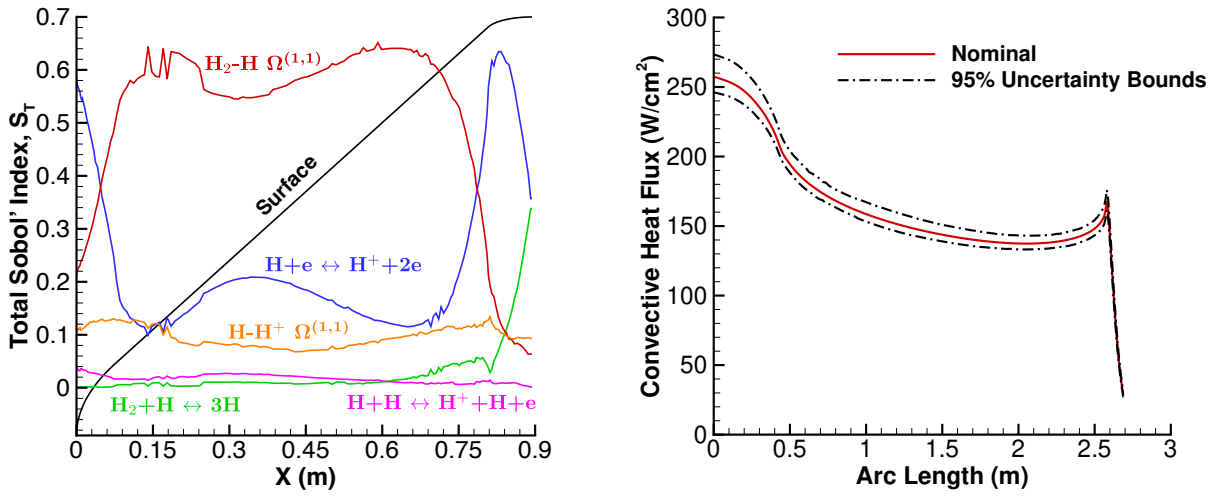

(a) Sobol' indices.

(b) Uncertainty bounds.

Figure 3.13: Total Sobol' indices (a) and epistemic uncertainty bounds (b) for the convective-heating distribution at the mid-altitude trajectory point.

At the stagnation point, the primary source of uncertainty in the convective heating is the hydrogen EII reaction-rate coefficient, which accounts for 58.6% of the total uncertainty. The remainder of the stagnation-point convective heat flux uncertainty is attributed to collision-integral uncertainty, with the largest contribution coming from the H_2-H $\Omega^{(1,1)}$ collision integral at 20.9%. This trend is the opposite of what is found for the maximum-heating condition, in which the primary source of the stagnation-point convective heat flux uncertainty is the H_2-H $\Omega^{(1,1)}$ collision

integral. Compared to the maximum-heating condition, the mid-altitude trajectory point has a lower freestream density and a higher freestream velocity. As a result, the EII reaction heavily affects the shock standoff distance and, therefore, temperature distribution throughout the stagnation region. This is emphasized in Figure 3.14, which shows the stagnation-line temperature and electron number density profiles for the 200 cases used to construct the PCE. A substantial variation in shock standoff distance is observed, which explains the large contribution of the uncertainty in the hydrogen EII rate coefficient to the uncertainty in the stagnation-point convective heating. The mid-altitude freestream conditions are in somewhat of an intermediate zone, as the ionization rate is either quite slow or quite fast, depending on the value of k_f for Equation 3.1. If k_f for Equation 3.1 is sufficiently large, then EII reactions will proceed faster causing the average stagnation line temperature and shock standoff distance to decrease. Conversely, if the k_f for Equation 3.1 is sufficiently small, the net rate of ionization is slower, and the shock-layer translational temperature and shock standoff distance are larger. Such a high level of variability is observed because the freestream density is just high enough to produce sufficient collisions in the post-shock flow to the point where the ionization is relatively fast (if k_f is greater than nominal), but it is also just low enough to produce relatively slow ionization (if k_f is lower than nominal). Furthermore, for EII to be active, the heavy-particle ionization reactions must first produce a sufficient density of free electrons, which explains the smaller contribution of the rate coefficient for the reaction $\mathrm{H} + \mathrm{H} \leftrightarrow \mathrm{H}^+ + \mathrm{e}^-$ to the convective-heating uncertainty.

The contribution from the hydrogen EII reaction-rate coefficient is less pronounced moving downstream from the stagnation point. Along the frustum, up to 64% of the convective-heating uncertainty is attributed to the $\mathrm{H_2}$-H $\Omega^{(1,1)}$ collision integral, and up to 21% is attributed to the hydrogen EII reaction rate. There is also a smaller allocation of uncertainty to the H-H^+ $\Omega^{(1,1)}$ collision integral, which accounts for approximately 9-10% of the convective-heating uncertainty. The collision integrals and hydrogen EII reaction rate drive the transport properties within the boundary-layer and the temperature at the boundary-layer edge, respectively, which explain their large contribution to the uncertainty in the convective heat flux.

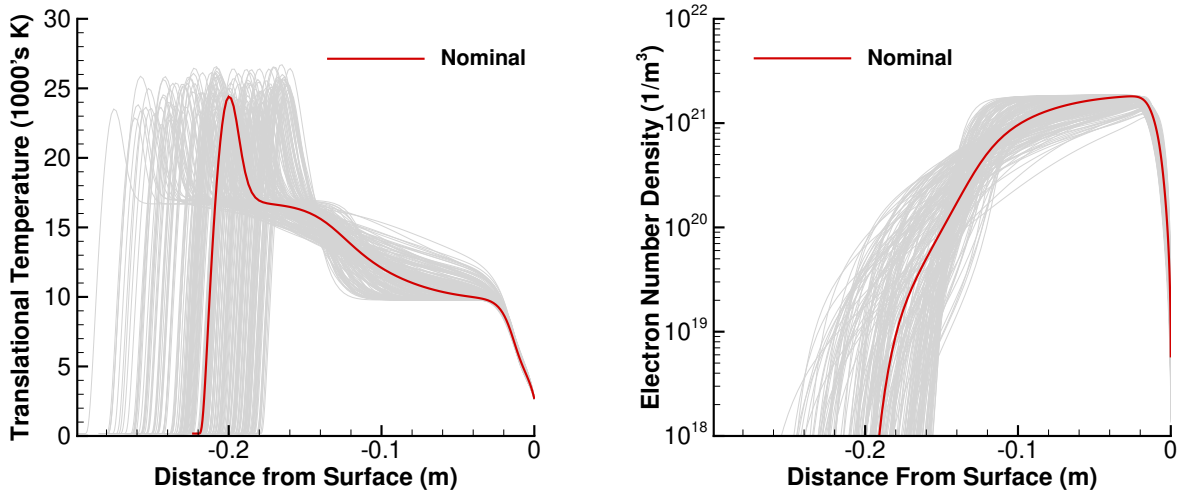

(a) Temperature.

(b) Electron number density.

Figure 3.14: Profiles of the stagnation line translational temperature (a) and electron number density (b) at the mid-altitude trajectory point.

Rounding the shoulder, a similar phenomenon as the maximum-heating condition is observed, where the majority of the convective heat-flux uncertainty is attributed to the hydrogen EII reaction-rate coefficient. Right before the shoulder ($x = 0.83$ m), 63.5% of the uncertainty in the convective heat flux is attributed to the hydrogen EII rate coefficient. Similar to the maximum-heating condition, this is again due to the shoulder environment being driven by recombination phenomena. Additionally, because the lower freestream density produces less overall ionization, a more pronounced sensitivity to the reaction-rate coefficient of $H_2 + H \leftrightarrow 3\,H$ is observed. At the shoulder, the uncertainty in this rate coefficient accounts for 33.5% of the convective-heating uncertainty.

Figure 3.15 shows the total Sobol' indices and epistemic uncertainty intervals for the radiative heating along the vehicle surface at the mid-altitude condition. Uncertainty in the radiative heating is high, with a maximum of 96.9% above and -49.1% below the nominal prediction. The magnitude of the radiative heating along much of the surface is large, but it is smaller around the shoulder. At the stagnation point, the radiative heating is nominally 168 W/cm^2 with a 95% confidence interval of 90.6–318 W/cm^2. At the shoulder, the peak radiative heating is nominally 78.0 W/cm^2 with a

95% confidence interval of 43.5–152 W/cm^2.

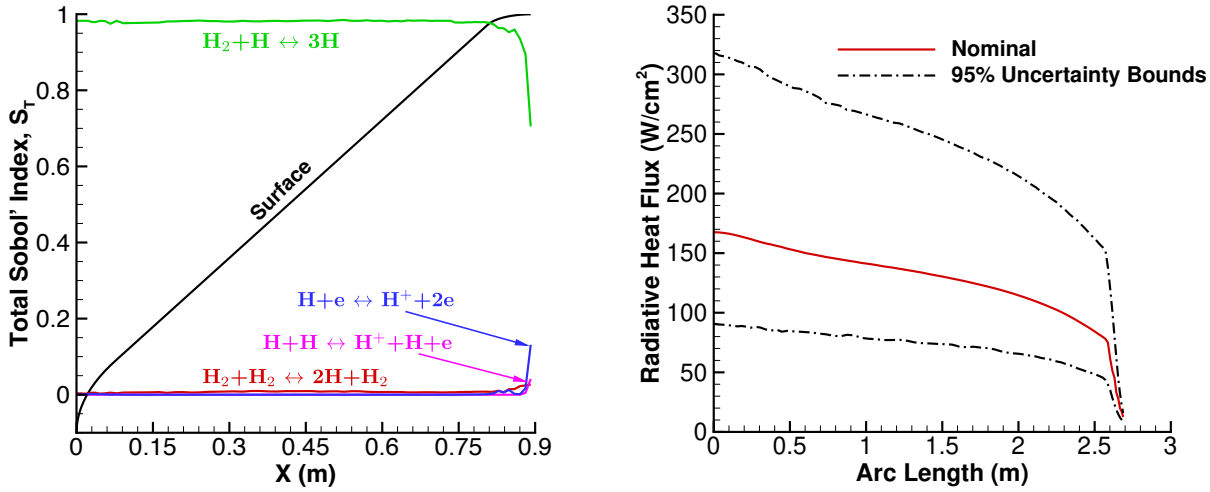

(a) Sobol' indices.

(b) Uncertainty bounds.

Figure 3.15: Total Sobol' indices (a) and epistemic uncertainty bounds (b) for the radiative-heating distribution at the mid-altitude trajectory point.

Similar to the maximum-heating condition, the majority of the radiative-heating distribution uncertainty is attributed to the uncertainty in the rate coefficient for the reaction $H_2 + H \leftrightarrow 3\,H$. For most of the surface ($x < 0.8$ m), uncertainty in this rate coefficient accounts for more than 98% of the variability in the radiative heating. Because the freestream density is lower, molecular hydrogen persists deeper into the shock layer which contributes significantly to the radiative heating. The presence of molecular hydrogen deeper within the shock layer is also emphasized by the minor contribution ($\sim 1\%$) of the reaction-rate coefficient for the dissociation of H_2 with partner H_2, which is not identified as a primary source of uncertainty at the maximum-heating condition.

At the shoulder, a similar trend to the maximum-heating condition is observed, where the contribution of the hydrogen dissociation reaction-rate coefficient to the radiative-heating uncertainty decreases, and the contribution from hydrogen EII reaction-rate coefficient increases. The magnitude of this decrease, however, is significantly smaller than what is observed at the maximum-heating condition. Here, 71.4% of the uncertainty in the radiative heating is attributed to uncertainty in the

hydrogen dissociation rate coefficient, while only 12.6% is attributed to uncertainty in the hydrogen EII rate coefficient. This is explained by the lower overall levels of ionization in the flowfield, so the shoulder heating phenomenon discussed in the previous section is less pronounced. It is further emphasized by the minor contribution (4.00%) of the rate coefficient for the heavy-particle ionization reaction $H + H \leftrightarrow H^+ + H + e^-$ to the shoulder radiative heating. This reaction is a precursor reaction to the hydrogen EII reaction, so it has a more pronounced effect when EII does not dominate the ionization process. Unlike the maximum-heating condition, the radiative heating is not observed to be sensitive to the freestream CH_4 concentration. At higher altitudes, the mole fraction of CH_4 in the atmosphere is lower, so there is far less contribution to the net radiative heating from carbon-based species.

3.5.3 High-Altitude Condition Heating ($t = 110$ s)

Figure 3.16 shows the total Sobol' indices and epistemic uncertainty intervals for the convective heating along the vehicle surface at the high-altitude condition. There is a noticeable increase in the convective-heating uncertainty along the vehicle frustum, with a maximum of 5.09% above and -11.7% below the nominal prediction. At the stagnation point, the convective heating is nominally 102 W/cm^2 with a 95% confidence interval of 98.5–106 W/cm^2. At the shoulder, the convective heating is nominally 58.4 W/cm^2 with a 95% confidence interval of 52.7–60.3 W/cm^2.

A substantial change in the processes dominating the convective-heating uncertainty is observed relative to the previous two trajectory points. At the stagnation point, uncertainty in collision integrals accounts for 58.4% of the convective-heating uncertainty, with 39.8% of that coming from the H_2-H $\Omega^{(1,1)}$ collision integral. Due to the much lower freestream density, H_2 does not completely dissociate along the stagnation line, which explains why 24.4% of the convective-heating uncertainty is attributed to the uncertainty in rate coefficient for the dissociation of H_2 with partner H. As there is far less ionization in the stagnation region, this reaction largely dictates the stagnation-line species concentration distribution, which explains why a larger variety of collision integrals is affecting the convective heating variability. For the little ionization that does occur in the stagnation region, the

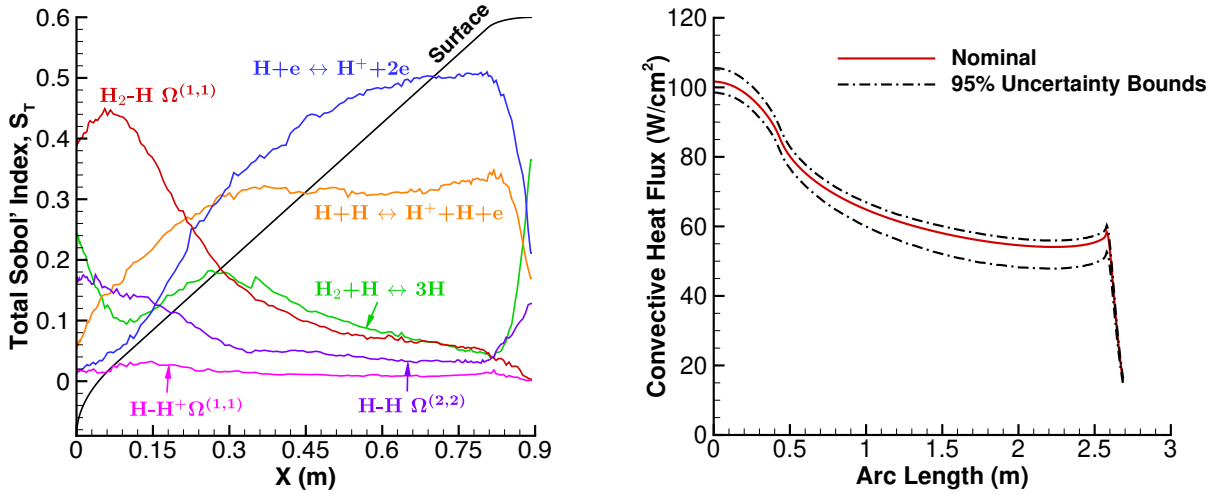

(a) Sobol' indices.

(b) Uncertainty bounds.

Figure 3.16: Total Sobol' indices (a) and epistemic uncertainty bounds (b) for the convective-heating distribution at the high-altitude trajectory point.

hydrogen heavy-particle ionization reaction $\mathrm{H}+\mathrm{H} \leftrightarrow \mathrm{H}^+ + \mathrm{H} + \mathrm{e}^-$ affects the convective heating the most, contributing about 5.4% to the overall convective-heating uncertainty. Very little sensitivity of the convective heating to the hydrogen EII reaction-rate coefficient is observed at the stagnation point because insufficient ionization has occurred along the stagnation line for this reaction to be particularly active.

Moving downstream from the stagnation point, the sensitivity of the convective heating to the uncertainty in the collision integrals decreases, and the sensitivity to the hydrogen ionization reaction-rate coefficients increases. Because the density at the high-altitude trajectory point is so low, the majority of atomic hydrogen ionization occurs downstream of the stagnation region, as shown in Figure 3.17. Note that, for $0 \leq x \leq 0.3$ m, the rate coefficient for the ionization of atomic hydrogen with partner atomic hydrogen contributes more to the overall uncertainty in convective heating than the EII rate coefficient does, whereas, for $x > 0.3$ m, the opposite is true. This is because a sufficient amount of free electrons must first be produced before the EII reaction can become active, and the primary source of free electrons is initially from the heavy-particle

ionization reaction $H + H \leftrightarrow H^+ + H + e^-$. The delayed ionization of hydrogen also explains why the lower epistemic uncertainty bound in Figure 3.16b is further from the nominal prediction downstream of the stagnation point, as the hydrogen EII reaction causes more cooling of the flow near the boundary layer. From Figure 3.17, it is interesting to note that the hydrogen EII reaction is primarily moving in the forward direction (i.e. ionizing), with a much smaller region near the surface where it is net moving in the reverse direction (i.e. recombination). This is in contrast to the previous two trajectory points, where the hydrogen EII is almost always moving in the reverse direction downstream from the stagnation region.

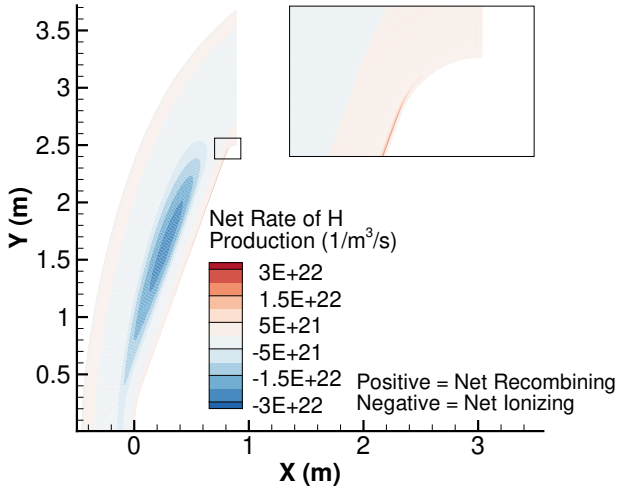

Figure 3.17: Net rate of production of hydrogen atoms from the reaction $H + e^- \leftrightarrow H^+ + 2\,e^-$ at the high-altitude condition using nominal aerothermal models.

Figure 3.18 shows the total Sobol' indices and epistemic uncertainty intervals for the radiative heating along the vehicle surface at the high-altitude condition. The uncertainty in the radiative heating is high, with a maximum of 161% above and -60% below the nominal prediction. The magnitude of the radiative heating along much of the surface is lower than the previous two trajectory points, almost vanishingly so around the shoulder. At the stagnation point, the radiative heating is nominally 68.7 W/cm^2 with a 95% confidence interval of 29.9–151 W/cm^2. At the shoulder, the peak radiative heating is nominally 24.6 W/cm^2 with a 95% confidence interval of 10.0–64.2 W/cm^2.

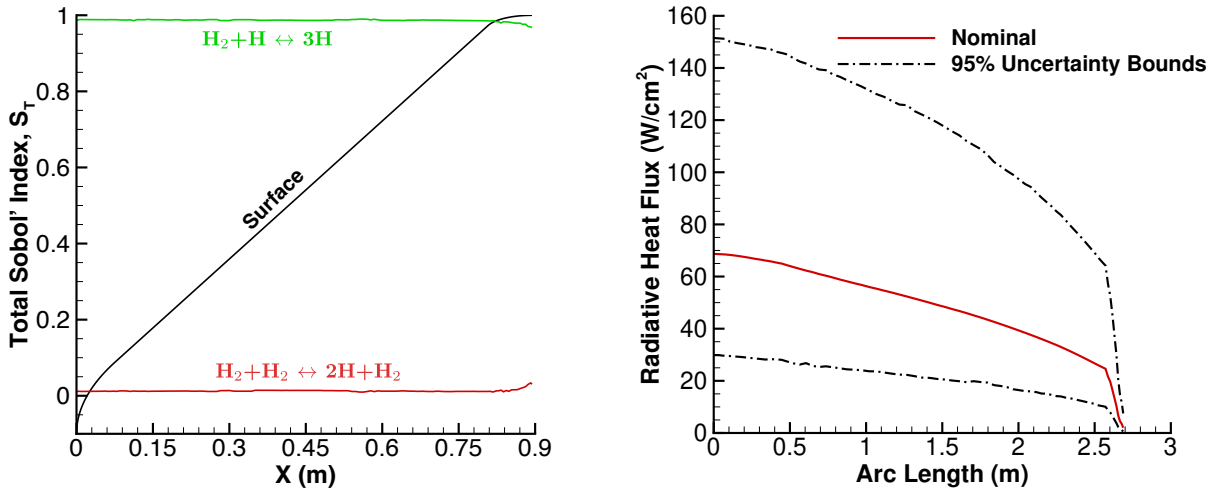

(a) Sobol' indices.

(b) Uncertainty bounds.

Figure 3.18: Total Sobol' indices (a) and epistemic uncertainty bounds (b) for the radiative-heating distribution at the high-altitude trajectory point.

All of the uncertainty in the radiative heating is attributed to the uncertainty in the rate coefficients for hydrogen dissociation. At the stagnation point, the uncertainties in the rate coefficients for H_2 dissociation with partners H and H_2 account for 98.7% and 1.10% of the radiative-heating uncertainty, respectively. Compared to the mid-altitude trajectory point, a slightly larger overall allocation of uncertainty is observed for the dissociation of H_2 with partner H_2 because of the even lower freestream density. Unlike either of the other two trajectory points, the radiative heating does not appear to be sensitive to any ionization reactions. The amount of ionization here is relatively small, and, as shown in Figure 3.17, the hydrogen EII recombination phenomena observed at the previous two trajectory points is not apparent at this trajectory condition. Therefore, it is not expected that the radiative heating, particularly at the shoulder, be affected by ionization.

3.5.4 Lower-Speed Condition Heating ($t = 250$ s)

Figure 3.19 shows the total Sobol' indices and epistemic uncertainty intervals for the convective heating along the vehicle surface at the $t = 250$ s condition. The uncertainty in the convective

heating remains similar across the majority of the surface with a maximum of 4.90% above and -4.10% below the nominal prediction. At the stagnation point, the convective heating is nominally 362 W/cm^2 with a 95% confidence interval of 349–374 W/cm^2. At the shoulder, the convective heating is nominally 241 W/cm^2 with a 95% confidence interval of 235–249 W/cm^2.

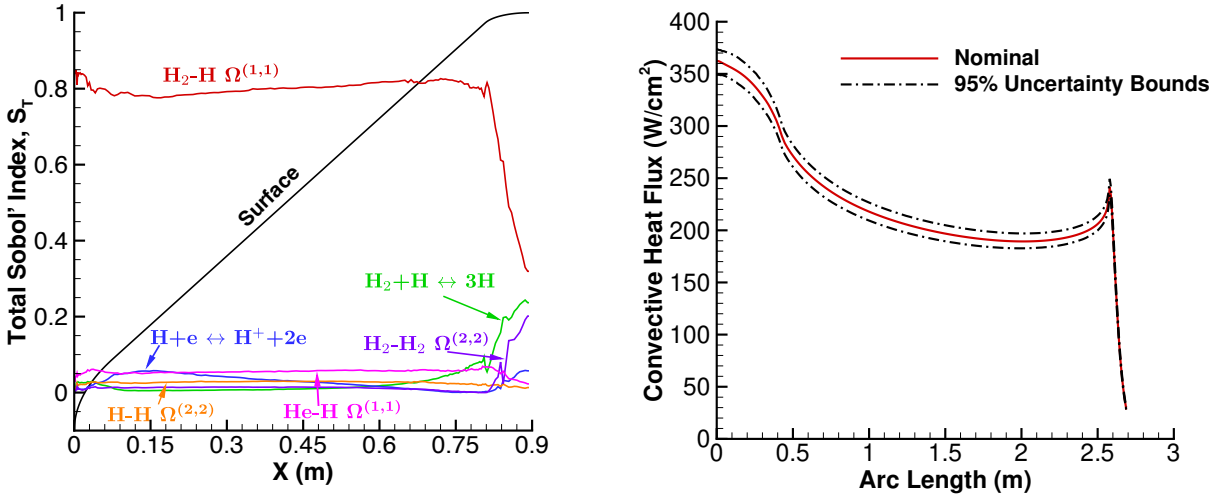

(a) Sobol' indices.

(b) Uncertainty bounds.

Figure 3.19: Total Sobol' indices (a) and epistemic uncertainty bounds (b) for the convective-heating distribution at the $t = 250$ s trajectory point.

Similar to the maximum-heating ($t = 198$ s) condition, the variability in the stagnation-point convective heating is primarily attributed to the uncertainty in the H_2-H $\Omega^{(1,1)}$ collision integral, which accounts for 84% of the uncertainty. Unlike the maximum-heating condition, however, little sensitivity of the stagnation-point convective heating to the hydrogen EII rate coefficient is observed. This is because the lower freestream velocity does not produce significant ionization within the shock layer. The sensitivity levels for collision integrals He-H $\Omega^{(1,1)}$ and H-H $\Omega^{(2,2)}$ are similar to those at the maximum convective heating condition.

Moving downstream from the stagnation point, the convective-heating variability becomes driven almost exclusively by the uncertainty in the collision integrals, with a minor contribution from the hydrogen EII rate coefficient. The small contribution from the hydrogen EII rate coefficient

is explained by the small amount of ionization that occurs downstream from the stagnation region. As the flow cools and begins recombining near the shoulder, the sensitivity of the convective heating to the hydrogen EII reaction-rate coefficient decreases, and an increased sensitivity to the rate coefficient for the reaction $H_2 + H \leftrightarrow 3\,H$ is observed. Right at the shoulder, this reaction accounts for 24.4% of the convective heating-uncertainty.

To more clearly assess the effects of recombination versus dissociation on the resulting heat transfer, the chemical source terms for the dissociation of hydrogen with partner atomic hydrogen (i.e. $H_2 + H \leftrightarrow 3\,H$) are computed. Similar to Equation 3.2, the net rate production of hydrogen molecules from this reaction is computed as

$$r_{\text{net}} = r_{\text{form}} + r_{\text{loss}} = n_{\text{H}} n_{\text{H}} n_{\text{H}} k_b(T_b) - n_{\text{H}_2} n_{\text{H}} k_f(T_f) \quad . \tag{3.3}$$

Figure 3.20 shows the net production of H_2 computed using Equation 3.3 and the nominal values for all CFD parameters. Due to the high temperatures, there is a significant loss of hydrogen molecules (i.e. dissociation) in the majority of the shock layer. As the flow expands around the vehicle shoulder, the shock-layer temperatures decrease, which causes a net formation (i.e. recombination) of hydrogen molecules. Along the surface, the transition from dissociation to recombination does not occur until directly at the shoulder, where the increase in the total Sobol' index for the $H_2 + H \leftrightarrow 3\,H$ rate coefficient is observed. This indicates that it is the variability in the recombination processes which is affecting the uncertainty in the convective heating at the shoulder, not dissociation. This is consistent with the findings of previous sections. The recombination of H into H_2 also explains the increased contribution of the H_2-H_2 $\Omega^{(2,2)}$ collision integral to the uncertainty in the convective heating at the shoulder.

Figure 3.21 shows the total Sobol' indices and epistemic uncertainty intervals for the radiative heating along the vehicle surface at the $t = 250$ s condition. The uncertainty in the radiative heating is high, with a maximum of 112% above and -65.8% below the nominal prediction. The magnitude of the radiative heating along the entire surface is far lower than the previous trajectory points. At the stagnation point, the radiative heating is nominally 32.0 W/cm^2 with a 95% confidence interval

of 16.4–65.1 W/cm^2. At the shoulder, the peak radiative heating is nominally 9.92 W/cm^2 with a 95% confidence interval of 4.61–20.2 W/cm^2.

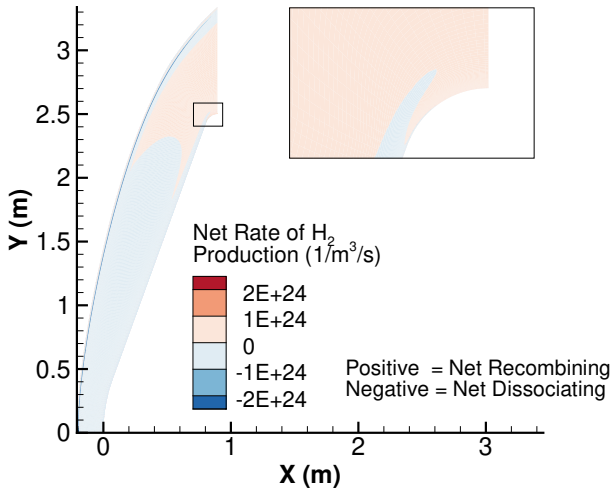

Figure 3.20: Net rate of production of hydrogen molecules from the reaction $H_2 + H \leftrightarrow 3\,H$ at the $t = 250$ s condition using nominal aerothermal models.

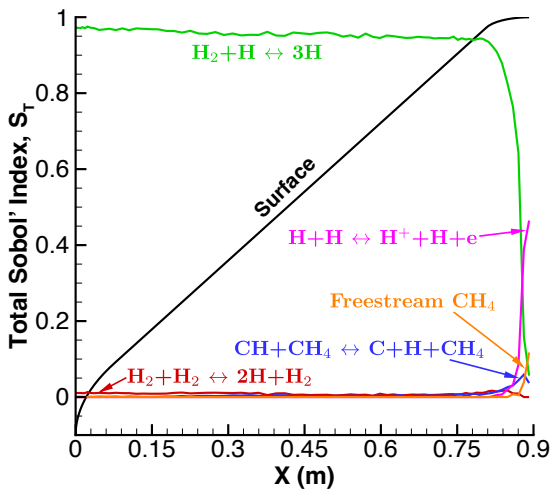

(a) Sobol' indices.

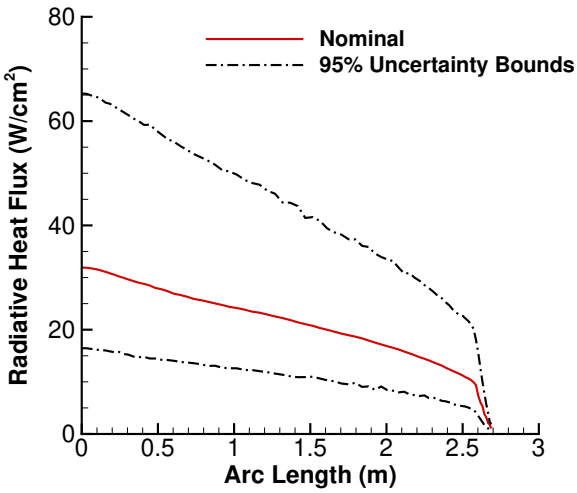

(b) Uncertainty bounds.

Figure 3.21: Total Sobol' indices (a) and epistemic uncertainty bounds (b) for the radiative-heating distribution at the $t = 250$ s trajectory point.

The majority of the uncertainty in the radiative heating is attributed to the rate coefficient for H_2 dissociation. At the stagnation point, the uncertainties in the rate coefficient for H_2 dissociation with partners H and H_2 account for 97.1% and 1.06% of the radiative-heating uncertainty, respectively. Compared to the maximum-heating trajectory point, a slightly larger overall allocation of uncertainty is observed for the dissociation of H_2 with partner H_2 because of the lower freestream velocity. Approaching the shoulder, the ionization of H via heavy-particle and electron impact start to contribute more to the radiative-heating uncertainty. Instead of the H EII reaction dominating the shoulder radiative-heating uncertainty, it is the heavy-particle ionization of H with partner H that contributes most. The heavy-particle ionization reactions of H are precursor reactions to the H EII reactions, and they are proportionally more active at this lower-speed condition. Similar to the maximum-heating condition, it is the reverse direction of these reactions that is primarily influencing the variability of the radiative heating at the shoulder. Finally, a slight sensitivity to the rate coefficient for the reaction $CH + CH_4 \leftrightarrow C + H + CH_4$ is observed at the shoulder. This reaction is an additional source of C and H, of which H is the dominant source of radiation at this location.

3.5.5 Axial-Force Coefficient

Finally, consideration is given to the effect of the uncertainty in the CFD input parameters on the variability of the aerodynamic coefficients of the entry vehicle. The aerodynamics will directly influence the range of feasible aerocapture trajectories, so quantifying the uncertainty in them is useful for overall mission design. Figure 3.22 shows the convergence of the total Sobol' indices for the axial-force coefficient at all four trajectory points. The total Sobol' index profiles for all trajectory points display a rapid level of convergence, except for the high-altitude trajectory point, which requires more samples to converge.

Overall, the primary sources of uncertainty in the axial-force coefficient are the hydrogen ionization reaction-rate coefficients, with a slight contribution of the rate coefficient for the dissociation of H_2 with partner H at the high-altitude condition. These reactions dictate the shape of

the shock wave around the vehicle and, therefore, aerodynamic coefficients. Table 3.5 summarizes the allocation of uncertainty in the axial-force coefficient for all four trajectory points, and Table 3.6 summarizes the epistemic uncertainty bounds computed from the PCE. The uncertainty in the axial-force coefficient is small, with a maximum value of only a couple percent. Current aerodynamic margins for the aerodynamic coefficients of Ice Giant aerocapture entry vehicles are on the order of 3-5% for the axial-force coefficient and 10-20% for the normal force-coefficient [124]. While this work cannot estimate the normal-force coefficient uncertainty due to the axisymmetric flow assumption, these uncertainties are generally greater than the uncertainty predicted in this work. No sensitivity to the freestream CH_4 concentration is observed, as expected.

Table 3.5: Allocation of uncertainty in axial-force coefficient to uncertain CFD input parameters.

Parameter	Contribution to C_x Uncertainty (%)			
	$t = 110\ s$	$t = 150\ s$	$t = 198\ s$	$t = 250\ s$
$\mathrm{H} + \mathrm{e} \leftrightarrow \mathrm{H}^+ + 2\,\mathrm{e}^-$	66.7	91.5	91.4	52.1
$\mathrm{H} + \mathrm{H} \leftrightarrow \mathrm{H}^+ + \mathrm{H} + \mathrm{e}^-$	25.9	8.38	8.83	45.6
$\mathrm{H}_2 + \mathrm{H} \leftrightarrow 3\,\mathrm{H}$	11.3	0.00	0.00	0.0

Table 3.6: Epistemic uncertainty bounds for the axial-force coefficient at each trajectory point.

Trajectory Point	95% Epistemic Uncertainty Bounds
$t = 110$ s	1.615–1.637
$t = 150$ s	1.616–1.663
$t = 198$ s	1.618–1.652
$t = 250$ s	1.618–1.629

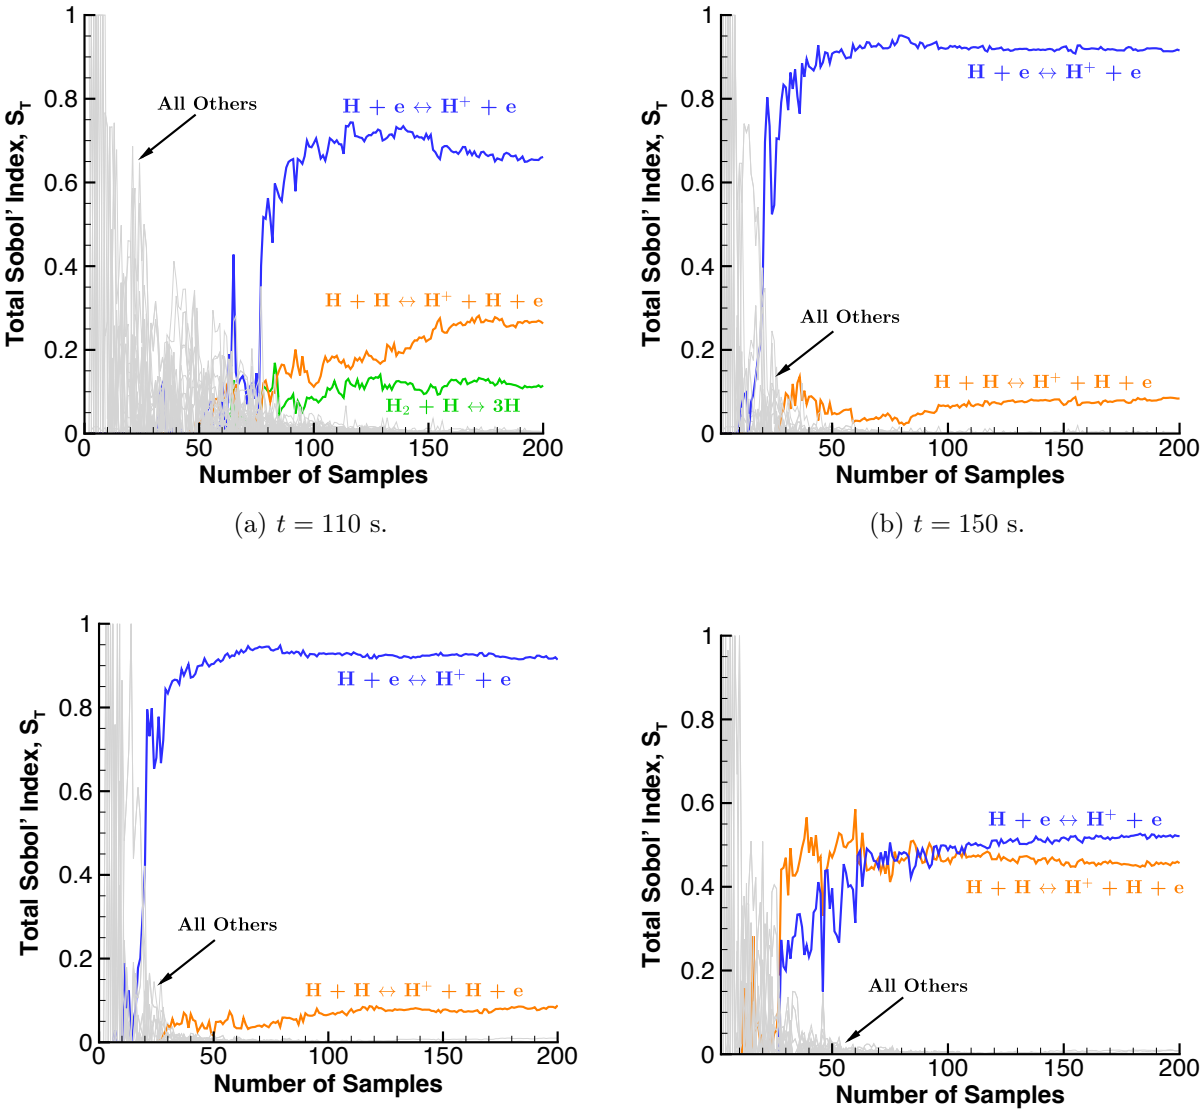

(a) $t = 110$ s.

(b) $t = 150$ s.

(c) $t = 198$ s.

(d) $t = 250$ s.

Figure 3.22: Convergence of total Sobol' indices for the axial-force coefficient at all four trajectory points.

3.6 Effects of H_2^+

As shown in the previous section, the primary ionization reaction to which the radiative-heating distributions are sensitive is the hydrogen EII reaction. In order for an EII reaction to proceed, a sufficient number of free electrons must first be produced. In the 12-species thermochem-

ical model used in this work, the only source of these initial free electrons are from the heavy-particle ionization of H. Therefore, the point at which EII becomes active depends heavily on the rate at which the heavy-particle ionization proceeds. An alternative path to the production of free electrons is associative ionization (AI). While AI is the dominant source of the initial free electrons in air [125], it has historically been omitted for hydrogen-helium flows because it typically only occurs if one of the neutral collision partners is in an electronically excited state [29]. Unlike oxygen and nitrogen in air, hydrogen does not have low-lying metastable excited electronic states, motivating the belief that an insufficient amount of electronically excited hydrogen is present to drive the AI reaction. However, Section 3.5.1 demonstrates that there is electronic excitation of atomic hydrogen, so it may be that AI is non-negligible. Carroll et al. [40] investigated the effects of including H_2^+ in a state-specific simulation of a 27.66 km/s shock tube experiment using an H_2/He mixture. They concluded that the inclusion of H_2^+ was not necessary because it did not significantly alter the post-shock species distributions or radiance calculations in the specific conditions they studied. However, because this work considers a broader range of freestream conditions, neglecting H_2^+ may have an impact on the aerothermodynamic environments.

Figure 3.23 compares a Boltzmann-weighted rate using the data in [126] to the hydrogen AI rate proposed by Lino da Silva [118]. More details on the Boltzmann-weighted rate may be found in Appendix A.2. The rate of Lino da Silva is derived from dissociative recombination (DR) data [127] and the principle of detailed balance. For comparison, the hydrogen heavy-particle ionization rates from Leibowitz et al. [30] are also shown. The Boltzmann-weighted rate overpredicts the rate of hydrogen AI relative to the rate of Lino da Silva because of the underlying Boltzmann state distribution assumption. Additionally, both the Boltzmann-weighted rate and the rate of Lino da Silva are greater than the heavy-particle hydrogen ionization reaction-rate coefficients, indicating that the inclusion of AI would result in an earlier production of electrons and ionized species.

The remainder of this section quantifies the effect of including H_2^+ on the nominal aerothermodynamic environments and aerothermodynamic uncertainty. To do so, four additional reactions are added to the existing rate set, the details of which are summarized in Appendix A.1. Note that

these reactions are specifically chosen because they are the same reactions modeled by Carroll et al. [40] to assess the effects of H_2^+.

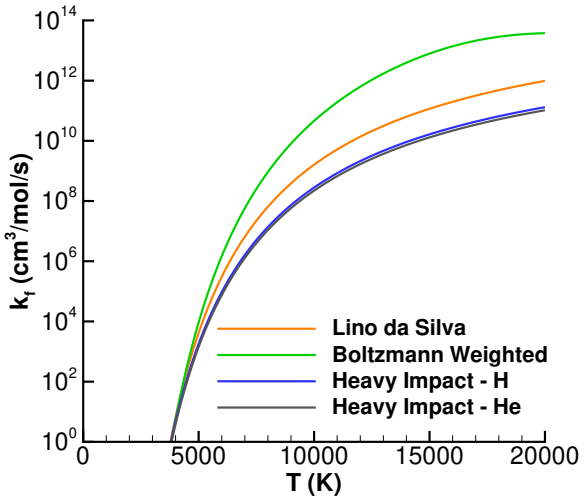

Figure 3.23: Comparison of hydrogen associative ionization rates.

3.6.1 Effect on Nominal Aerothermodynamic Environments

Figure 3.24 compares the convective- and radiative-heating distributions for all four trajectory points using the 12- and 13-species thermochemical models (i.e. with and without H_2^+). The convective-heating distributions are relatively unchanged. The inclusion of H_2^+ slightly increases the convective heating at all trajectory points, with the exception of the point at $t = 110$ s. The decrease in the convective heating with the addition of H_2^+ at the $t = 110$ s condition is due to the decreased shock-layer temperatures produced by slightly greater levels of ionization. In contrast, the radiative-heating profiles are significantly affected by the inclusion of H_2^+. Specifically, the inclusion of H_2^+ decreases the radiative heating at the stagnation point and along most, if not all, of the frustum for all four cases. Table 3.7 quantifies the effect of H_2^+ on the stagnation-point radiative heat flux for all four trajectory points. This decrease is most pronounced for the $t = 150$ s trajectory point because, as was shown in the previous section, the radiative heating produced at this trajectory condition is particularly sensitive to the rate of the ionization just behind the shock.

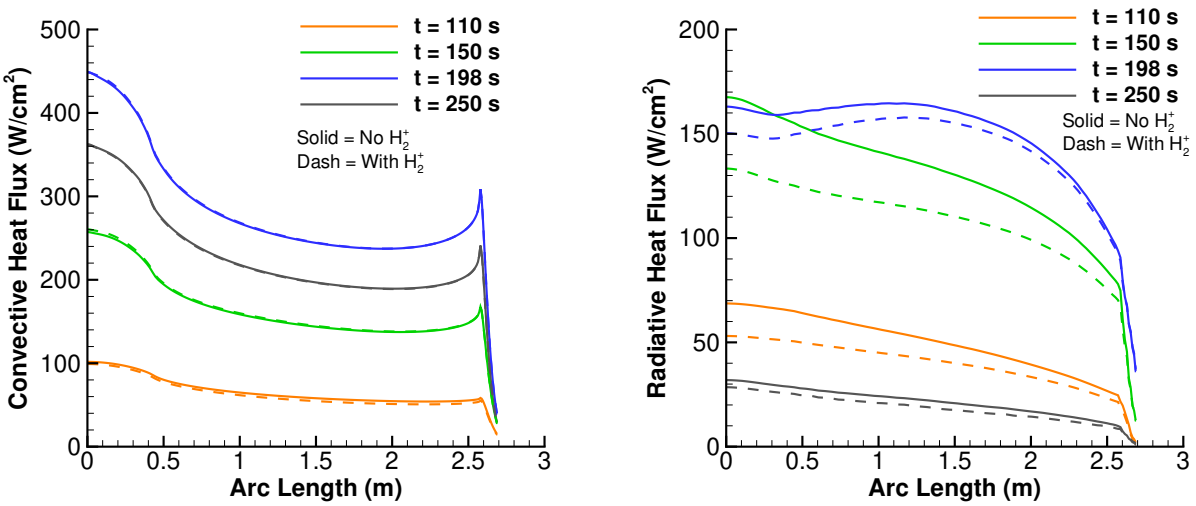

(a) Convective heating.

(b) Radiative heating.

Figure 3.24: Nominal convective (a) and radiative (b) heating profiles with and without $H_2{}^+$.

As shown in Figure 3.23, the rate of hydrogen AI is faster than the corresponding heavy-particle ionization rates. Therefore, the inclusion of $H_2{}^+$ should cause a faster rise in the electron number density behind the shock. This is confirmed by Figure 3.25a, which compares the stagnation-line profiles of the vibrational-electronic-electron temperature, T_v, and electron number density, n_e, produced with the 12- and 13-species thermochemical models at the $t = 150$ s trajectory point. A faster rise in electron number density allows the hydrogen EII reactions to dominate the ionization process more quickly behind the shock. As a result, the energy lost from translation due to ionization cools the gas more quickly and produces a decrease in shock standoff distance. The more rapid initiation of the hydrogen EII reaction is more clearly illustrated in Figure 3.25b, which shows the nominal net rate of the hydrogen EII reaction along the stagnation line at the $t = 150$ s point with and without $H_2{}^+$. Here, the x-axis is normalized by the shock standoff distance for each case to more clearly visualize the relative location behind the shock at which EII becomes active. Because surface-incident radiative heating is a function of the volume of the surrounding radiating gas, a smaller shock layer will result in a more optically thick gas and produce lower overall radiative

Table 3.7: Effect of H_2^+ on stagnation-point radiative heating.

Trajectory Point	$q_{stag,rad}$ (W/cm^2) No H_2^+	With H_2^+	Percent Change (%)
$t = 110$ s	68.7	53.1	-22.7
$t = 150$ s	168	133	-20.8
$t = 198$ s	163	150	-7.98
$t = 250$ s	32.0	28.5	-10.9

heating. Along the stagnation line, the decrease in radiation is primarily attributed to a decrease in the H_2 Lyman-α band. While the spectral locations of the radiative features are nearly identical in the 12- and 13-species models, the decreased temperatures produced by the 13-species models result in the dominant radiative features being slightly lower in magnitude than those produced by the 12-species model.

While including H_2^+ produces a clear decrease in radiative heating along the entire surface for three trajectory points, its effects diminish along the vehicle surface for the maximum-heating ($t = 198$ s) condition. Because, in this work, the highest levels of ionization are produced at this trajectory point, the effects of any additional/faster ionization introduced by H_2^+ are less pronounced. This particularly becomes the case moving downstream for the stagnation region, where the slower heavy-particle ionization reactions have had more time to contribute to the overall ionization level. The blunted effect of H_2^+ on the radiative heating at the maximum-heating condition is also emphasized by the fact that the stagnation-point radiative heating is decreased by a proportionally smaller amount when compared to the other three trajectory points.

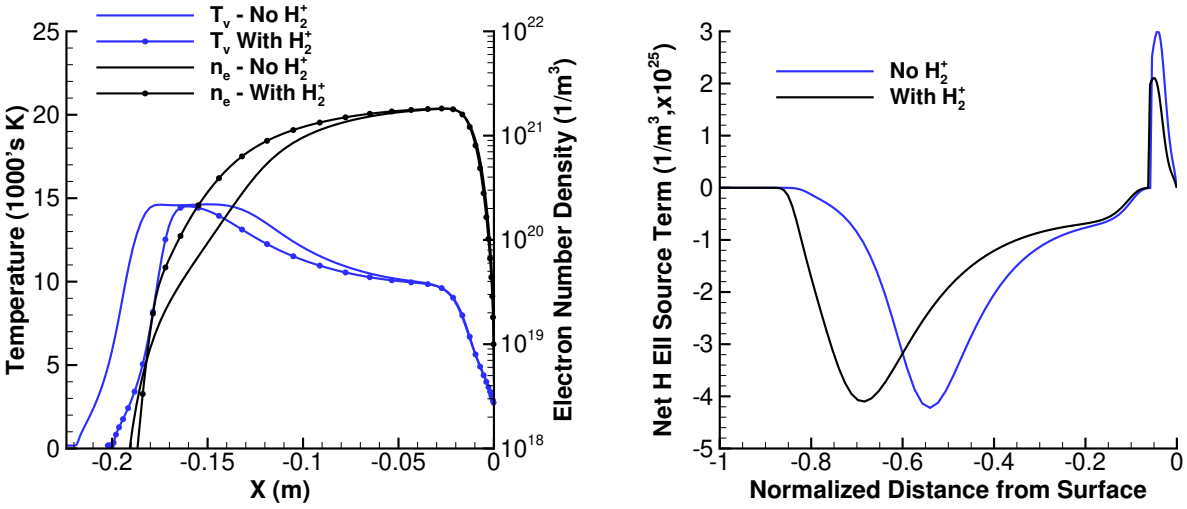

(a) T_v and n_e.

(b) Net rate of hydrogen EII.

Figure 3.25: Nominal stagnation-line temperature and electron number density profiles (a) and net H EII rate (b) at the mid-altitude ($t = 150$ s) trajectory point with and without $H_2{}^+$.

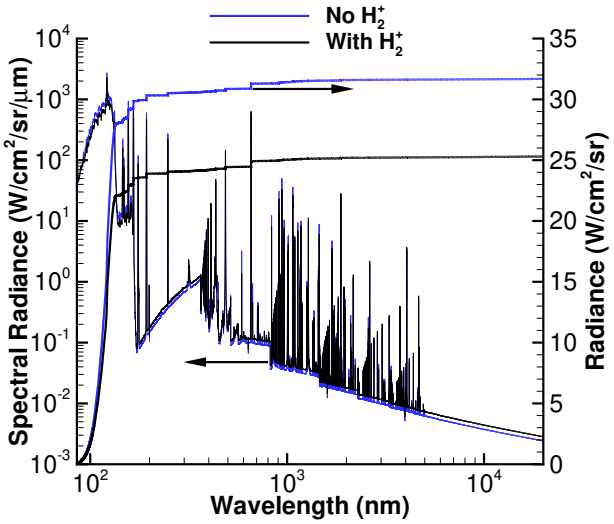

Figure 3.26: Comparison of radiance profiles produced at the $t = 150$ s condition with and without $H_2{}^+$.

3.6.2 Effect on Aerothermodynamic Uncertainty and Sensitivities

Uncertainty intervals for the reaction-rate coefficients involving H_2^+ are taken from Table A.2. The effect of including H_2^+ on the 95% epistemic uncertainty bounds of the stagnation-point convective heating is summarized in Table 3.8. Similar to the nominal convective-heating distributions, the inclusion of H_2^+ produces little change in the convective-heating uncertainty, as the majority of the convective-heating uncertainty is driven by uncertainties in collision integrals. The $t = 110$ s condition shows the only decrease in the convective-heating-uncertainty bounds, while the other three conditions produce a slight increase. At the $t = 110$ s condition, the low freestream densities cause much of the shock layer to be in thermal nonequilibrium. The introduction of H_2^+ thus results in a faster thermal relaxation process which produces less severe gradients near the surface. This effect is less pronounced at the other trajectory points because of the higher freestream densities. The relative differences shown in Table 3.8 are consistent across the majority of the vehicle surface.

Table 3.8: Effect of H_2^+ on stagnation-point convective heating epistemic uncertainty bounds.

Trajectory Point	95% Bounds (W/cm^2)		Percent Change (%)
	No H_2^+	With H_2^+	
$t = 110$ s	98.5–106	95.3–104	-3.25/-1.89
$t = 150$ s	246–273	250–278	+1.63/+1.83
$t = 198$ s	432–467	435–471	+0.69/+0.86
$t = 250$ s	349–374	351–377	+0.57/+0.80

Figure 3.27 displays the total Sobol' indices for the convective-heating distribution at the $t = 110$ s and $t = 150$ s trajectory points. The total Sobol' index profiles for the $t = 198$ s and $t = 250$ s trajectory points look similar to the results without H_2^+ and are therefore omitted for brevity. While the level of uncertainty in the convective heating is relatively unchanged with the addition of H_2^+, the parameters contributing most to that uncertainty change substantially. This is particularly the case for the high-altitude ($t = 110$ s) condition, where the less dense gas is more sensitive to changes in chemical reaction rates. For the $t = 110$ s condition, the convective heating

is no longer sensitive to the rate coefficient for the heavy-particle ionization of H with partner H. Because the hydrogen AI rate is faster than this heavy-particle ionization rate, this result is expected. Instead, as shown in Figure 3.27a, nearly 50% of the convective-heating uncertainty in the stagnation region is attributed to the rate coefficients for the dissociation and exchange reactions involving H_2^+. Similar to the case with no H_2^+, the hydrogen EII rate coefficient begins to dominate the convective-heating uncertainty approaching the shoulder, and the recombination of H into H_2 dominates the uncertainty at the shoulder.

Moving to the $t = 150$ s trajectory point, differences in the total Sobol' indices for the convective heating are observed again. Specifically, the stagnation-point convective heating is less sensitive to the hydrogen EII rate coefficient with the inclusion of H_2^+. Previously, variability in the hydrogen EII rate coefficient accounted for nearly 60% of the convective-heating uncertainty, but Figure 3.27b indicates that the inclusion of H_2^+ drops its contribution to approximately 35%. In turn, the rate coefficient for the exchange reaction $H_2 + H^+ \leftrightarrow H_2^+ + H$ contributes to just over 10% of the convective-heating uncertainty. Moving downstream from the stagnation point, the sensitivities look similar to the case without H_2^+, with most of the uncertainty along the frustum being produced by collision-integral uncertainty.

At the maximum-heating condition ($t = 198$ s), the total Sobol' indices again look similar to the results obtained without H_2^+. A slight decrease in the hydrogen EII rate coefficient total Sobol' index is observed at the stagnation point, however, the large majority of the convective-heating uncertainty is still driven by variability in the H_2-H $\Omega^{(1,1)}$ collision integral. The reverse direction of the hydrogen EII reaction also drives the uncertainty in the shoulder convective heating. Notably, unlike the results with no H_2^+, a slight sensitivity (∼1.0-2.5%) of the convective heating to the freestream CH_4 concentration is observed. Because the inclusion of H_2^+ produces a more rapid rise in free electrons behind the shock, the carbon EII reaction will be more active. Because carbon is only present in trace quantities, the rate at which carbon ions are produced will be particularly sensitive to any changes in the amount of CH_4 in the freestream. This also explains why the observed sensitivity to the freestream CH_4 concentration is small.

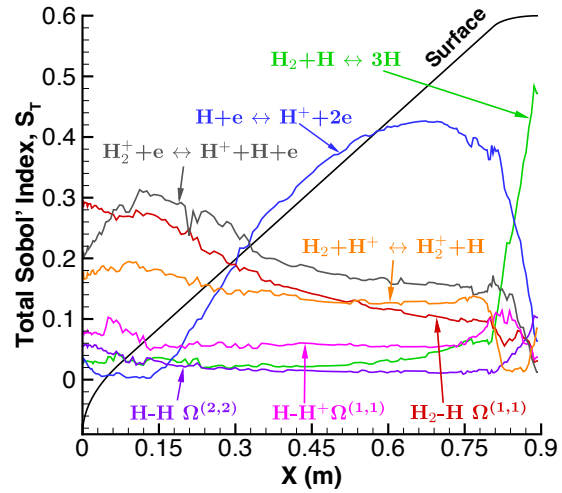

(a) $t = 110$ s.

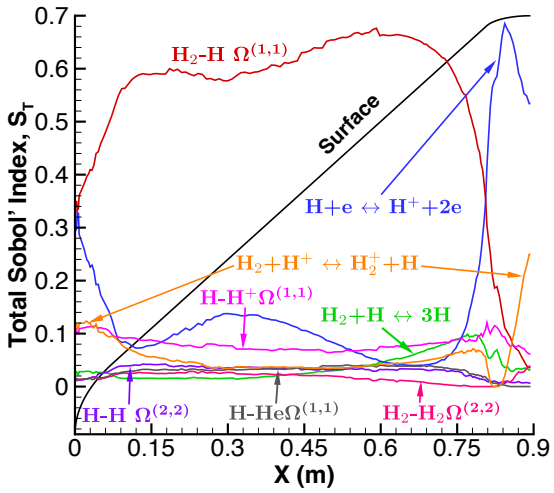

(b) $t = 150$ s.

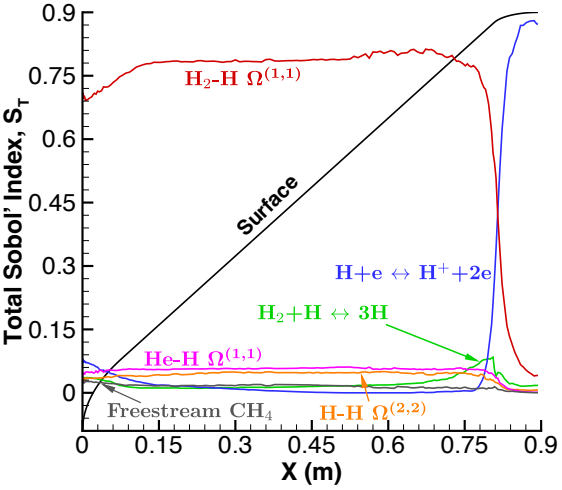

(c) $t = 198$ s.

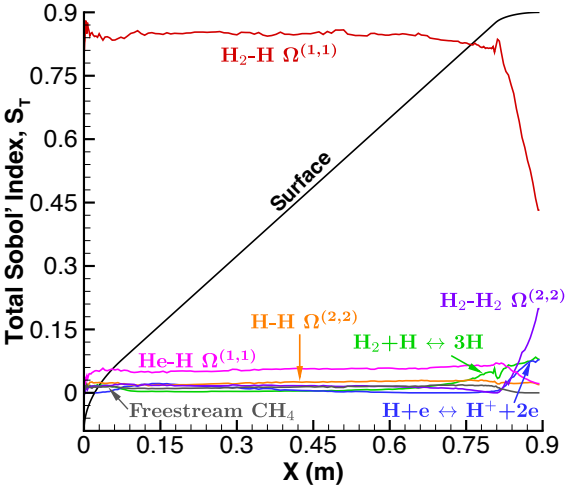

(d) $t = 250$ s.

Figure 3.27: Total Sobol' indices for the convective-heating distribution including H_2^+.

Finally, at the $t = 250$ s condition, nearly identical total Sobol' index profiles are obtained when considering H_2^+ to those obtained without H_2^+. Similar to the $t = 150$ s and $t = 198$ s trajectory points, a slight blunting of the total Sobol' index for the hydrogen EII rate coefficient is observed. A small sensitivity ($\sim 1\%$) of the convective heating to the freestream CH_4 concentration is also observed, but, like the case with no H_2^+, most of the convective-heating uncertainty is attributed to variability in the H_2-H $\Omega^{(1,1)}$ collision integral.

The effect of including H_2^+ on the 95% epistemic uncertainty bounds of the stagnation-point radiative heating is summarized in Table 3.9. Unlike the convective-heating uncertainty, the inclusion of H_2^+ produces a substantial decrease in the upper bound of the 95% confidence interval, in excess of 40% at the $t = 110$ s condition. These findings indicate that the inclusion of H_2^+ not only decreases the nominal radiative heating, but it also shrinks the total uncertainty in the radiative-heating distributions (at least for the conditions considered in this work). This is particularly impactful for the $t = 150$ s trajectory point which has the most severe upper bound on radiative heating. A general decrease in shock standoff distance is produced with the inclusion of H_2^+, which lowers the post-shock temperatures and resultant radiative heating.

Table 3.9: Effect of H_2^+ on stagnation-point radiative heating epistemic uncertainty bounds.

Trajectory Point	95% Bounds (W/cm^2)		Percent Change (%)
	No H_2^+	With H_2^+	
$t = 110$ s	29.9—151	23.1–85.3	-22.7/-43.4
$t = 150$ s	90.6—318	79.7–195	-12.0/-38.7
$t = 198$ s	118–254	117–197	-0.85/-22.4
$t = 250$ s	16.4–65.1	14.5–48.8	-11.6/-25.0

Figure 3.28 displays the total Sobol' indices for the radiative-heating distribution at the $t = 198$ s and $t = 250$ s trajectory points. The differences in the total Sobol' index profiles produced with and without H_2^+ are more pronounced for these two conditions, so the total Sobol' index profiles for the other two trajectory points are omitted for brevity. Starting with the $t = 198$ s condition, a decrease in the total Sobol' index of the rate coefficient for dissociation of H_2 with

partner H, and an increase in the total Sobol' index of the rate coefficient for the dissociation of H_2 with partner He is observed. The inclusion of H_2^+ further reduces the amount of H_2 and H just behind the shock, so the next most prominent species with which H_2 can collide and dissociate is He. Along the frustum and shoulder, the lowering of the H_2-H dissociation total Sobol' index is more pronounced. In turn, there is a marked increase in the total Sobol' indices for the hydrogen EII rate coefficient and freestream CH_4 concentration. As described previously, the inclusion of H_2^+ results in an earlier initiation of EII reactions. It follows that the additional/faster ionization that occurs in the stagnation region will have a rippling effect on the species number densities and temperatures downstream. Similarly, if the carbon EII can begin sooner because of the increased electron number density produced by H_2^+ reactions, the amount of carbon ions will be more strongly influenced by the freestream CH_4 concentration, and the radiation calculation will be affected accordingly.

The increase in the total Sobol' index for the hydrogen EII rate coefficient is also observed for the $t = 250$ s case, but the increase in the freestream CH_4 total Sobol' index along the frustum is not. Due to the lower velocity, the level of ionization predicted by the 12-species rate mechanism is relatively low, and no sensitivity of the radiative heating to ionization reactions is observed at this condition when H_2^+ is excluded. The faster AI and charge exchange reactions introduced with the 13-species rate mechanism therefore have a more pronounced effect on the ionization levels for the lower-speed conditions. Furthermore, the decrease in the sensitivity of the radiative heating to the rate coefficient for H_2 dissociation is more drastic than for the other trajectory points. Without H_2^+, variation in the rate coefficient for H_2 dissociation accounts for 80-95% of the uncertainty along most of the surface. Figure 3.28d, on the other hand, indicates that the inclusion of H_2^+ drops the contribution of H_2 dissociation rate coefficient to the radiative-heating uncertainty by a factor of $\sim$2 near the shoulder. As a result, the sensitivity to the hydrogen ionization and ionization/charge exchange reactions involving H_2^+ is increased. Overall, the inclusion of H_2^+ amplifies any effects of ionization on the radiative-heating uncertainty. There is also a small ($\sim$ 3%) contribution of the CH_4 dissociation reaction with partner H_2^+ to the radiative-heating uncertainty. This reaction produces additional H, which is the second most contributing species to the radiative heating along

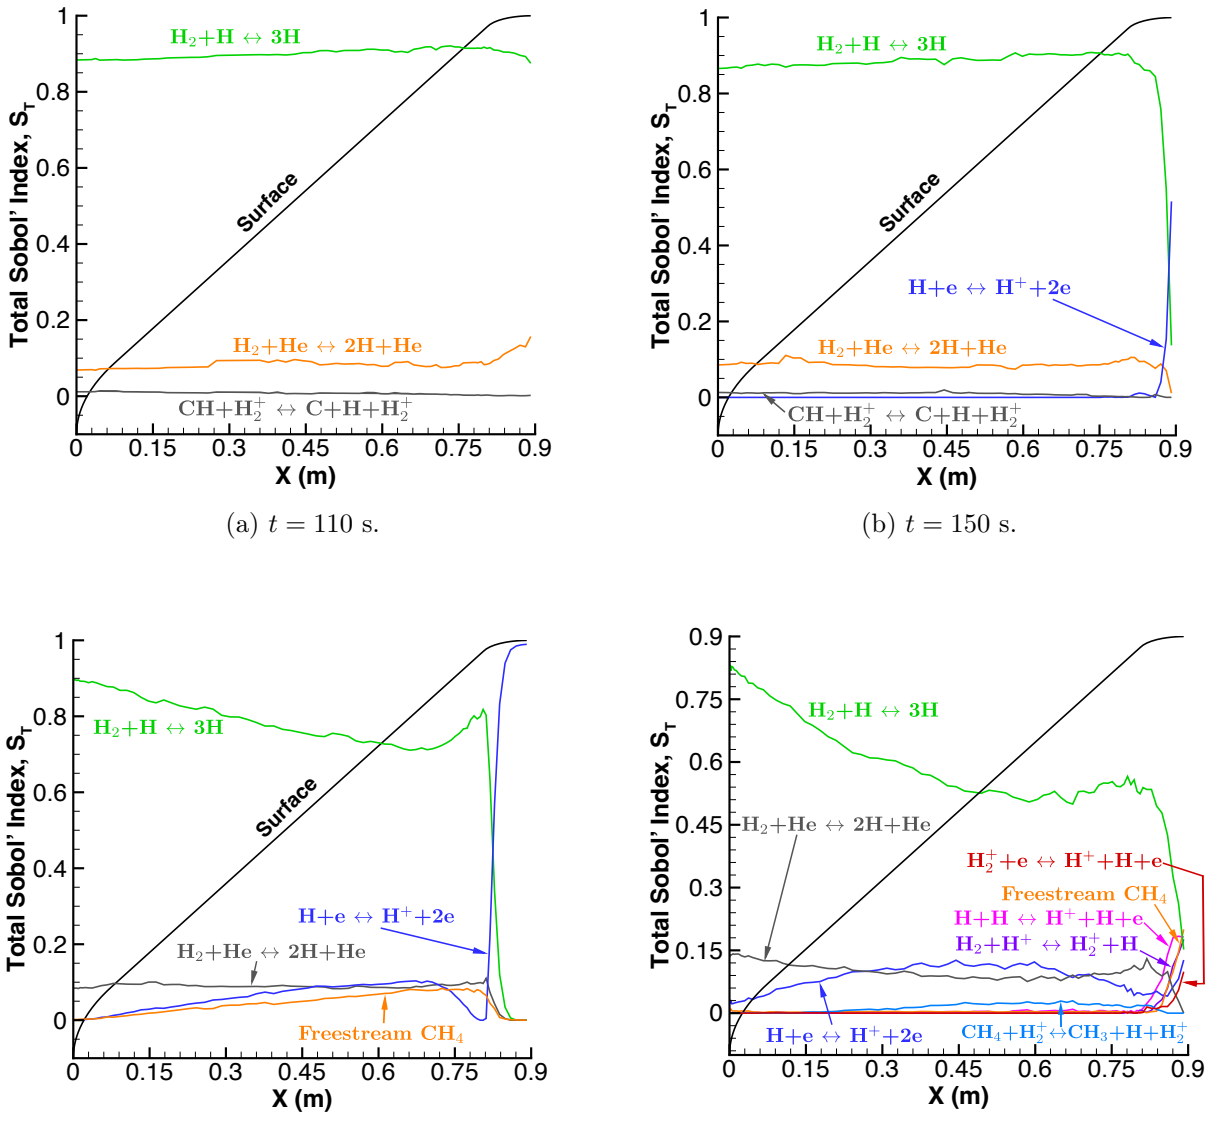

(a) $t = 110$ s.

(b) $t = 150$ s.

(c) $t = 198$ s.

(d) $t = 250$ s.

Figure 3.28: Total Sobol' indices for the radiative-heating distribution including $H_2{}^+$.

the frustum at this condition. The total Sobol' index for this reaction with partner H_2^+ (instead of a neutral collider) is a product of the uncertainty intervals used for the hydrogen AI reaction, as they produce a highly uncertain H_2^+ population.

For the $t = 110$ s and $t = 150$ s trajectory points, most of the radiative-heating uncertainty is still attributed to the rate coefficient for H_2 dissociation with partner H. Nearly all of the radiation produced at these two conditions comes from H_2, so it is expected that the primary source of radiative-heating uncertainty is still related to H_2 dissociation. Similar to the lower-altitude conditions, variability in the rate coefficient for the dissociation of H_2 with partner He (not H_2) contributes to approximately 8-10% of the total uncertainty. Furthermore, a small sensitivity ($\sim$1-2%) is observed for the rate coefficient for the reaction $CH + H_2^+ \leftrightarrow C + H + H_2^+$. This reaction produces both C and H, which contribute more to the radiative heating than the reactant CH. A specific sensitivity to this dissociation reaction with partner H_2^+, instead of a neutral collider, appears for the same reason as the lower-altitude trajectory points.

3.7 Summary

This Chapter conducted an uncertainty quantification (UQ) and global sensitivity analysis (GSA) of the aerothermodynamic environment for an entry capsule performing aerocapture at the Ice Giants. Most existing work on the modeling of Ice Giant aerothermodynamics consider a freestream composed only of hydrogen and helium. However, there is also speculation about the effect of freestream methane (CH_4) on the aerothermodynamic environment, particularly the radiative heating. To investigate these effects, this work considered uncertainties in chemical reaction-rate coefficients, transport property fitting parameters, vibrational-relaxation rates, and the freestream CH_4 concentration. A non-intrusive, point-collocated polynomial chaos expansion (NIPC-PCE) method was used to compute Sobol' indices and uncertainty bounds for convective- and radiative-heating distributions, as well as the axial-force coefficient. The least-angle regression algorithm was applied to promote sparsity in the polynomial expansion coefficient vector when solving the NIPC-PCE problem.

The convective heating was high at all four trajectory points. For the $t = 110$, 150, 198, and 250 s trajectory points, the stagnation point convective heat flux was nominally predicted to be 102, 257, 449, and, 362 W/cm^2, respectively. The shoulder convective heat flux was nominally predicted to be 58.4, 168, 308, and 241 W/cm^2, respectively. For all four conditions, the uncertainty in the convective-heating distributions was relatively small, ranging from 4.10–11.7%. At the stagnation point, the primary sources of convective-heating uncertainty were determined to be the variation in collision integrals and hydrogen electron-impact-ionization (EII) rate coefficient. The $t = 150$ s condition exhibited the highest level of sensitivity to the EII rate coefficient at the stagnation point because of its large effect on the shock standoff distance. For the $t = 150$, 198 and 250 s conditions, the convective-heating uncertainty along the frustum of the vehicle was primarily attributed to collision integrals, due to the surface being modeled as fully catalytic. At the $t = 110$ s condition, the frustum heating was largely sensitive to the hydrogen EII rate coefficient, because the ionization of hydrogen was delayed to downstream of the stagnation region. For all four conditions, convective heating at the shoulder was highly sensitive to hydrogen ionization rate coefficients, particularly in the reverse direction, suggesting that recombination phenomena largely governed the shoulder environments. In contrast, convective heating showed no sensitivity to the freestream CH_4 concentration.

Radiative-heating environments were substantial for freestream velocities in excess of 25 km/s. For the $t = 110$, 150, 198, and 250 s trajectory points, the stagnation point radiative heat flux was nominally predicted to be 68.7, 168, 163, and 32.0 W/cm^2, respectively. The peak shoulder radiative heat flux was nominally predicted to be 24.6, 78.0, 93.9, and 9.92 W/cm^2, respectively. For all four conditions, the uncertainty in the radiative heating was high, with epistemic uncertainty bounds exceeding 160%. At the $t = 198$ and 250 s condition, the primary source of radiative-heating uncertainty at the stagnation point and along the frustum was determined to be the rate of H_2 dissociation with partner H. Near the shoulder, there was a slight sensitivity of the radiative heating to the freestream CH_4 concentration, due to ionization of carbon being delayed to downstream of the stagnation point. Spectral features of atomic carbon were present, but they did not contribute

significantly to the total radiative heating. For the $t = 198$ s trajectory point, the primary source of radiative-heating uncertainty at the shoulder was the hydrogen EII reaction rate coefficient, where a decrease in the rate of this reaction increased the radiation. This was attributed to the increased population of the second electronic state of hydrogen H($n = 2$) which produced significant differences in the radiation from the H Lyman-α line. Similarly, at the $t = 250$ s trajectory point, the ionization of atomic hydrogen via heavy-particle impact was the largest source of uncertainty for the shoulder radiation, and a minor contribution from the freestream CH_4 was observed. At the $t = 110$ and 150 s conditions, the radiative-heating uncertainty was dominated by the rate coefficient for the dissociation of H_2 with partner H, due to the decreased freestream density and lower ionization levels. No sensitivity of the radiative heating to the freestream CH_4 concentration was found at the $t = 110$ or 150 s conditions. Overall, the high levels of uncertainty observed in the radiative heating environments would require substantial margin in the TPS design, which motivates further investigation, both experimentally and computationally, into the rate of dissociation of molecular hydrogen at aerocapture-relevant conditions.

The axial-force coefficient exhibited low levels of uncertainty, with a maximum value of a few percent. The largest sources of uncertainty in C_x were determined to be the rate coefficients for the ionization reactions $\mathrm{H} + \mathrm{e}^- \leftrightarrow \mathrm{H}^+ + 2\,\mathrm{e}^-$, and $\mathrm{H} + \mathrm{H} \leftrightarrow \mathrm{H}^+ + \mathrm{H} + \mathrm{e}^-$, which dominated at lower-altitude trajectory conditions. The rate coefficient for molecular hydrogen dissociation with partner atomic hydrogen was also found to affect the axial-force coefficient at the highest-altitude trajectory point. These three reactions drove the shape of the shock-wave around the vehicle and, therefore, aerodynamic forces.

Finally, the effects of including $H_2{}^+$ in the thermochemical model were investigated. The GSA and UQ were repeated, and the results were compared to those obtained without $H_2{}^+$. Inclusion of $H_2{}^+$ had little effect on the convective-heating distributions in both the nominal and uncertain scenarios, although, some changes to key sensitivities driving the convective heating were observed. Specifically, the lowest-density trajectory point demonstrated increased sensitivity to the additional ionization reactions introduced by $H_2{}^+$. While the sensitivities for the convective heating changed,

the overall level of convective-heating uncertainty only varied by a few percent at most when compared to the results produced with no H_2^+. The radiative-heating distributions and sensitivities, on the other hand, changed drastically with the inclusion of H_2^+. A decrease in the nominal radiative heating and the 95% confidence intervals for the radiative-heating distribution was observed for all four trajectory points, in some cases by up to 40%. The inclusion of H_2^+ produced a decreased sensitivity of the radiative heating to the rate coefficient for H_2 dissociation and an increased sensitivity to the H EII rate coefficient and freestream CH_4 concentration. As such, it is recommended that the rate at which reactions involving H_2^+ proceed at conditions relevant to Ice Giant aerocapture be studied more thoroughly.

Chapter 4

Probabilistic Reduced-Order Aerothermodynamic Modeling

"Everything should be made as simple as possible, but not simpler."
— Albert Einstein

4.1 Introduction

This chapter presents a method for developing efficient Gaussian Process Regression (GPR) surrogate models for hypersonic aerothermodynamics and aerodynamics, with a focus on approximating the high-fidelity uncertainty quantification data from Chapter 3 along an entire aerocapture trajectory. Three models are constructed to predict key quantities, including stagnation-point heat flux, aerodynamic coefficients, and full surface heat-flux distributions, demonstrating strong predictive accuracy and generalization. A multi-fidelity GP framework is introduced, leveraging data from both low- and high-fidelity aerodynamic models to improve efficiency and reduce computational cost while preserving model predictive capabilities. Additionally, constrained GPR techniques are explored to enforce physical properties like symmetry and monotonicity, further enhancing model reliability. Finally, the surrogate models are integrated into a trajectory simulation, illustrating the advantages of probabilistic aerodatabases over traditional lookup tables for quantifying the hysteresis effects of aerothermodynamic uncertainty.

4.2 Background

Generation of high-fidelity aerothermodynamic and aerodynamic databases (referred to as aerodatabases for brevity) is an integral part of the design of aerocapture systems. They are used to inform the design of almost all vehicle subsystems such as guidance, navigation and control, thermal protection system, and structural design. An aerothermodynamic database, consisting of both convective and radiative heating estimates, is used to design and size the thermal protection system (TPS) that protects the vehicle from the severe aerodynamic heating experienced during hypersonic flight [122, 123]. Obtaining an accurate prediction, and the associated uncertainty in that prediction, allows for a more robust TPS design and higher confidence levels for mission success. Similarly, a high-fidelity aerodynamic database is required for various elements of the mission design phase including six degree-of-freedom trajectory simulation [128], stability analysis, and aerodynamic loading predictions.

A standard aerodatabase consists of two components: a nominal value for some quantity of interest (QoI), and an associated uncertainty about the nominal value [129]. The level of uncertainty in a given aerothermodynamic QoI is essential for assessing the robustness of the aerocapture system to modeling uncertainties. In Chapters 6 and 7, extensive Monte Carlo campaigns are conducted to quantify the robustness of two novel aerocapture guidance algorithms. These Monte Carlo analyses require estimates of the uncertainty in QoIs, and a method of generating dispersed quantities of interest that are consistent with the underlying model/process. Furthermore, the number of entries in an aerodatabase can be quite large, making their construction a time-consuming process. For example, in addition to an extensive amount of experimental data, the Orion Crew Vehicle hypersonic aerodynamics database required over 1000 high-fidelity computational fluid dynamics simulations [130]. Assembling this large quantity of data into a coherent and useful format is nontrivial, and the complexity only grows as the geometry of the vehicle changes throughout the mission design cycle. Assignment of uncertainty levels consistent with the underlying data is crucial, yet existing approaches to impose bounds and distributions of QoIs use more ambiguous *ad hoc* approaches

[129, 130]. Therefore, there is a definite need for a more robust and efficient method to generate aerodatabases for nominal and off-nominal scenarios.

Recent studies have focused on improving the development of aerodatabases, with particular emphasis on statistical and probabilistic approaches rather than traditional methods relying on complex lookup tables [131, 132, 133, 134, 135, 136]. Scoggins et al. [137] investigated the use of a multihierarchy Gaussian process (GP) model to efficiently fuse aerodynamic data from multiple sources, producing a surrogate capable of predicting both mean and dispersed aerodynamic coefficients for the Orion Crew Vehicle. Using the same Orion dataset, Nakamura-Zimmerer et al. [138] developed structured covariance neural networks to parameterize the mean and covariance functions of a GP, demonstrating improved scalability with larger training datasets compared to traditional GP regression. Later work from Scoggins et al. employed GP-based surrogate models [122], and Bayesian regression [139] to develop aerodatabases for a NASA Ice Giant aerocapture concept. These probabilistic techniques have proved to be an effective means of integrating data from multiple sources into a single model capable of generating statistically consistent aerodatabases. In turn, the explicit probabilistic treatment of the aerodatabase lends itself to more informative conclusions, such as uncertainty in the entry vehicle trim conditions [137], when compared to heritage methods.

This chapter builds on these probabilistic methods to create surrogate models for the hypersonic aerodynamics and aerothermodynamics and an aerocapture vehicle [140]. In addition to the canonical GP-techniques presented in Chapter 2, it utilizes both multi-output and multi-fidelity Gaussian process techniques to lessen the need for extensive high-fidelity computational fluid dynamics simulations across the full trajectory range. Multi-output models enhance prediction accuracy by identifying correlations among different variables of interest [141], while multi-fidelity strategies decrease the quantity of high-fidelity data points necessary to build an adequately precise surrogate model [142, 143, 144, 145]. In particular, an efficient multi-output and autoregressive multi-fidelity approach [146, 147] is implemented to develop surrogate models for key parameters such as surface heating and aerodynamic coefficients. These models enable the rapid creation of probabilistic aerodatabases, delivering prediction uncertainties that align with those of the underlying training data.

Although a Neptune aerocapture mission serves as the example, this framework is adaptable to any hypersonic flight system.

4.3 Beyond Vanilla Gaussian Process Regression

The fundamental Gaussian Process Regression (GPR) theory presented in Chapter 2 provides a powerful foundation for constructing surrogate models, offering flexibility, incorporation of uncertainty into the predictions, and the ability to interpolate sparse data. However, for advanced aerodatabase design, standard GPR techniques alone are insufficient. This section details the advanced GPR techniques employed throughout this chapter and dissertation.

4.3.1 Multi-Output Gaussian process Regression

The quantities of interest in this work (aerodynamic coefficients, heating rates, etc.) all have some level of underlying physical correlation. Therefore, it may be advantageous to design a model that attempts to learn the relationship between multiple outputs at once. In the case of vector-valued predictions, the unknown mapping becomes a multidimensional function $\boldsymbol{f} : \mathbb{R}^D \mapsto \mathbb{R}^M$ with $M > 1$. The GP correspondingly has a mean function $\boldsymbol{\mu} : \mathbb{R}^D \mapsto \mathbb{R}^M$, and the covariance matrix becomes a block-structured matrix $\boldsymbol{K} \in \mathbb{S}_+^{MN \times MN}$, where M is the dimension of each output/observation and N is the number of training samples. This framework is often referred to as "multi-task" Gaussian process regression, where the goal is to predict M different functions simultaneously. For multi-task prediction, Equation 2.60 is written as [141]

$$\boldsymbol{f}(\boldsymbol{x}) \sim \mathcal{GP}\left(\boldsymbol{\mu}(\boldsymbol{x}), \boldsymbol{K}(\boldsymbol{x}, \boldsymbol{x}')\right) \tag{4.1}$$

where the covariance matrix is more explicitly written as

$$\boldsymbol{K}(\boldsymbol{X}, \boldsymbol{X}) = \begin{bmatrix} (K(\boldsymbol{X}_1, \boldsymbol{X}_1))_{1,1} & \dots & (K(\boldsymbol{X}_1, \boldsymbol{X}_2))_{1,M} \\ (K(\boldsymbol{X}_2, \boldsymbol{X}_1))_{2,1} & \dots & (K(\boldsymbol{X}_2, \boldsymbol{X}_M))_{2,M} \\ \vdots & \dots & \vdots \\ (K(\boldsymbol{X}_M, \boldsymbol{X}_1))_{M,1} & \dots & (K(\boldsymbol{X}_M, \boldsymbol{X}_M))_{M,M} \end{bmatrix}_{MN \times MN} . \tag{4.2}$$

Here, $\boldsymbol{X}_m \in \mathbb{R}^N$ for $m = 1, \ldots, M$ is the set of inputs corresponding to the m^{th} output of $\boldsymbol{f}$ and $(K(\cdot,\cdot))_{m,m'}$ is an $N \times N$ matrix where $m,\ m' \in [1, M]$. The posterior predictive mean and covariance are then obtained using the same analysis applied to the scalar-output problem.

One can potentially improve the posterior predictions in a multi-output setting if the dimensions of the output target are correlated by incorporating the correlation into the structure of the GP kernel. The simplest method of correlating the outputs of a multi-output GP is the intrinsic model of coregionalization (IMC) [148, 149, 141], which is sometimes also called the linear model of coregionalization (LMC). In IMC, the covariance of the GP outputs $f_m(x)$ for $i = 1, \ldots, M$ is defined via

$$cov(f_i(\boldsymbol{X}), f_j(\boldsymbol{X'})) = k(\boldsymbol{X}, \boldsymbol{X'}) \cdot B[i, j] \tag{4.3}$$

such that the covariance of the i^{th} function at $\boldsymbol{X}$ and the j^{th} function at $\boldsymbol{X'}$ is a kernel $k(\boldsymbol{X}, \boldsymbol{X'})$ times the i, j^{th} entry of a positive definite matrix $B \in \mathbb{R}^{M \times M}$. To ensure that the matrix B is positive definite, the model training procedure uses a separate matrix W such that

$$B = WW^T + \text{diag}(\kappa) \tag{4.4}$$

where $\kappa \ll 1$ is added for numerical stability and additional positive definiteness guarantees.

4.3.2 Multi-Fidelity Gaussian process Regression

Obtaining sufficient aerodynamic/aerothermodynamic data from a high-fidelity source such as computational fluid dynamics (CFD) is a time-consuming process. In many cases, numerous methods, with varying levels of fidelity, may be available to predict the same quantity of interest. For example, a low-fidelity estimate of the aerodynamic coefficients may rapidly be attained using modified-Newtonian theory, while a higher-fidelity estimate would require more computationally expensive three-dimensional CFD simulations. In this situation, the lower-fidelity, cheaper to obtain data can be used to supplement the higher-fidelity data in a multi-fidelity modeling framework. The denser, low-fidelity data is used to learn the general trends in the quantities of interest with respect

to the input variables (e.g. altitude, velocity), and the sparser, higher-fidelity data is used to anchor these trends when possible.

One of the most common approaches is the linear Auto-Regressive (AR1) method of Kennedy and O'Hagan [146], and it has been used extensively in the engineering world [131, 132, 133, 150, 151, 152, 144, 137]. Suppose there are s levels of information sources that each produce an output $y_t(\boldsymbol{x}_t)$ for $t = 1, \ldots, s$. If the input and output of each data set is organized in pairs by increasing fidelity, then $\boldsymbol{y}_1$ represents the cheapest and least accurate data available, and $\boldsymbol{y}_s$ represents the most expensive and most accurate data available. The AR1 method defines the autoregressive scheme

$$f_t(\boldsymbol{x}) = \rho_{t-1} f_{t-1}(\boldsymbol{x}) + \delta_t(\boldsymbol{x}) \tag{4.5}$$

where $f_t(\boldsymbol{x})$ and $f_{t-1}(\boldsymbol{x})$ are Gaussian processes priors defined for the data at fidelity level t and $t-1$, respectively. In this setting, $\rho \in \mathbb{R}$ is a constant that quantifies the level of correlation between the outputs of each model, and $\delta_t(\boldsymbol{x})$ is a Gaussian process that is assumed to be independent from the GP used to model f_t. Training a model of the form in Equation 4.5 requires the GP's describing each fidelity level to be trained simultaneously. In this setting, the inversion of the GP kernel matrix is required to minimize Equation 2.69, which may become computationally prohibitive if the training data set is excessively large. To reduce this computational cost, Le Gratiet [147] proposed a recursive formulation of the AR1 model in which the GP prior in Equation 4.5 is replaced with the GP posterior of the previous inference level

$$f_t(\boldsymbol{x}) = \rho_{t-1} \hat{f}_{t-1}(\boldsymbol{x}) + \delta_t(\boldsymbol{x}) \tag{4.6}$$

where the $\hat{(\cdot)}$ notation is used to represent the posterior of the GP. This modification allows each fidelity level t to be trained independently, greatly reducing the computational cost. For example, in the case of two levels of fidelity, the first level is trained directly on the lowest-fidelity data set, and the subsequent GP is trained on the difference between the training data at that level and the corresponding subset of observations from the lower-fidelity data set. Note that this assumes that the training inputs for each level of fidelity $\boldsymbol{X}_t$ are a subset of the previous level's training data set

such that $\boldsymbol{X}_t \subseteq \boldsymbol{X}_{t-1} \subseteq \ldots \subseteq \boldsymbol{X}_1$. The posterior distribution at any test point $\boldsymbol{x}$ for each level is then obtained as

$$\hat{f}_t(\boldsymbol{x}) = \mathcal{GP}(\hat{\mu}_t, \hat{K}_t) \tag{4.7}$$

$$\hat{\mu}_t(\boldsymbol{x}) = \rho_{t-1}\hat{\mu}_{t-1}(\boldsymbol{x}) + K_t(\boldsymbol{x}, \boldsymbol{X}_t)\left[K_t(\boldsymbol{X}_t, \boldsymbol{X}_t) + \sigma_{n,t}^2 I\right]^{-1}\left(\boldsymbol{Y}_t - \rho_{t-1}\hat{\mu}_{t-1}(\boldsymbol{X}_t)\right) \tag{4.8}$$

$$\begin{aligned}\hat{K}_t(\boldsymbol{x}, \boldsymbol{x}') = {} & \rho_{t-1}^2\hat{K}_{t-1}(\boldsymbol{x}, \boldsymbol{x}') + K_t(\boldsymbol{x}, \boldsymbol{x}') \\ & - K_t(\boldsymbol{x}, \boldsymbol{X}_t)\left[K_t(\boldsymbol{X}_t, \boldsymbol{X}_t) + \sigma_{n,t}^2 I\right]^{-1} K_t(\boldsymbol{X}_t, \boldsymbol{x}')\end{aligned} \tag{4.9}$$

where the correlation coefficient is computed via [137]

$$\rho_{t-1} = \frac{\boldsymbol{Y}_t^T \boldsymbol{Y}_t}{\boldsymbol{Y}_t^T \boldsymbol{Y}_{t-1}(\boldsymbol{X}_t)} \quad . \tag{4.10}$$

It is noted that a more general, non-linear model, introduced by Perdikaris [153] et al. in 2017, called the non-linear autoregressive Gaussian process (NARGP), has been shown to capture non-linear correlations between each fidelity level more accurately than AR1 [142]. However, this work demonstrates that the AR1 model is sufficiently accurate for predicting the desired quantities of interest, so the additional complexity of the NARGP model is not required.

4.4 Application to Neptune Aerocapture

This section presents an application of the GPR framework described in Chapter 2 and Section 4.3 to develop surrogate models to predict aerothermodynamic and aerodynamic quantities of interest for an aerocapture mission at Neptune. The atmospheric entry state is derived from Lockwood et al. [13] that defines a retrograde entry with an inertial entry velocity of $V_I = 29$ km/s and initial altitude of $h_0 = 1000$ km. For consistency with the analysis in Chapter 3, the entry vehicle is modeled as a 5 m diameter, 70-degree sphere cone with a nose radius of 1.25 meters.

For all CFD simulations, the chemical kinetics model of Leibowitz [26] and Leibowitz and Kuo [30] are used, and the vehicle surface is assumed to be fully catalytic and in radiative equilibrium with a surface emissivity of 0.9. For all GP models, the input space includes altitude-velocity pairs such that $\boldsymbol{x}_i = [h_i,\ V_i]$. To bound the values of h_i and V_i at which to collect heating and aerodynamic

data, the theoretical entry corridor for the entry vehicle is determined to produce the bounding lift-up (bank angle $\sigma_b = 0$ deg.) and lift-down (bank angle $\sigma_b = 180$ deg.) trajectories as shown in Figure 4.1a. These trajectories produce velocity bounds of $V_i \in [15,\ 30]$ km/s and altitude bounds of $h_i \in [190,\ 480]$ km, where the upper bound on the altitude is chosen such that the Knudsen number is approximately 0.01, when considering a length scale of 5 m, so that the continuum assumption is valid [154]. LeMANS is used to simulate 50 altitude-velocity pairs that are shown as the closed symbols in Figure 4.1a, which are generated from a Latin Hypercube sampling algorithm using a "minimax" criterion [155]. Each simulation performs an axisymmetric computation at zero angle of attack, and the resulting surface heat-flux and pressure/shear distributions are used to compile the training and test data sets for the various GP surrogate models. The training dataset is also augmented with the four data points used in the detailed UQ analysis in Chapter 3. An example of flow field pressure contours computed by LeMANS for a freestream altitude and velocity of 193 km and 25.8 km/s, respectively, is shown in Figure 4.1b.

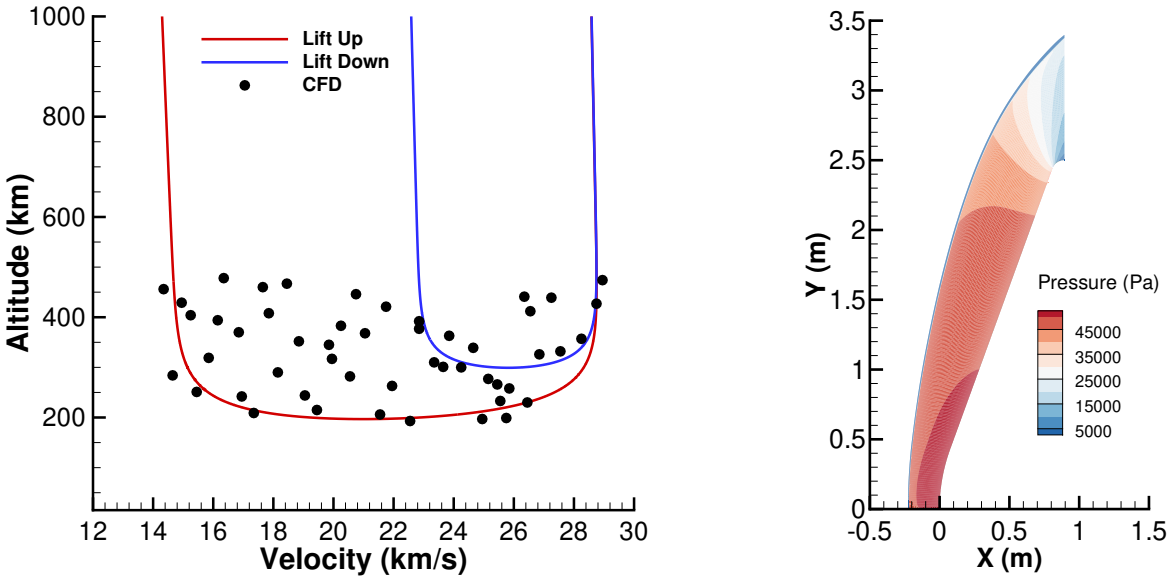

(a) Database points.

(b) Pressure contours.

Figure 4.1: Bounding lift-up and lift-down trajectories with CFD database points (a) and example pressure contours computed by LeMANS (b).

4.4.1 Simultaneous Prediction of Stagnation-Point Heating and Axial-Force Coefficient

The first task of this work is to generate a single model capable of predicting both the convective heating and aerodynamic coefficients of the entry vehicle. To do so, the outputs of the GP are correlated using the LMC described in Section 4.3.1, and the output space is the log of the stagnation-point heat flux and axial force coefficient such that $\boldsymbol{y}_i = \left[\log_{10}(q_{\text{stag,i}}),\ C_{x,i}\right]$, where $i = 1, \ldots, N$. The log of the heat flux is chosen because it produces a more linear relationship between the model inputs and outputs and improves model convergence.

Gaussian processes are highly capable at modeling noisy/uncertain data. While training a GP typically involves optimizing the posterior likelihood variance σ_n^2, it is also possible to enforce a particular value of σ_n^2. This is typically done in scenarios where the noise level of the data is known *a priori*; for example, if a sensor has a known measurement noise, it is possible to infuse that uncertainty directly into the GP's predictions. Chapter 3 demonstrated that uncertainty in both the heating and aerodynamic exists due to parametric uncertainty in the CFD inputs. Therefore, it is desirable for the trained GP model to account for these uncertainties in its predictions. To first order, the computed uncertainty in Chapter 3 defines a known measurement noise σ_n^2 in the general regression problem in Equation 2.59, so the model no longer has to learn this value during training. Instead, the values of σ_n^2 for the stagnation-point heating and axial-force coefficient are set to be Gaussian random variables with a mean located at the posterior predictive mean of the GP and a standard deviation consistent with the uncertainties calculated in Chapter 3.

A 90/10 training-validation split is used to partition the available data, meaning 90% of the data points are used to train the model, and the remaining 10% are reserved for assessing the trained model's predictive capabilities. The input space is standardized by using z-score normalization (i.e. zero mean and unit variance) [156] to ensure that all input factors operate on the same scale. The model hyperparameters are optimized by minimizing Equation 2.69 using the Adam [103] optimizer with an initial learning rate of 0.1. Training is terminated once the loss averaged over some window

stops decreasing, within some small tolerance (e.g. $\delta = 1 \times 10^{-6}$), from one iteration to the next. The optimized hyperparameters for the multi-output GP are summarized in Table 4.1, which have been left in non-dimensional form as a result of the z-score normalization.

Table 4.1: Optimized hyperparameters for multi-output GP trained to predict stagnation-point heat flux and axial-force coefficient.

Hyperparameter	Variance, σ^2	Altitude Length Scale, l_h	Velocity Length Scale, l_V
Value	7.76×10^{-2}	5.258	5.478

It is possible to learn a significant amount from the optimized model hyperparameters. The variance σ^2 is an estimate of the uncertainty in the trained GP's predictions and is directly tied to the variance in the data used to train the model. The units of the length scale parameters correspond to the input dimension and indicate the distance in that dimension in which significant variations in the output variable are expected. Recall that, because of the z-score normalization, the input dimensions are technically unitless, but they can be put into dimensional units by inverting the z-score normalization procedure. By standardizing the training input space, a more consistent comparison between input dimensions may be made. Input dimensions with larger length scales signify that less variability is expected in the output, while smaller length scales indicate the function varies more rapidly over that dimension. Note that the optimized altitude and velocity length scales are of the same order of magnitude. This indicates that the GP outputs are expected to have a similar level of variation with changes in altitude and velocity. At first, such a trend is inconsistent with existing models for both stagnation-point heat flux and aerodynamic coefficients. Fay & Riddell [157], and later Sutton & Graves [158], derived expressions for the stagnation-point convective heating that show cubic dependence on the freestream velocity, but a square-root dependence on the freestream density (and therefore altitude). However, recall that the GP targets contained the log of the stagnation-point convective heat flux. In log-space, the Fay & Riddell/Sutton & Graves relationships between stagnation-point convective heating and altitude/velocity becomes linear, which is consistent with the results in Table 4.1. A similar conclusion may be made for the

axial force coefficient when considering Newtonian theory [159].

To certify that the trained model is capable of producing accurate predictions across the entire input space, it must be evaluated at altitude-velocity pairs that are included in the training data set and at test points that are not included in the training data set. Table 4.2 summarizes the trained model's ability to predict the stagnation-point convective heat flux and axial-force coefficient at the 45 altitude-velocity pairs in the training data set, and the remaining 5 pairs in the test data set. The net posterior-predictive relative error is defined as

$$\text{Error} = \left\| \frac{\hat{f}(\boldsymbol{X}) - \boldsymbol{Y}(\boldsymbol{X})}{\boldsymbol{Y}(\boldsymbol{X})} \right\| \tag{4.11}$$

where $\hat{f}(\boldsymbol{X})$ is the trained GP's mean prediction at the input $\boldsymbol{X}$, and $\boldsymbol{Y}(\boldsymbol{X})$ is the corresponding training/test data collected at the input $\boldsymbol{X}$. Here, the input $\boldsymbol{X}$ is either the training set or the validation set.

Table 4.2: Accuracy statistics for multi-output GP trained to predict stagnation-point heat flux and axial-force coefficient as computed by Equation 4.11.

	Error over training set (%)	Error over validation set (%)
q_{stag}	5.78	10.1
C_x	2.55	2.70

Here, a difference between the accuracy in the stagnation-point convective heat flux and axial-force coefficient is observed. While the model is able to predict the axial-force coefficient accurately on both the training and validation data sets, the accuracy in prediction is lower for the stagnation-point convective heat flux. This is explained by the fact that observation-noise values consistent with the uncertainties calculated in Chapter 3 were enforced on the heat-flux and axial-force coefficient data in order to account for the parametric uncertainty in the CFD results. Recall that the heat-flux uncertainties are substantially larger than the aerodynamic uncertainties, so, it is expected that the posterior mean of the GP does not perfectly reproduce the heat-flux data that was used to train the model. Instead, the posterior-predictive distribution should *contain* the training data.

Figures 4.2 and 4.3 show the trained GP's posterior predictive mean and 100 random samples

for the stagnation-point heating and axial-force coefficient, respectively, for altitudes in the range $h \in [150,\ 500]$ km at a velocity of $V = 27.0$ km/s. This velocity is intentionally chosen to be consistent with the velocity of the maximum-heating condition in Chapter 3. Note that the upper and lower bounds of this altitude range are outside of the altitude values used during training. The nominal stagnation-point convective heat flux and axial-force coefficient for the maximum-heating condition in Chapter 3 are plotted for reference, where the error bars represent the uncertainty from the PCE surrogate.

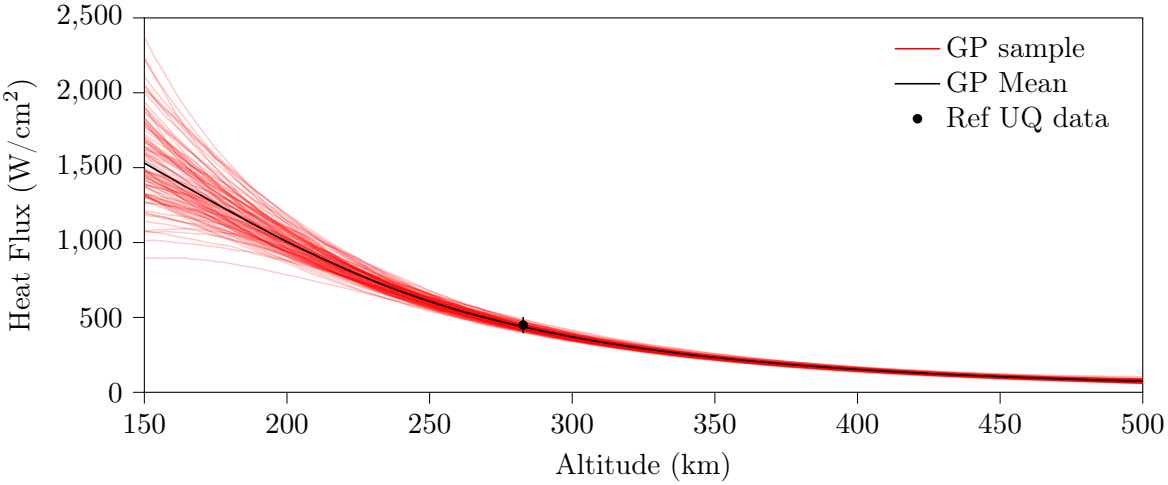

Figure 4.2: Posterior predictions of stagnation-point convective heating at $V = 27.0$ km/s.

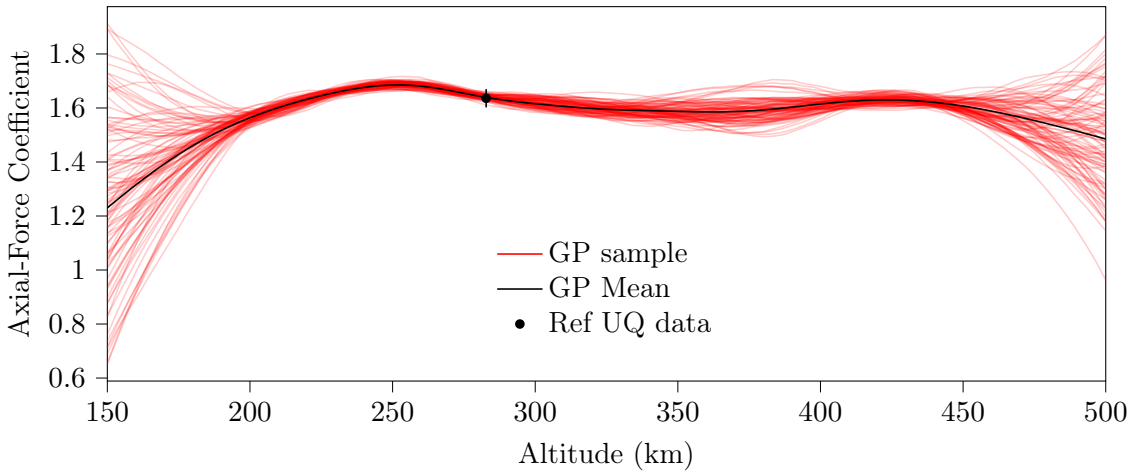

Figure 4.3: Posterior predictions of axial-force coefficient at $V = 27.0$ km/s.

The GP does an excellent job at predicting the trends in both the stagnation-point heat flux and the axial-force coefficient. Additionally, as shown in Figure 4.2, the GP surrogate is successfully able to incorporate the uncertainty in the stagnation-point heating computed by the PCE surrogate from Chapter 3 into its predictions across the entire altitude space. The model slightly overpredicts the stagnation-point convective heating at higher velocities, with a maximum relative error of 3.9% between the posterior-predictive mean and the CFD data. Again, this difference is expected due to the specified noise level in the CFD data, but to reduce the predictive error, it would be necessary to include additional training data in the high-velocity region of parameter space. This is straightforward to achieve, as the data currently allocated for model testing can be used to augment the current training data set, and additional training data can be generated to augment the original set. There is a noticeable increase in the variation of the GP's predictions for altitudes lower than 200 km, as these altitudes are lower than the lowest-altitude data point included in the training procedure. The increase in function variation away from training data is a well-known characteristic of Gaussian process regression and allows for the "out-of-distribution" (i.e. extrapolating instead of interpolating) uncertainties to be estimated. The predictions at altitudes where more training data is available are more accurate than that shown in Figure 4.2, as expected, with a maximum relative error of approximately 2.0%. The training and test data are well contained within the uncertainty level determined by the trained GP.

The model also does an excellent job at predicting the axial-force coefficient as shown in Figure 4.3. From Newtonian theory, aerodynamic coefficients for hypersonic blunt bodies should vary slightly with Mach number, so it is encouraging that the profile in Figure 4.3 is relatively flat. Even though predictions for a constant velocity of 27 km/s are shown, the Mach number changes with altitude due to the variation in the speed of sound. Due to the slightly sparser altitude training data at $V = 27.0$ km/s, the model shows larger predictive uncertainties at higher altitudes. These higher altitudes lie outside the theoretical entry corridor shown in Figure 4.1a, so the increased aerodynamic uncertainty in this region has no implication for trajectory/guidance design. An increase in predictive uncertainty away from training data is a known phenomenon of

GPR, and it indicates that including more data where the predictive uncertainty is higher would improve the model's overall predictive capabilities. There is both slight over- and under-prediction of the axial-force coefficient by the GP, but the maximum relative difference between the posterior predictive mean and the training and test data is only 1.09%. Similar to the stagnation-point heating results, the GP's predictions at conditions near more training data are more accurate, with errors on the order of 0.5%.

4.4.2 Prediction of Convective Heat-Flux Surface Distributions

It is desirable to predict the heat-flux across the surface of the vehicle, as entire surface-heating profiles are required during the design phase and TPS design [123]. Additionally, the ability to efficiently and accurately model the vehicle's aerothermodynamic environment, along with its associated uncertainty, could enable the development of more advanced guidance algorithms that utilize the full surface-heating distribution to inform the resulting guidance law. For blunt-body capsules, the heat flux at the shoulder can be comparable in magnitude to the stagnation-point heat flux. Expanding the aerothermodynamic model within a guidance algorithm beyond a simplified stagnation-point heating model would ensure that heating across the entire vehicle is accounted for during trajectory planning.

The input space is expanded to be $\boldsymbol{x}_i = [h_i,\ V_i, s_i]$, where the added variable s is a normalized arc length along the vehicle surface. This is chosen because an input of $s = 0$ produces a prediction at the stagnation point, and an input of $s = 1$ produces a prediction at the shoulder. The output space is taken to be the surface heat flux at the input arc-length value such that $y_i = q_i(s_i)$, where $i = 1, \ldots, N$. Instead of using the log of the heat flux, as was done in Section 4.3.1, both the input and output space are normalized using z-score normalization. This is chosen because the relationship between the freestream altitude/velocity and the entire surface-heating distribution is not as well-defined as the relationship between altitude/velocity and stagnation-point heating. A 90/10 training-validation split is used to partition the available data, such that 45 of the 50 CFD data sets are used to train the model, and the remaining 5 are reserved for model validation. Using

the raw CFD data would produce a training data set size of $N = 45 \times n_{\text{surface}}$, where n_{surface} is the number of wall-tangential points along the vehicle surface used in the CFD simulations. To prevent the training data set from becoming too large, the surface profiles are interpolated onto a linearly spaced grid of 100 points, such that $n_{\text{surface}} = 100$. The optimized hyperparameters for the variance and length scales corresponding to the altitude, velocity, and non-dimensional arc length are summarized in Table 4.3. Recall that these values are also non-dimensional because of the z-score normalization.

Table 4.3: Optimized hyperparameters for single-output GP trained to predict convective heat-flux distributions.

Hyperparameter	Variance, σ^2	Length Scale, l_h	Length Scale, l_V	Length Scale, l_s
Value	10.8	0.845	2.54	0.117

Compared to the hyperparameters obtained in Section 4.3.1, a different trend is observed in the values shown in Table 4.3. In this case, it is shown that $l_s \leq l_h \leq l_V$, which indicates that the GP outputs are most sensitive to changes in the non-dimensional arc length and are the least sensitive to changes in the velocity. It is expected that the heat-flux distribution will be most sensitive to the position along the body, as regions with sharp curvature (e.g. the shoulder) produce augmentations in surface heating. However, the results in Table 4.3 show that the heat-flux distribution varies more with changes in freestream altitude than changes in freestream velocity. A normalized heat flux, known as the Stanton number, may be introduced to remove the dependency of the surface heating on the freestream conditions

$$St = \frac{q}{0.5\rho V^3} \quad . \tag{4.12}$$

Figure 4.4a shows the Stanton number surface profiles for four difference freestream conditions. The blue lines indicate conditions that have similar freestream altitude (but different velocities), and the red lines indicate conditions that have similar freestream altitudes (but different velocities). Clearly, there is a much larger variation in the Stanton number when the freestream altitude is changed than there is when the freestream velocity changes. This is likely because the heat-flux

distribution is also a strong function of the local surface shear stress, which depends more on the freestream density (and therefore altitude) than it does the freestream velocity. Surface shear-stress profiles computed for each of the four cases are plotted in Figure 4.4b, confirming that a much larger difference in viscous effects are observed by varying the freestream altitude than by varying the freestream velocity.

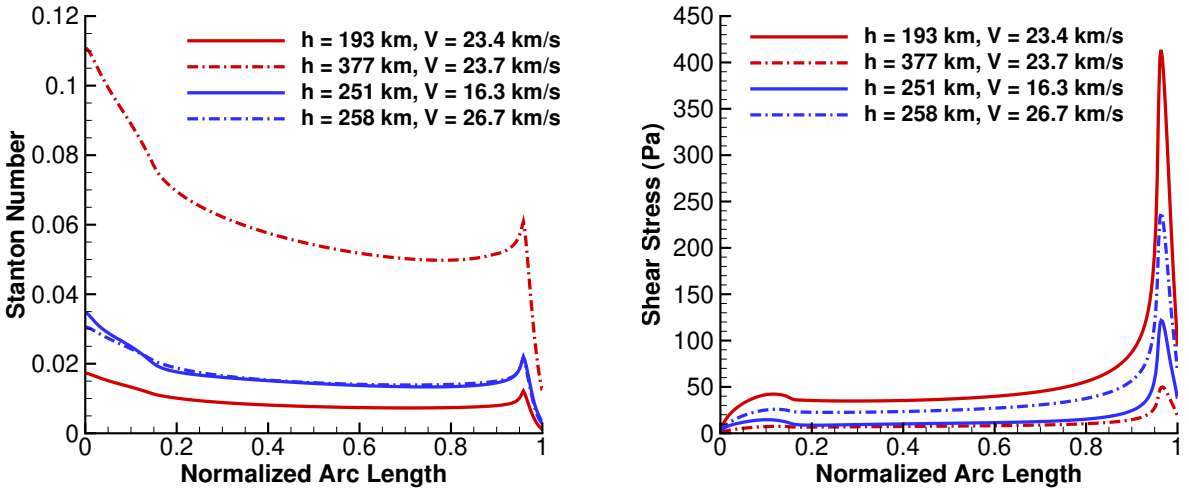

(a) Stanton number.

(b) Shear stress.

Figure 4.4: Bounding lift-up and lift-down trajectories with CFD database points (a) and example Pressure contours computed by LeMANS (b).

To demonstrate the trained GP's ability to rapidly predict entire convective heat-flux distributions, the model is evaluated at altitude-velocity conditions that are in both the training and validation data sets. Figure 4.5 illustrates the surface heat-flux distributions predicted by the GP for $(h, V)_{\text{training}} = (368 \text{ km}, \ 21.9 \text{ km/s})$, which is a condition included directly within the training data set. The solid blue lines show the results obtained from LeMANS, while the dashed red lines and shaded red area show the posterior predictive distribution from the trained GP. Excellent agreement between the posterior predictive mean and the data taken directly from LeMANS is observed, as expected. Similar predictions are made at an off-training point, $(h, V)_{\text{validation}} = (215 \text{ km}, \ 20.3 \text{ km/s})$, as shown in Figure 4.6. The uncertainty in the posterior predictions is larger for the off-training

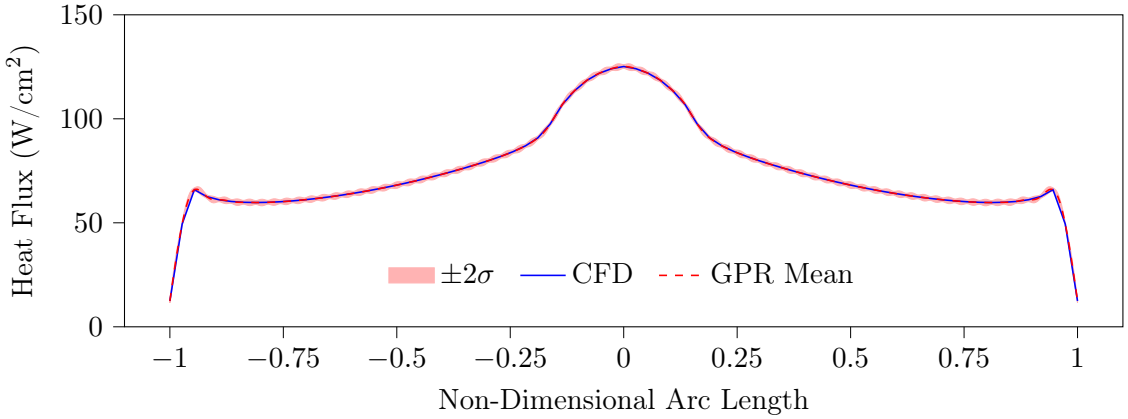

Figure 4.5: Posterior predictions of surface heating at $h = 368$ km and $V = 21.9$ km/s. These conditions are included in the model training data set.

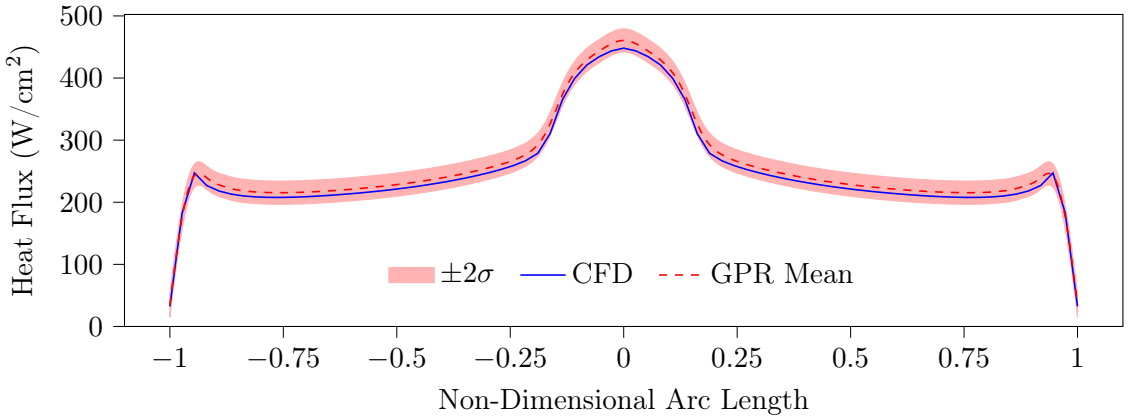

Figure 4.6: Posterior predictions of surface heating at $h = 215$ km and $V = 20.3$ km/s. These conditions are not included in the model training data set.

point, as expected, but the posterior predictive mean of the GP still shows excellent agreement with the data obtained directly from LeMANS. It is interesting to note that, near the shoulder where $|s| \approx 0.95$, there is a slight oscillation in the heating magnitudes predicted by the GP. The explanation for such behavior is that the optimizer is attempting to find a value for l_s that is able to capture both the sharp changes in heat flux near the shoulder and the more gradual changes in heat flux near the frustum and stagnation region.

4.4.3 Three-Dimensional Aerodynamic Model

In addition to altitude and velocity, a full aerodatabase must be able to account for variation in aerodynamic coefficients as a function of angle of attack. Consequently, it is necessary to develop a GP-based model that includes angle of attack as an input. To achieve this, 50 three-dimensional CFD simulations are conducted, where altitude and freestream velocity values correspond to the points shown in Figure 4.1a, while angle-of-attack values are selected using a Latin-hypercube sampling approach. The angle-of-attack bounds are chosen to be $\alpha \in [-20,\ 0]$ degrees, as blunt bodies produce positive lift at negative angles of attack. Additionally, because the entry vehicle is geometrically axisymmetric, its aerodynamic characteristics should be symmetric in angle of attack, such that $C_L(-\alpha) = -C_L(\alpha)$ and $C_D(-\alpha) = C_D(\alpha)$, so it is not necessary to collect data at positive angle-of-attack values.

The GP input space is defined to be $\boldsymbol{x}_i = [h_i,\ V_i, \alpha_i]$, where the added variable α is the vehicle angle of attack, and the output space is the lift and drag coefficients, such that $\boldsymbol{y_i} = [C_{L,i},\ C_{D,i}]$, where $i = 1, \ldots, N$. Instead of using a simple 90/10 training-validation split of the available data, the training data set is gradually increased in size. By gradually increasing the training data set size, a clearer measure of model convergence can be obtained, and it is of interest to see how the trained model's predictions change as more data is added to the training set. Specifically, models are generated with N = 10, 20, 30, 40, and 45 3D-CFD data points, where the training points are randomly selected from the entire data set. Note that, by taking this approach, at least 5 data points are reserved for model validation. When training each model, the feature vectors are normalized using z-score normalization, and the model outputs are correlated using the intrinsic model of coregionalization described in Section 4.3.1. The optimized hyperparameters for all five GPs are summarized in Table 4.4, where the presented values have been given dimensional units by inverting the z-score normalization process for a more intuitive interpretation.

The length scale corresponding to the angle-of-attack dimension is much smaller than the range of angle-of-attack values used during training, whereas the length scales for the altitude

Table 4.4: Optimized hyperparameters for multi-output GP trained to lift and drag coefficients as a function of altitude, velocity, and angle of attack.

N	Variance, σ^2	Length Scale, l_h	Length Scale, l_V	Length Scale, l_α
10	1.03×10^{-3}	605 km	62.8 km/s	3.50 deg
20	7.36×10^{-4}	653 km	37.0 km/s	6.40 deg
30	2.39×10^{-3}	632 km	31.4 km/s	10.0 deg
40	1.06×10^{-2}	672 km	31.4 km/s	11.7 deg
45	9.33×10^{-3}	645 km	31.4 km/s	11.8 deg

and velocity are larger than the range of altitude/velocity values used in training. This indicates that, for the variation in inputs of interest, more variation in aerodynamic coefficients with changes in angle of attack is expected. This is consistent with the underlying physics of the model, as aerodynamic coefficients are expected to be most sensitive to angle of attack. The model trained with only 10 data points has the smallest angle-of-attack length scale, which is likely because the 10 data points used displayed large amounts of variation over the angle-of-attack values considered. A similar conclusion may be made for the velocity length scale being significantly larger than models with more training data, where the data instead showed little variation over the velocity values considered. Additionally, it is interesting that the velocity length scale does not change for $N \geq 30$, indicating that the velocity dependence of the aerodynamic coefficients is more easily learned than others. Overall, each length shows a relative degree of convergence by increasing the training data set size, as increasing the data set size samples the input space more thoroughly.

Recall that the training data set contains angle-of-attack values in the range $\alpha_{\text{train}} \in [-20,\ 0]$ degrees. Because the entry vehicle is geometrically axisymmetric, its aerodynamic characteristics should be also symmetric in angle of attack, such that $C_L(-\alpha) = -C_L(\alpha)$ and $C_D(-\alpha) = C_D(\alpha)$. Using this information allows for the out-of-distribution accuracy to be assessed rapidly by querying the trained GP models for $\alpha_{\text{test}} > 0$. Figures 4.7–4.9 show the lift, drag, and lift-to-drag ratio response surfaces, respectively, predicted by the posterior mean of GPs trained with 10, 20, and 45 data points. All surfaces are generated for $V \in [15,\ 30]$ km/s and $\alpha \in [-20,\ 20]$ degrees at an altitude of $h = 300$ km.

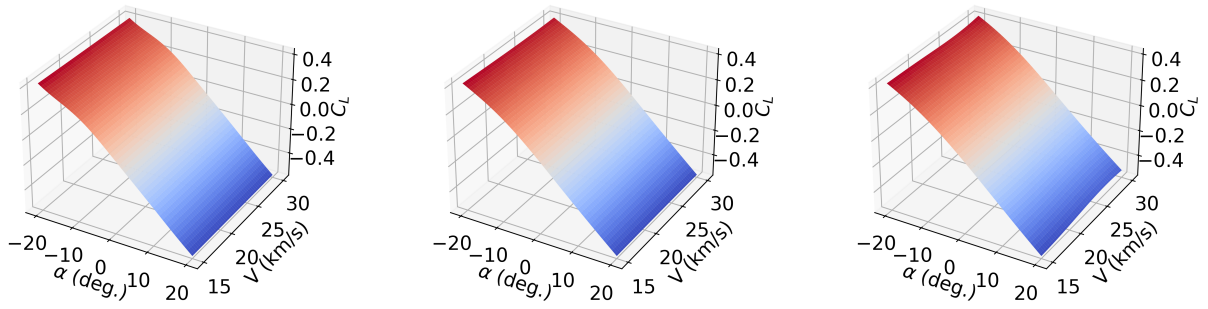

(a) 10 points. (b) 20 points. (c) 45 points.

Figure 4.7: Lift-coefficient response surfaces at $h = 300$ km generated from GP trained using 10 data points (a), 20 data points (b), and 45 data points (c).

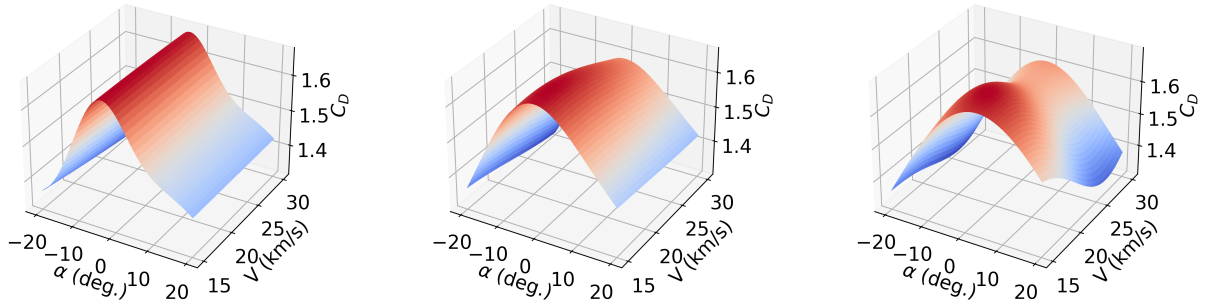

(a) 10 points. (b) 20 points. (c) 45 points.

Figure 4.8: Drag-coefficient response surfaces at $h = 300$ km generated from GP trained using 10 data points (a), 20 data points (b), and 45 data points (c).

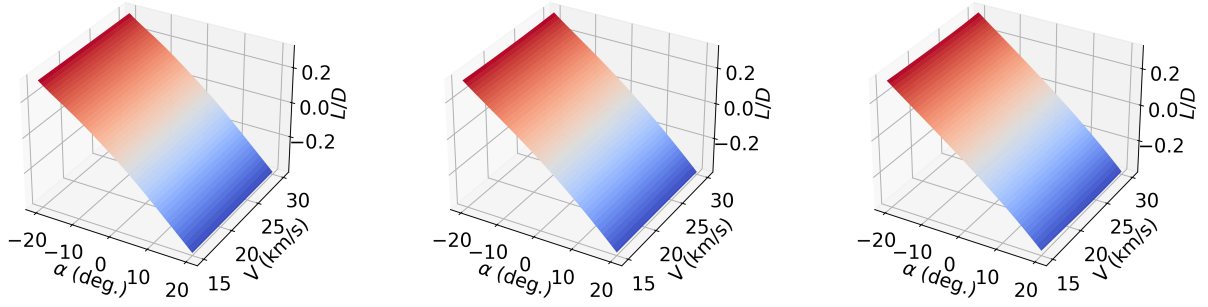

(a) 10 points. (b) 20 points. (c) 45 points.

Figure 4.9: Lift-to-drag ratio response surfaces at $h = 300$ km generated from GP trained using 10 data points (a), 20data points (b), and 45 data points (c).

The models trained with 10 data points produce some adherence to aerodynamic symmetry

for $\alpha > 0$, but only for the lift coefficient and lift-to-drag ratio. Generally, the lift coefficient is linear in angle of attack, and the drag coefficient is quadratic in angle of attack, so it is expected that the out-of-distribution predictions of the drag coefficient are less consistent than those for the lift coefficient and lift-to-drag ratio. However, as the training data set size increases, the response surfaces begin to show much more consistent aerodynamic symmetry in angle of attack. To more clearly assess the GP's out-of-distribution predictive capabilities, slices of the response surfaces in Figures 4.7–4.9 are shown in Figures 4.10–4.12 for a freestream velocity of $V = 30$ km/s.

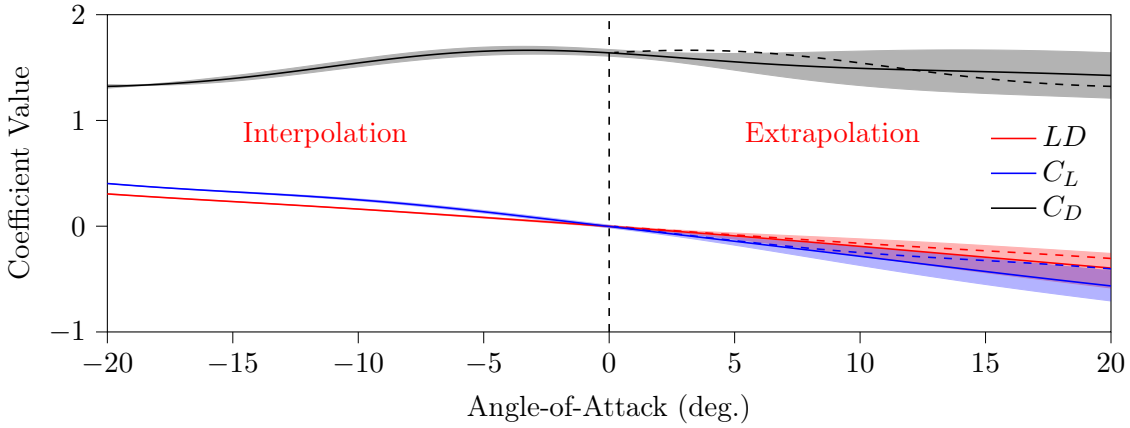

Figure 4.10: Slices of lift/drag coefficients and lift-to-drag ratio at $h = 300$ km and $V = 30$ km/s obtained from GP trained using 10 3D-CFD data points.

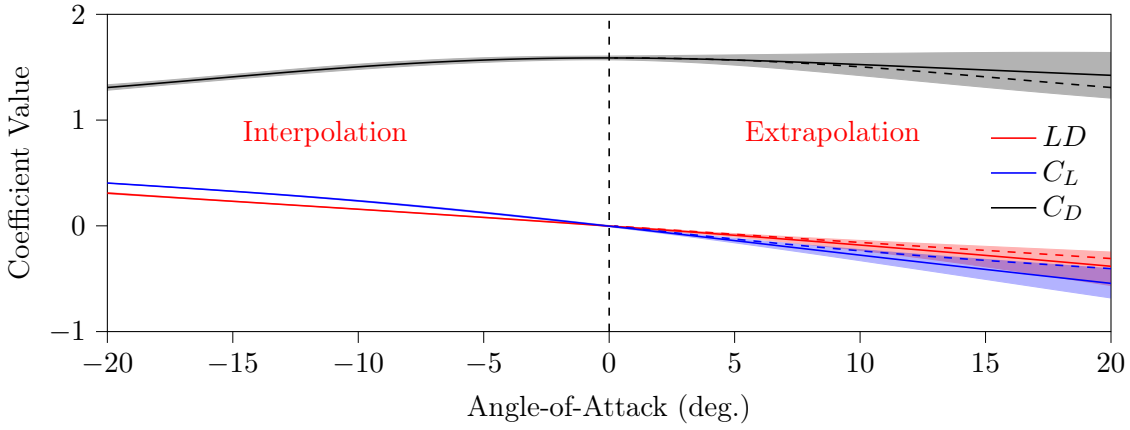

Figure 4.11: Slices of lift/drag coefficients and lift-to-drag ratio at $h = 300$ km and $V = 30$ km/s obtained from GP trained using 20 3D-CFD data points.

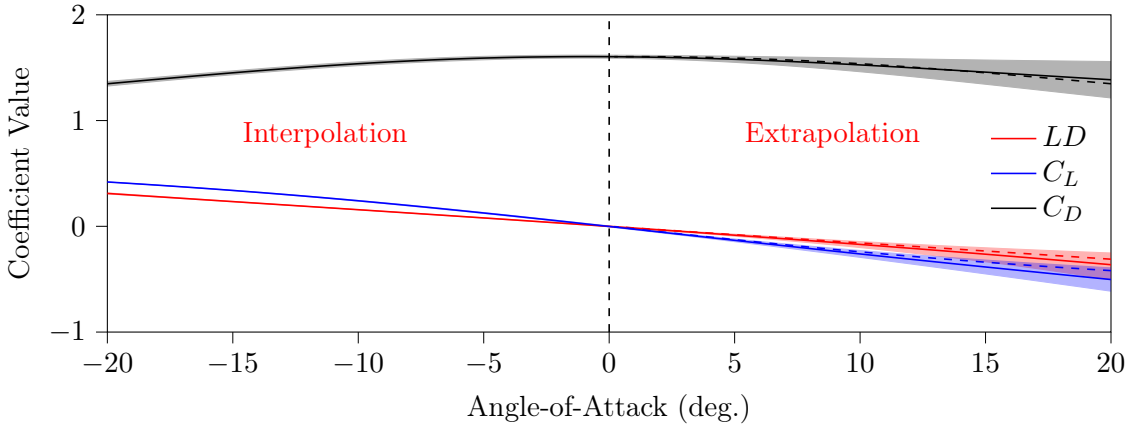

Figure 4.12: Slices of lift/drag coefficients and lift-to-drag ratio at $h = 300$ km and $V = 30$ km/s obtained from GP trained using 45 3D-CFD data points.

The solid lines in Figures 4.10–4.12 represent the GP mean profile, and the shaded regions are the corresponding 3σ bounds for each quantity of interest. In addition to the posterior distributions predicted by each GP, dashed lines are plotted for $\alpha > 0$ that represent the expected aerodynamic coefficient value due to aerodynamic symmetry (i.e. $C_L(-\alpha) = -C_L(\alpha)$ and $C_D(-\alpha) = C_D(\alpha)$) based on the GP's predictions for $\alpha < 0$. Note that using the predictions for $\alpha < 0$ from the GP as the expected values for $\alpha > 0$ is more valid when more training data is included, where the GP produces excellent predictions of all aerodynamic coefficients. However, for lower training data set sizes, it also illustrates the relative degree of aerodynamic symmetry of the response surfaces generated by the GP. As the number of data points used to train the GP increases, the predicted posterior distributions become more symmetric in angle of attack. The out-of-distribution ($\alpha > 0$) predictions by the GP show a remarkable level of accuracy by the time that 45 data points are included in the training process, as the posterior mean profiles for each aerodynamic coefficient show good agreement with their respective expected profiles. Generally, the GP mean underpredicts the lift coefficient and lift-to-drag ratio for $\alpha > 0$ and shows very good agreement with the drag coefficient. However, the expected profiles of all aerodynamic quantities of interest are contained within the posterior distribution predicted by the GP. Overall, the results in Figures 4.10–4.12 demonstrate that the GP-based model is able to extrapolate beyond its training data set with a notable level

of accuracy – a phenomenon that is not typically observed in GP regression. This is primarily attributed to the use of the LMC to correlate the lift and drag coefficients. If this aerodatabase had been constructed using a standard lookup-table technique, a standard linear extrapolation approach for extrapolating the aerodynamic coefficients for $\alpha > 0$ would be significantly less accurate than what is shown in Figure 4.12. Additionally, the usage of heritage look-table techniques would not produce uncertainty estimates that are consistent with the uncertainty in the data used to construct the model, whereas the uncertainty in the GP-based predictions are statistically consistent with the data used to train the model.

Three-dimensional CFD is required to generate an aerodynamic database that is capable of modeling aerodynamic coefficients as a function of angle of attack. However, the computational expense of a 3D-CFD simulation is significantly larger than that of an axisymmetric simulation, which increases the expense required to produce a sufficiently accurate aerodatabase. Therefore, it is of interest to investigate alternative methods to decrease the computational expense associated with generating a full 3D aerodatabase. In particular, this task also seeks to investigate the ability of the multi-fidelity modeling technique outlined in section 4.3.2 to decrease the number of 3D-CFD data sets necessary to assemble an aerodatabase.

To construct a multi-fidelity model, a specific fidelity level must be assigned to all available data. Specifically, this model uses the axisymmetric CFD data collected in Section 4.3.1 as the "low-fidelity" data, and the 3D-CFD data as the "high-fidelity" data. This fidelity hierarchy is chosen as an axisymmetric CFD is significantly less computationally expensive compared to a 3D-CFD simulation (e.g. an axisymmetric solution may be obtained in a few hours, whereas a 3D solution may require days to converge). As such, the multi-fidelity GP consists of two levels of fidelity where

$$f_1(\boldsymbol{x}) = \delta_1(\boldsymbol{x}) \tag{4.13}$$

$$f_2(\boldsymbol{x}) = \rho_1 \hat{f}_1(\boldsymbol{x}) + \delta_2(\boldsymbol{x}) \quad . \tag{4.14}$$

Here, $\boldsymbol{x} = [h,\ V,\ \alpha]$, the subscript $(\cdot)_1$ is the fidelity level of the axisymmetric CFD, and the subscript $(\cdot)_2$ is the fidelity level of the 3D-CFD. Recall that the multi-fidelity architecture described

in Section 4.3.2 assumes a nested structure to the data in each fidelity level. Fundamentally, this would require evaluating the low-fidelity model at nonzero angle-of-attack values. However, since an axisymmetric CFD simulation assumes $\alpha = 0$, it is presumed that the low-fidelity model returns the $\alpha = 0$ result even when queried at nonzero angles of attack. Therefore, by using the axisymmetric CFD as the low-fidelity data, the high-fidelity model is essentially trained to learn an angle-of-attack correction to the aerodynamic coefficients obtained by the axisymmetric model.

For each GP in the multi-fidelity model, the target vector consists of the lift and drag coefficients such that $\boldsymbol{y}_i = [C_{L,i},\ C_{D,i}]$, where the outputs are correlated using the IMC. Due to the axisymmetric assumption, the low-fidelity model returns $C_{L,i} = 0$ for all inputs. The low-fidelity GP is first trained using all 50 axisymmetric CFD data points. To assess the effect of varying the high-fidelity data-set size, high-fidelity models are then generated with $N = 10$, 20, 30, 40, and 45 3D-CFD data points. Figure 4.13 compares the relative error in the lift and drag coefficients predicted by the multi-fidelity model to the corresponding predictions obtained using the single-fidelity model generated using only 3D-CFD data.

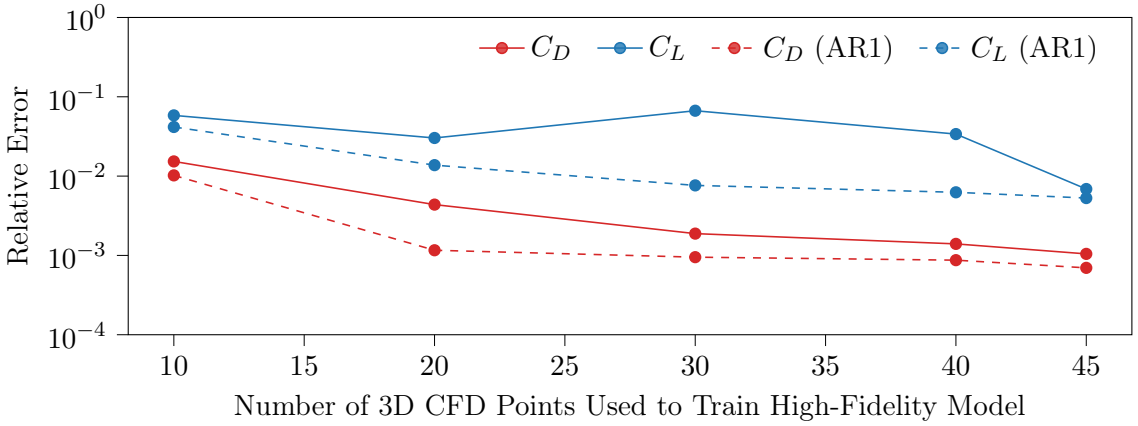

Figure 4.13: Relative error of GP trained to predict three-dimensional aerodynamic coefficients for training data sets that include various numbers of 3D-CFD data.

For all 3D-CFD data-set sizes, the multi-fidelity model produces more accurate predictions of both the lift and drag coefficients. Additionally, the multi-fidelity model converges to a low relative error more rapidly than the single-fidelity model, as only 20 3D-CFD data points are required to

produce predictions as accurate as the most accurate single-fidelity model. The multi-fidelity model is able to learn the general trends in the aerodynamics by using the axisymmetric CFD to train the low-fidelity model, and then anchor its predictions in the much sparser high-fidelity data set. Interestingly the single-fidelity model produces an increase in the lift-coefficient error for $N = 30$ training data points. This is likely due to the random sampling of the input space for training each model and greater variability in lift-coefficient values relative to the drag coefficient. The results in Figure 4.13 demonstrate that a multi-fidelity modeling architecture could significantly reduce the number of high-fidelity CFD simulations required to generate a full-scale aerodatabase, which would allow for more rapid vehicle analysis and trajectory design.

4.4.4 Using Gaussian-Process Based Models in Trajectory Simulations

As a final demonstration of the advantages associated with Gaussian-process based surrogate models for hypersonic entry vehicle design, the models trained in Sections 4.4.1 and 4.4.2 are used to evaluate the aerothermodynamic environment along the representative Neptune aerocapture trajectory in Figure 3.1 of Chapter 3.

The altitude and velocity values from Figure 3.1 are first used to evaluate the GP trained to predict the uncertain stagnation-point convective heat flux. Figure 4.14a shows 100 samples from the posterior distribution of the GP and Figure 4.14b shows the corresponding integrated heat load, $Q = \int q(t)\mathrm{d}t$. The black lines in Figure 4.14 are the corresponding stagnation-point heat flux and heat load predicted by the standard Sutton-Graves correlation for a hydrogen-helium atmosphere [158]. The predictions from the GP-based model reflect the uncertainty in the CFD simulations used to train the model, whereas the Sutton-Graves model only returns a deterministic estimate of the aerothermodynamic environment of the vehicle. To ensure that a vehicle's thermal protection system is sufficiently robust to its aerothermodynamic environment, the range of heat-flux and heat-load values shown in Figure 4.14 must be considered. The use of a GP-based model allows for the aerothermodynamic uncertainty to be propagated along the entire trajectory, allowing for a rapid assessment of the aerothermodynamic environment and its associated uncertainties. A

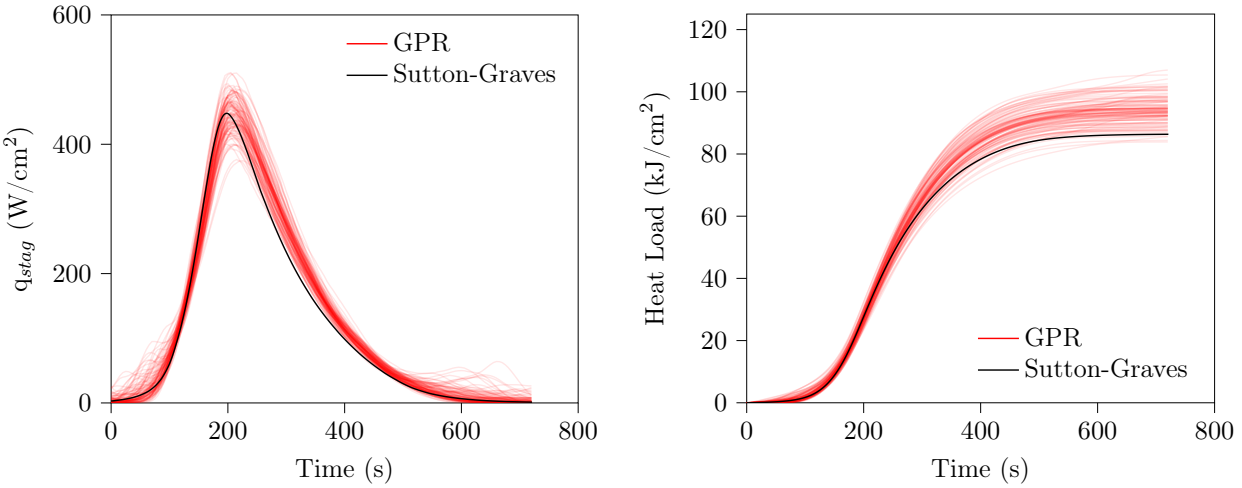

(a) Heat flux.

(b) Heat load.

Figure 4.14: Stagnation-heat flux (a) and integrated heat load (b) computed using GP-based model and Sutton-Graves correlation.

difference between the heat-flux magnitudes predicted by the GP-based model and the Sutton-Graves is observed, but this is expected, as the Sutton-Graves model is generally used as a low-fidelity engineering model to quickly estimate the stagnation-point heat flux.

Next, the model developed in Section 4.4.2 is used to predict the aerothermodynamic environment of multiple locations on the vehicle heat shield along the trajectory in Figure 3.1. Using the altitude and velocity values from the representative trajectory, the convective heat flux and heat load are predicted at the stagnation point ($s = 0$), the shoulder ($s = 0.95$), and a point midway along the frustum ($s = 0.5$). Figure 4.15 shows the ability of the GP-based model to quickly evaluate the aerothermodynamic environment of multiple locations along the vehicle heat shield. Estimates of the heat flux and heat load are necessary for TPS material selection and sizing, so a model that is capable of efficiently predicting entire heating distributions and corresponding heat-load profiles could significantly accelerate the TPS design phase.

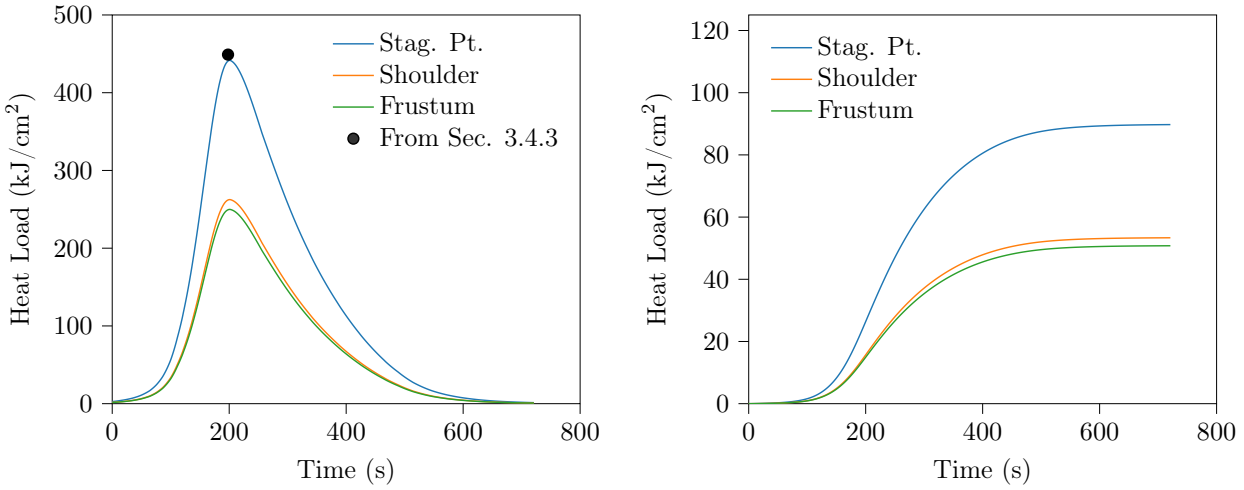

(a) Heat flux.

(b) Heat load.

Figure 4.15: Stagnation-heat flux (a) and integrated heat load (b) computed using GP model trained to predict convective heat-flux distributions.

4.5 Incorporation of Physical Constraints

In the previous section, the baseline GP models were trained without explicitly enforcing physical constraints, such as symmetries and monotonicity, resulting in increased prediction uncertainty in regions of sparse training data. This issue is particularly evident in the prediction of aerodynamic coefficients at positive angles of attack, where the GP model fails to capture the expected symmetric behavior about zero degrees for smaller training set sizes. To address these limitations, this section explores methods for incorporating physical constraints directly into the GP framework. Intentional manipulation of the GP kernel function's structure provides a means to improve predictive accuracy and enforce physically consistent trends, enhancing the reliability of the GP-based models for aerothermodynamic applications.

4.5.1 Encoding of Symmetry

A Gaussian process (GP) is completely defined by its mean and covariance functions, and it is typically assumed that the mean function of the GP prior is zero. As such, the form of the functions produced from a GP depend solely on the structure of the covariance (kernel) function.

The previous analyses in the chapter use "off-the-shelf" kernel functions, meaning no modification to them has been made, which makes the uncertainty in their predictions increase the further the model is sampled from the training data. Reducing this uncertainty typically requires collecting additional data in the uncertain region, however, this can be avoided in scenarios where the underlying function is symmetric (e.g. $f(x) = x^2$ is symmetric about $x = 0$). This symmetry can be exploited in the kernel design which not only produces more accurate model predictions, but it also reduces the amount of training data necessary to construct the model.

Expressing symmetry in a GP kernel is an example of a GP that has an "invariant" prior. This implies that there exists some finite set of transformations of the input space $\{g_1, g_2, \dots \ : \ g_i \in G\}$ to which the function should remain invariant

$$f(\boldsymbol{x}) = f(g(\boldsymbol{x})) \quad \forall \ \boldsymbol{x}, \quad \forall g \in G \quad . \tag{4.15}$$

For the functions produced by a GP to obey the set of invariances, the kernel function must also respect the set of invariances [160]

$$k(\boldsymbol{x}, \boldsymbol{x}') = k(g(\boldsymbol{x}), g(\boldsymbol{x}')) \quad \forall \ \boldsymbol{x}, \quad \forall g \in G \quad . \tag{4.16}$$

One technique for constructing an invariant kernel is "summing over the orbit" [161], where the "orbit" of $\boldsymbol{x}$ with respect to some group G is the set $\{g(\boldsymbol{x}) \ : \ g \in G\}$. In this approach, the invariance is expressed in the kernel as a double sum over the orbits of $\boldsymbol{x}$ and $\boldsymbol{x}'$

$$k_{\text{invariant}}(\boldsymbol{x}, \boldsymbol{x}') = \sum_{g \in G} \sum_{g' \in G} k(g(\boldsymbol{x}), g'(\boldsymbol{x}')) \tag{4.17}$$

where $k(\cdot, \cdot)$ is a standard kernel function (e.g. SE). As an example, consider the one-dimensional symmetry about $x = 0$ (i.e. reflection about the y-axis)

$$f(x) = f(-x) \quad . \tag{4.18}$$

Here, the set of transformations is $g_1(x) = x$ and $g_2(x) = -x$, which yields an invariant kernel structure of

$$k_{\text{reflect}}(x, x') = k(x, x') + k(x, -x') + k(-x, x') + k(-x, -x') \quad . \tag{4.19}$$

For stationary kernels like the SE and Matérn,[1] Equation 4.19 reduces to

$$k_{\text{reflect}}(x, x') = k(x, x') + k(x, -x') \quad . \tag{4.20}$$

This form of symmetry is particularly relevant to the blunt, axisymmetric aeroshells studied in this dissertation, as the drag coefficient should be symmetric about an angle of attack of zero. Before implementing Equation 4.20 for the entire aerodatabase, however, it is useful to demonstrate its capabilities on a simpler problem. Consider the noisy function $f(x) = x^2$, which is symmetric about $x = 0$. However, assume that measurements of this function are noisy such that

$$f(x) = x^2 + \epsilon \qquad \epsilon \sim \mathcal{N}(0, \sigma_n^2) \tag{4.21}$$

and that observations of $f(x)$ only exist for $x \geq 0$. Two cases are considered. The first trains a GP using the standard SE kernel, and the second uses the symmetric kernel defined in Equation 4.20, where the baseline kernel is an SE kernel. Figure 4.16 shows 50 functions sampled from the posterior distribution of both models. The invariant kernel maintains its predictive accuracy across the entire domain, while the standard kernel does not.

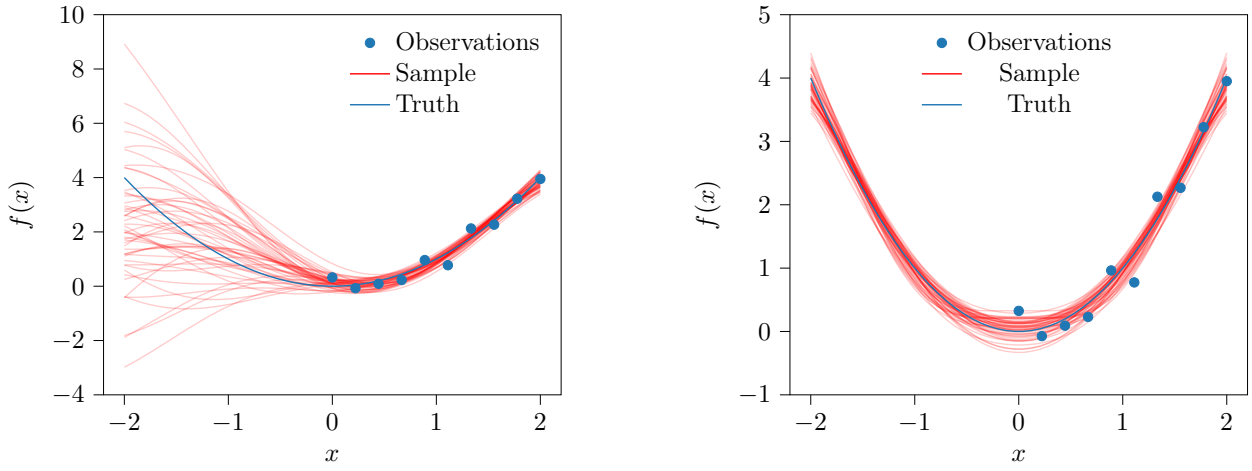

(a) Standard SE kernel.

(b) Invariant SE kernel.

Figure 4.16: Posterior distributions for GPs trained using a standard squared-exponential kernel (a) and a kernel designed to enforce symmetry about $x = 0$ (b).

[1] A *stationary* kernel only depends on the distance between the input locations $\|\boldsymbol{x} - \boldsymbol{x}'\|$

4.5.1.1 Quasi-Symmetry

An additional case considers functions that are "quasi-symmetric" about $x = 0$, such as $f(x) = x^3$. In reality, these functions are reflected about the x and y axis, but the approach outlined here is technically enforcing an invariance of the GP kernel about $x = 0$. Specifically, the mapping to be encoded within the structure of the GP kernel is

$$f(x) = -f(-x) \quad . \tag{4.22}$$

Because this is not an exact invariance, these mappings are referred to as "quasi-symmetric". For the axisymmetric aeroshells studied in this dissertation, the lift coefficient is quasi-symmetric in angle of attack (i.e. $C_L(\alpha) = -C_L(-\alpha)$). To incorporate the quasi-symmetry into the kernel, the sum over the orbit becomes

$$k_{\text{quasi}}(x, x') = k(x, x') - k(x, -x') \tag{4.23}$$

where it is assumed that the baseline kernel $k(\cdot, \cdot)$ is stationary.

To demonstrate the effect of using Equation 4.23 over a standard GP kernel, consider the noisy function

$$f(x) = x^3 + \epsilon \qquad \epsilon \sim \mathcal{N}(0, \sigma_n^2) \tag{4.24}$$

where training data is only available for $x \geq 0$. Figure 4.17 shows samples drawn from a GP trained using a standard SE kernel and its corresponding covariance matrix, and Figure 4.18 shows the same information obtained when using Equation 4.23 for the kernel. Note how the GP trained using Equation 4.23 produces superior predictions for $x < 0$ than the GP trained using a standard kernel. This improvement is reflected in the covariance matrices of both GPs. For the GP trained using a standard SE kernel, the covariance matrix has a single positive band along the diagonal (this is typical for SE kernels). However, the covariance matrix obtained using Equation 4.23 has a positive band along one diagonal, and a negative band along the opposite diagonal. This negative correlation is what enforces the mapping $f(x) = -f(-x)$ in the trained GP.

The framework described in this section may easily be expanded to multi-variate functions that have symmetry/quasi-symmetry about 0 in one (or more) of the input dimensions by using

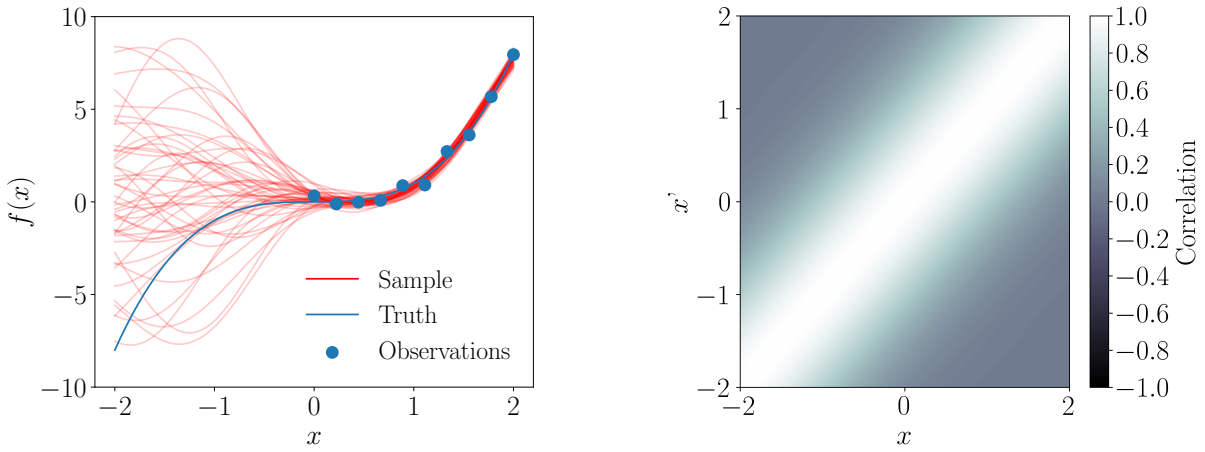

(a) Samples (b) Covariance.

Figure 4.17: Sampled functions (a) and resulting covariance matrix (b) for GP trained using a standard SE kernel.

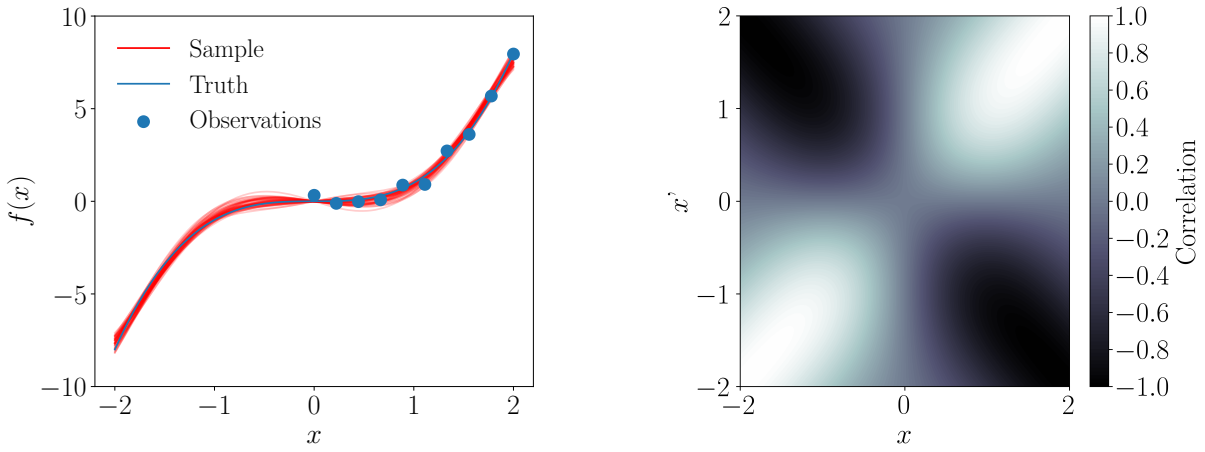

(a) Samples (b) Covariance.

Figure 4.18: Sampled functions (a) and resulting covariance matrix (b) for GP trained using a quasi-symmetric kernel function.

separate kernels for each input dimension. For example, if a function takes a two-dimensional input $\boldsymbol{x} = [x, y]$, and the function is symmetric only about $x = 0$, the invariant kernel is of the form

$$k(\boldsymbol{x}, \boldsymbol{x}') = k_{\text{reflect}}(x, x') + k_{\text{SE}}(y, y') \quad . \tag{4.25}$$

A similar approach can be taken for the quasi-symmetric case (or cases where the function has symmetries *and* quasi-symmetries). For instance, while the lift/drag coefficients for an axisymmetric aeroshell are quasi-symmetric/symmetric in angle of attack, they are not necessarily quasi-symmetric/symmetric in other input dimensions like altitude or velocity. Because the surrogates developed in this dissertation take multidimensional inputs, the invariances are constructed using kernels similar to the form of Equation 4.25.

4.5.2 Application to Aerodatabase

The first objective is to enhance the predictive capabilities of the three-dimensional aerodynamic GP model at positive angles of attack. As observed in the previous section, the baseline GP model, trained without enforcing physical constraints, exhibits increased uncertainty in this region. This behavior arises due to the lack of symmetry enforcement and the limited availability of training data at positive angles of attack. To address these issues, an invariant kernel is designed to encode aerodynamic symmetry and quasi-symmetry in the angle of attack coordinate. Because the symmetry is addressed directly in the kernel, the LMC is not used to correlate the outputs of the GP. Instead, two separate GPs are trained to predict the lift and drag coefficients, where the kernel function for each is

$$k_{C_D}(\boldsymbol{x}, \boldsymbol{x}') = k_{\text{reflect}}(\alpha, \alpha') + k_{\text{SE}}([h, v], [h', v']) \tag{4.26}$$

$$k_{C_L}(\boldsymbol{x}, \boldsymbol{x}') = k_{\text{quasi}}(\alpha, \alpha') + k_{\text{SE}}([h, v], [h', v']) \quad . \tag{4.27}$$

Here, k_{reflect} and k_{quasi} are expressed using Equations 4.20 and 4.23, respectively, and k_{SE} is a standard squared-exponential kernel.

Similar to the approach in Section 4.4.3, surrogates with varying training data-set sizes are generated. Figure 4.19 compares the predictions of the GP trained with the LMC and standard kernel to the GP trained with the symmetry-informed kernels for $N = 45$. The predictions at negative angles of attack are nearly identical for the two approaches, however, substantial differences are observed at positive angles of attack. Because of the symmetric invariance encoded within the kernel,

the symmetry-informed GP maintains its predictive confidence for $\alpha > 0$. Furthermore, it produces perfect symmetry in the angle-of-attack dimension, while the baseline model from Section 4.4.3 does not. This is further visualized in Figure 4.20, which compares the drag-coefficient response surfaces produced by the posterior mean of each GP model. Here, the symmetry is observed across the entire input space, emphasizing the capability of the invariant kernel to encode physical constraints without requiring additional training data. Generating data for use in hypersonic aerodatabases, whether computationally or experimentally, is expensive and a time-consuming process, so any way of accelerating the time to produce a sufficiently accurate surrogate is beneficial.

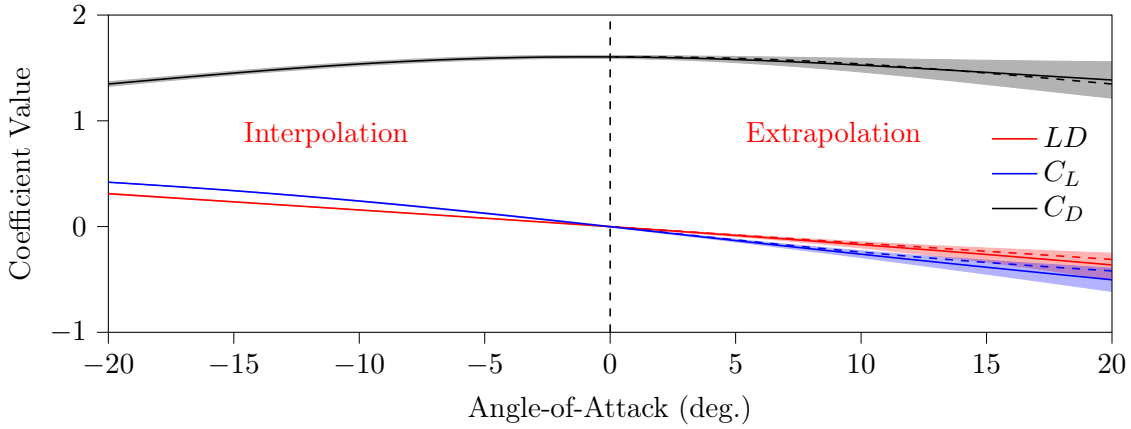

(a) Without symmetry.

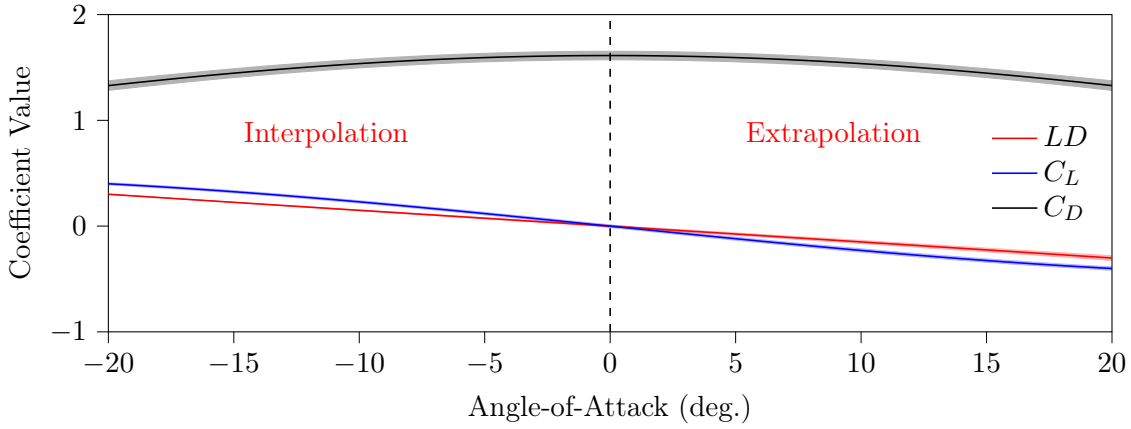

(b) With symmetry.

Figure 4.19: Comparison of the posterior distributions for aerodynamic coefficients at $h = 300$ km and $V = 30$ km/s without (a) and with (b) symmetry encoded in GP kernel.

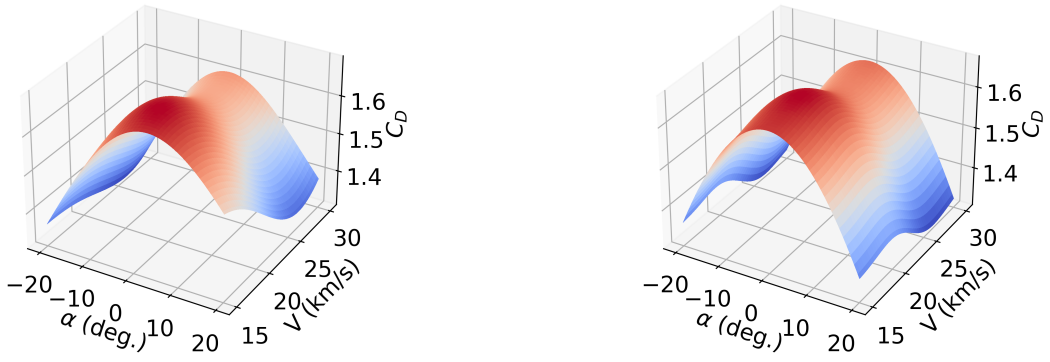

(a) Without symmetry.

(b) With symmetry.

Figure 4.20: Comparison of the posterior distributions for aerodynamic coefficients at $h = 300$ km and $V = 30$ km/s without (a) and with (b) symmetry encoded in GP kernel.

4.6 Summary

This chapter presented a Gaussian Process Regression (GPR) framework for developing efficient surrogate models to approximate high-fidelity aerothermodynamic uncertainty along an aerocapture trajectory. The constructed models successfully captured key aerodynamic and aerothermal quantities, including stagnation-point heat flux, aerodynamic coefficients, and full surface heat-flux distributions, demonstrating strong predictive accuracy and generalization. A multi-output GPR methodology was used to improve model predictions by identifying underlying correlations between data sets, and a multi-fidelity GPR approach was introduced to integrate low- and high-fidelity data, reducing computational cost while maintaining predictive performance.

Next, the surrogate models were incorporated into a trajectory simulation framework, highlighting the advantages of probabilistic aerodatabases in quantifying hysteresis effects due to aerothermodynamic uncertainty. The GPR models developed in this chapter will serve as a critical component in the stochastic guidance algorithm presented in Chapter 7. By integrating these probabilistic models into the guidance framework, the algorithm will be able to account for aerothermodynamic uncertainty explicitly in the trajectory design phase and acknowledge the inherent coupling between the vehicle aerodynamics and its trajectory.

Finally, to improve the physical consistency of the models, the GP kernel functions were carefully manipulated to enforce invariances in the GP posterior. Doing so ensured that the predictions adhered to expected aerodynamic trends, such as symmetry in the aerodynamic coefficients. These constraints enhanced the surrogate's predictive capabilities, particularly when extrapolating beyond the training dataset.

Chapter 5

Preliminaries - Trajectories and Guidance

"The missile knows where it is at all times. It knows this because it knows where it isn't."

— *Missile Guidance for Dummies*

5.1 Introduction

Similar to Chapter 2, this chapter introduces the fundamental tools and techniques used to generate aerocapture trajectories. It begins by describing the dynamical models used to simulate the motion of an aerocapture spacecraft, including methods for addressing non-spherical gravitational forces and uncertain aerodynamic forces. An overview of bank-angle modulation, the technique used in this dissertation for steering an aerocapture vehicle, is then provided. The necessity of aerocapture guidance is introduced by defining the aerocapture entry corridor, followed with a description of the current state-of-the-art full numerical predictor-corrector aerocapture guidance (FNPAG) algorithm.

Second, this chapter introduces the mathematical principles underpinning each of the guidance algorithms developed in this dissertation. Specifically, the basics of convex optimization and sequential convex programming (SCP) - a technique capable of addressing the limitations of the FNPAG algorithm - are described. An overview of SCP for trajectory optimization is given, including several general tutorials developed by the author for applying the principles of SCP to both deterministic and stochastic optimization problems. Finally, the chapter concludes with a description of the usage of chance constraints for constraining uncertain variables in stochastic optimization

problems.

5.2 Dynamical Models

This section describes the equations of motion used to model a hypersonic vehicle moving through a planet's atmosphere subject to atmospheric and gravitational forces. These are vital to characterize in the design of aerocapture trajectories, as they shape the overall flight profile of the spacecraft. Furthermore, the atmospheric models are often highly uncertain and are therefore necessary to study when designing robust aerocapture guidance algorithms.

5.2.1 Coordinate Frame

The motion of an aerocapture vehicle is described in a *planetocentric* coordinate system in which the origin is located at the center of the planet. A convenient coordinate frame for atmospheric flight is the planet-centered, planet-fixed frame expressed in spherical coordinates. In a planet-fixed frame, the coordinate system rotates with the central body. The position is described by the radial distance from the center of the planet r, the longitude θ, and the latitude ϕ, as shown in Figure 5.1. Longitude is measured from the x-axis and is defined positive to the east, and latitude is measured from the equatorial plane along a meridian and is defined positive to the north.

The velocity vector is described by the velocity magnitude and the two angles shown in Figure 5.2: the flight-path angle γ, and the heading angle ψ. As illustrated in Figure 5.2a, the flight-path angle γ is defined as the angle between the local horizontal plane (the plane passing through the vehicle that is orthogonal to the radius vector) and the velocity vector. It is positive when the velocity vector is above the horizontal plane. The heading angle ψ is defined as the angle between the local parallel of latitude and the projection of the velocity vector onto the horizontal plane. It is measured positive in the right-handed direction about the radius vector, as shown in Figure 5.2b.

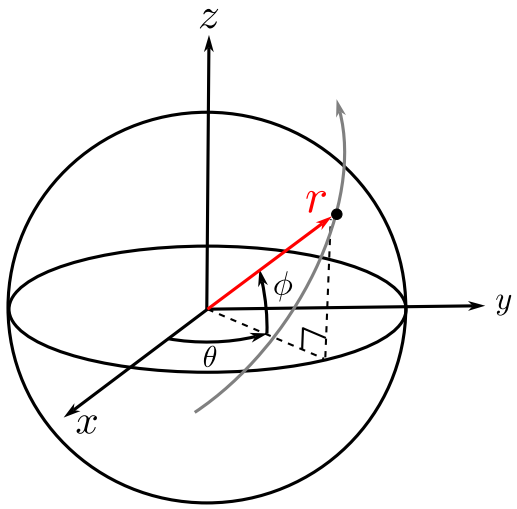

Figure 5.1: Description of position vector in spherical coordinates.

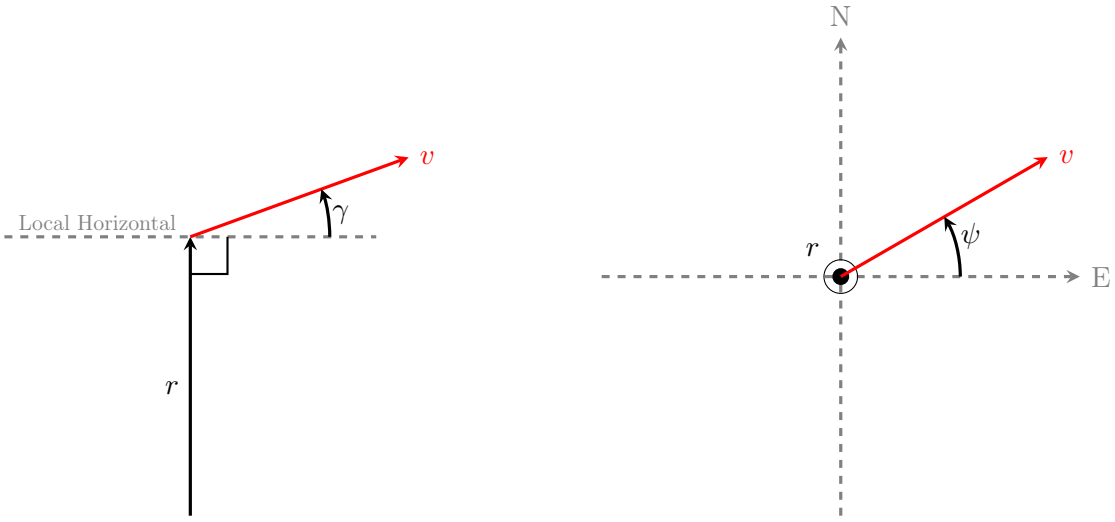

(a) Flight-path angle.

(b) Heading angle.

Figure 5.2: Definition of the flight-path angle (a) and heading angle (b).

5.2.2 Equations of Motion

The six coordinates used to describe the position and velocity are defined as the *states* of the vehicle. To describe the evolution of each state, the following system of equations is used [162]

$$\dot{r} = v \sin\gamma \tag{5.1a}$$

$$\dot{\theta} = \frac{v \cos\gamma \cos\psi}{r \cos\phi} \tag{5.1b}$$

$$\dot{\phi} = \frac{v \cos\gamma \sin\psi}{r} \tag{5.1c}$$

$$\dot{v} = -D - g_r \sin\gamma - g_\phi \cos\gamma \sin\psi + \Omega^2 r \cos\phi \left(\sin\gamma \cos\phi - \cos\gamma \sin\phi \sin\psi\right) \tag{5.1d}$$

$$v\dot{\gamma} = L \cos\sigma - g_r \cos\gamma + \frac{v^2}{r} \cos\gamma + g_\phi \sin\gamma \sin\psi + 2\Omega v \cos\phi \cos\psi + \tag{5.1e}$$

$$\Omega^2 r \cos\phi \left(\cos\gamma \cos\phi + \sin\gamma \sin\phi \sin\psi\right)$$

$$v\dot{\psi} = \frac{L \sin\sigma}{\cos\gamma} - \frac{v^2}{r} \cos\gamma \cos\psi \tan\phi - g_\phi \frac{\cos\psi}{\cos\gamma} + 2\Omega v \left(\tan\gamma \cos\phi \sin\psi - \sin\phi\right) \tag{5.1f}$$

$$- \Omega^2 r \frac{\sin\phi \cos\phi \cos\psi}{\cos\gamma}$$

$$\dot{\sigma} = u \qquad . \tag{5.1g}$$

Equation 5.1 is simply Newton's second law of motion applied to the motion of a vehicle about an oblate, rotating planet in spherical coordinates. Effects of planetary rotation are captured by the terms involving the rotation rate Ω. In this definition, the vehicle state has been augmented with the bank-angle σ, the rate of which is defined as the control input u. The actual control applied to the spacecraft would be a torque, for which the bank-angle rate serves as a proxy. As shown in Figure 5.3, the bank-angle is defined as the rotation of the lift vector about the velocity vector. A bank-angle of $\sigma = 0$ deg. is referred to as "lift-up", and a bank-angle of $\sigma = 180$ deg. is referred to as "lift-down". Modulation of the bank angle between these two limits is a common technique for steering hypersonic reentry vehicles, including the Space Shuttle [163], Mars Science Laboratory [164], Mars 2020 [165], and Orion spacecraft [53]. Changing the bank angle affects the amount of lift acting in the vertical and horizontal planes which allows for both the longitudinal/in-plane and lateral/out-of-plane motion to be controlled. Note that bank-angle modulation is only effective if

the spacecraft is capable of producing lift. For the axisymmetric, blunt reentry capsules studied in this dissertation, lift is typically produced by trimming the vehicle to a non-zero angle of attack, α, using a center-of-gravity offset. Determining the bank-angle profile along the aerocapture trajectory is the fundamental task of the guidance algorithms developed in this dissertation.

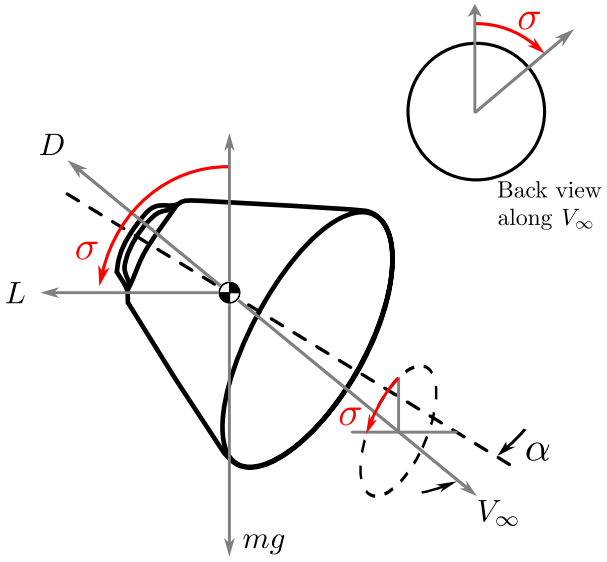

Figure 5.3: Bank-angle definition.

5.2.3 Force Models

The planet itself is modeled as an oblate spheroid which leads to non-uniform gravitational accelerations. Accounting for non-spherical gravitational fields is more important for aerocapture than other atmospheric flight applications because they influence the spacecraft's path not just during the atmospheric pass, but also in the resulting orbit. A widely used method for modeling these higher-order gravitational perturbations is spherical harmonics which represents the gravitation potential as an infinite sum. In practice, however, this sum is typically truncated to a finite set of terms because the higher-order terms are small or their effects are negligible. For large planetary bodies like those studied in this dissertation, the most dominant gravitational perturbation comes from the second zonal-harmonic or "J_2" which captures the effects of oblateness on the gravitational

field. Truncating the spherical-harmonic sum to only include terms with J_2 yields the following radial and latitudinal gravity components

$$g_r = \frac{\mu}{r^2}\left(1 + J_2\left(\frac{R_e}{r}\right)^2\left(1.5 - 4.5\sin^2\phi\right)\right) \tag{5.2a}$$

$$g_\phi = 3\frac{\mu}{r^2}\left(J_2\left(\frac{R_e}{r}\right)^2\sin\phi\cos\phi\right) \tag{5.2b}$$

where μ is the planet's gravitational parameter, and R_e is the planet's equatorial radius. Table 5.1 summarizes the gravitational force-model parameters for the three destinations considered in this research. Note that if $J_2 = 0$, Equation 5.2 reduces to the standard Keplerian expression for a point-mass gravity model.

Table 5.1: Planetary constants for each destination.

Planet	μ (km^3/s^2)	R_e (km)	J_2 $(\cdot)$	Ω (rad/s)
Neptune	6,835,100	24,764	3.411×10^{-3}	1.0834×10^{-4}
Mars	42,828	3,396.2	1.960×10^{-3}	7.088×10^{-5}
Earth	398,600	6,378.1	1.083×10^{-3}	7.292×10^{-5}

In Equation 5.1, L and D are the lift and drag accelerations acting on the vehicle, respectively. For a non-thrusting spacecraft they are given as

$$L = \frac{\rho v^2 L/D}{2\beta} \tag{5.3}$$

$$D = \frac{\rho v^2}{2\beta} \quad . \tag{5.4}$$

Here, ρ is the local freestream density, $L/D = \dfrac{C_L}{C_D}$ is the lift-to-drag ratio, and $\beta = \dfrac{m}{C_D A_{\text{ref}}}$ is the ballistic coefficient. The dynamics of the vehicle are highly dependent upon the atmospheric density, and accurate knowledge of the density is a driving factor in an aerocapture guidance algorithm's performance. The Global Reference Atmosphere Models (GRAMs) are used to capture the geospatial variation and uncertainty in the atmospheric properties of the aerocapture destinations studied in this research. GRAM is an engineering model for simulating the variation in atmospheric properties of various planets and has been used extensively for entry and aerocapture

studies [164, 166, 167, 168, 55, 169, 170, 74, 171, 172]. Using GRAM allows for both nominal and dispersed density profiles to be sampled, and Figure 5.4 illustrates the density variability produced by NeptuneGRAM, MarsGRAM, and EarthGRAM. The shaded regions denote the $\pm 3\sigma$ variability, and the colored lines denote three dispersed density profiles. The data in NeptuneGRAM is from a single Voyager flyby, while the data in MarsGRAM and EarthGRAM is much denser. This explains the higher density variability at low altitudes observed for NeptuneGRAM.

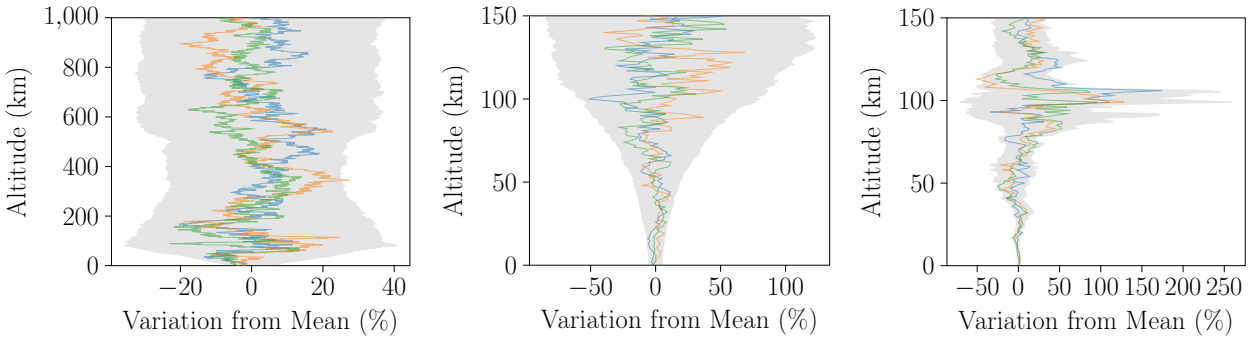

(a) Neptune. (b) Mars. (c) Earth.

Figure 5.4: Density variation produced by GRAM for Neptune (a), Mars (b), and Earth (c).

5.3 Aerocapture Corridor

An important metric when assessing the feasibility and mission requirements for aerocapture is the *entry corridor width*. Generally, the entry corridor width is defined as the range of entry flight-path angles that produce feasible aerocapture trajectories. There are two types of corridor widths: theoretical and "mission-specific," and the criteria for a feasible aerocapture trajectory depends on which corridor-width classification is used. For the theoretical entry corridor, a feasible aerocapture trajectory is one which places the spacecraft into *any* orbit (i.e. does not crash or escape), while the mission-specific corridor requires that the trajectory results in the spacecraft entering a *specific* orbit. For both classifications, the corridor width is determined by propagating full lift-up and full lift-down trajectories from atmospheric interface to atmospheric exit and determining the steepest and shallowest permissible entry flight-path angles, respectively. Because lift-up and lift-down produce

the maximum amount of vertical lift, these trajectories can be used to bound the trajectory space for a given entry vehicle and target orbit.

Figure 5.5 presents the trajectories for each entry corridor definition for a Neptune aerocapture scenario with an entry velocity of 29 km/s. The vehicle is assumed to have a constant lift-to-drag ratio of $LD = 0.3$ and ballistic coefficient of $\beta = 110$ kg/m^2, and the atmospheric density is taken to be the nominal NeptuneGRAM profile. For the mission-specific corridor, the target orbit is taken to have an apoapsis radius of $r_{a,\text{target}} = 100,000$ km.

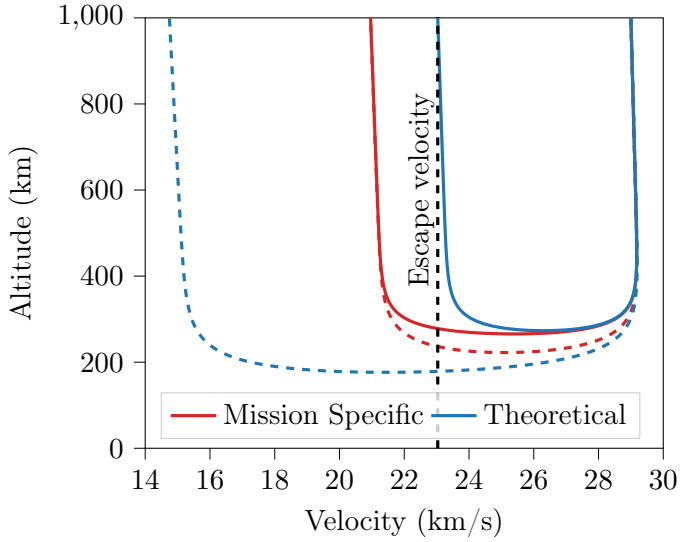

Figure 5.5: Comparison of theoretical and mission-specific entry corridor trajectories for Neptune aerocapture. The solid lines denote the lift-up trajectories, and the dashed lines denote the lift-down trajectories.

For this configuration, the theoretical corridor width is 2.27 degrees, and the mission-specific corridor width is 1.07 degrees. Characterization of the entry corridor width is necessary for designing the hyperbolic approach trajectory, as it should be wider than the 3σ delivery uncertainty of the entry flight-path angle. Because the lift-up and lift-down trajectories bound the flight envelope, they also bound the aerothermodynamic environment of the entry vehicle. In particular, the lift-up trajectory produces the maximum heat flux, while the lift-down trajectory produces the maximum integrated heat load. Note that the corridor widths reported here assume a nominal entry state and nominal models for the density and aerodynamics of the entry vehicle. In reality, each of these

are uncertain, and any uncertainty introduced into the system narrows the corridor width. For the Monte-Carlo analysis performed in Chapter 6, the nominal entry flight-path angle is chosen to be the center of the entry-corridor.

5.4 Aerocapture Guidance

Each of the models described in the preceding sections constitutes an input to the design and optimization of an aerocapture trajectory[1] . However, as highlighted by the quote at the beginning of Chapter 2, every model is inherently wrong and has some level of uncertainty associated with it. The entry-corridor calculations in the previous section assume the aerocapture vehicle has constant aerodynamic properties, yet the results of Chapters 3 and 4 indicate that this assumption is invalid. Consider the scenario where an engineer is tasked with designing the bank-angle profile along an aerocapture trajectory. Likely, he or she will begin by assuming that the entry vehicle has defined characteristics – lift-to-drag ratio, ballistic coefficient, etc.– that the atmosphere has a known density profile, and that the entry condition is fixed. From inspection of Figures 5.4 and 5.5, these are likely poor assumptions. So, what would happen if, when the spacecraft arrives to the target destination, the *actual* aerodynamic properties of the vehicle or the *actual* density profile differ from the models used to derive the bank-angle profile? Or the *actual* entry condition is different? The trajectory would deviate substantially from the desired path, and the vehicle would miss the target orbit (or even fail to capture into orbit). To avoid these catastrophic results, it is necessary to update the bank-angle profile in real time as the vehicle is progressing along its trajectory using information about where the vehicle is (i.e. current state) and where it is trying to go (i.e. target orbit). This is the role of the *guidance* system.

While many definitions for guidance exist, the following quote concisely summarizes the role of guidance in this research [173].

> "Guidance is about the determination of the maneuvering commands to steer the

[1] In reality, there are many more models such as mass properties, reaction control system (RCS), and six-degree-of-freedom aerodynamics, that go into a high-fidelity trajectory simulation, but that is beyond the scope of this discussion.

> vehicle to fly a trajectory that satisfies the specified terminal/targeting condition as well as other pertinent constraints, and, if required, optimizes a defined performance."
>
> — *Ping Lu, Editor-in-Chief - Journal of Guidance, Control, and Dynamics*

The guidance system can be thought of as the "brains" of the spacecraft. It makes decisions based on information provided from the navigation system and given models for the spacecraft and environment that will steer the vehicle to its desired state. Guidance is often conflated with control, so a brief distinction between the two is made here. The control system is more like the "muscles" of the spacecraft – it takes the maneuvering command output from guidance and then determines the necessary forces/torques or actions of the control effectors (thrusters, flaps, etc.) in order for the guidance command to be realized.

Guidance encompasses a wide range of concepts, and no single method defines the design of a guidance algorithm. Common approaches include reference-based guidance, target-condition-driven guidance, and computational guidance, among others [173]. The algorithms developed in this dissertation fall within the realm of computational guidance, characterized by their reliance on physics-based models and extensive numerical computation. As discussed in Chapter 1, a prominent example of such an approach is the Fully Numerical Predictor-Corrector Aerocapture Guidance (FNPAG) algorithm [55]. FNPAG has been used for numerous aerocapture studies [56, 57, 58, 59, 174] and is considered to be the state of the art. Given that the algorithms developed in this research are directly benchmarked against FNPAG, it is necessary to discuss its theoretical foundation and implementation details.

5.4.1 Fully Numerical Predictor-Corrector Aerocapture Guidance (FNPAG)

FNPAG is derived from optimal control theory, in which the performance objective (cost function) is the in-plane ΔV required to reach the final orbit after aerocapture. However, the optimal aerocapture guidance problem is highly non-convex, making it difficult to solve. To alleviate some of this difficulty, the FNPAG guidance law is derived considering only the longitudinal equations of motion and excludes path constraints like aerodynamic loading and heating. As such, the equations

of motion in Equation 5.1 are reduced to

$$\dot{r} = v \sin\gamma \tag{5.5a}$$

$$\dot{v} = -D - \frac{\mu \sin\gamma}{r^2} \tag{5.5b}$$

$$\dot{\gamma} = \frac{1}{v}\left[L\cos\sigma + \left(v^2 - \frac{\mu}{r}\right)\frac{\cos\gamma}{r}\right] \quad . \tag{5.5c}$$

Note these equations do not consider the effects of a rotating planet or non-spherical gravity. Because the cost function is defined to be a function only of the terminal longitudinal state, Pontryagin's Minimum Principle [175] may be used to express the form of the optimal bank angle

$$\sigma = \underset{0\leq\sigma\leq\pi}{\arg\min}\left\{p_r v\sin\gamma + p_v\left(-D - \frac{\mu\sin\gamma}{r^2}\right) + p_\gamma\left[\frac{L}{v}\cos\sigma + \left(v^2 - \frac{\mu}{r}\right)\frac{\cos\gamma}{rv}\right]\right\} \tag{5.6}$$

where p_r, p_v, and p_γ are the costate variables. Recognizing that Equation 5.6 is linear in $\cos\sigma$ and that $\frac{L}{v} > 0$, the optimal bank-angle structure is found to be bang-bang

$$\sigma^* = \begin{cases} \sigma_{\min} & \text{if } p_\gamma < 0 \\ \sigma_{\max} & \text{if } p_\gamma > 0 \end{cases} \tag{5.7}$$

where the minimum and maximum bank angles are $[\sigma_{\min}, \sigma_{\max}] = [0, \pi]$. There is an intermediate possibility for Equation 5.7, where $p_\gamma = 0$ over some finite time interval, implying that the optimal bank-angle value could lie anywhere between its bounds. This is known as "singular" optimal control, however, the possibility of such a case is proven to be impossible in an aerocapture scenario [55].

While it is impractical to solve the full optimal control problem onboard, the theoretical insights produced by the optimal-control analysis allow the FNPAG guidance law to be designed with optimality in mind. The bang-bang bank-angle profile produces a one-dimensional, parametric control profile that is assumed by the guidance system. This reduces the computational complexity of the guidance correction phase, as only a single parameter needs to be optimized.

FNPAG approaches the guidance problem using a two-phase algorithm, and each phase uses a numerical predictor-corrector (NPC) to solve a single-variable root-finding problem. In the first phase, the NPC solves for the time at which the bang angle switches from lift-up to lift down, denoted as the switching time t_s, that satisfies an apoapsis-targeting condition. During phase one,

a bang-bang bank-angle profile is enforced during the prediction component of the NPC. While theoretically optimal, assuming a pure bang-bang profile in which the bank-angle switches from 0 deg. to 180 deg. is not robust to dispersions, as it leaves no margin for correction later in the trajectory. Therefore, it is customary to use a larger lower bound and smaller upper bound during the prediction portion of phase one. In the FNPAG literature, these lower and upper bounds are denoted as σ_0 and σ_d, respectively [55]. Once the current $t > t_s$, the algorithm transitions to phase two in which the NPC seeks a constant bank-angle magnitude that satisfies the apoapsis targeting condition. For both phases, the apoapsis-targeting condition is expressed assuming Keplerian motion from atmospheric exit to apoapsis

$$r_a(\boldsymbol{x}_f) - r_{a,\text{target}} = 0 \qquad\qquad r_a(\boldsymbol{x}_f) = a\left(1 + \sqrt{1 - \frac{v_f^2 r_f^2 \cos^2\gamma_f}{\mu a}}\right) \tag{5.8}$$

where a is the semi-major axis

$$a = \frac{\mu}{2\mu/r_f - v_f^2} \quad . \tag{5.9}$$

Because FNPAG uses bank-angle modulation, a separate lateral logic algorithm is required to determine the sign of the bank angle (either positive or negative), as the longitudinal channel of the guidance system only computes the bank-angle magnitude [176]. For aerocapture, the lateral logic is used to help target an orbit with a specific inclination (or wedge angle[2] if a specific right ascension of the ascending node is known). If the bank-angle sign has to be reversed, the vehicle must fly an open loop control profile until the reversal is completed which makes it more susceptible to dispersions and accumulation of error. The version of FNPAG used in this research uses a predictive lateral-logic law developed by Smith [176] to manage inclination error. Simultaneously solving for the bank-angle magnitude *and* direction is prohibited by the fundamental assumption of longitudinal motion made in the optimal-control analysis.

FNPAG is a highly-capable aerocapture guidance algorithm, however, it does have its limitations. These have been mentioned throughout this section, but it is worth highlighting them explicitly. Namely, these are the inability to consider system-level constraints such as aerothermo-

[2] The wedge angle defines the angle between current and desired orbit angular-momentum vector.

dynamic heating and loading, and the inability to explicitly solve for the bank-angle direction. Each of these shortcomings is a product of the simplifications made during the derivation of the FNPAG guidance law, as it is computationally intractable to solve the full, nonlinear, constrained optimal control problem onboard a spacecraft.

5.5 Convex Optimization

An attractive characteristic of an optimization problem is its *convexity* [177]. Optimization problems that are *convex* can be solved both reliably and efficiently – two key characteristics of a robust guidance algorithm. While nonlinear, non-convex optimal control problems can require substantial computational effort to solve, modern interior-point algorithms (IPMs) are capable of solving convex-optimization problems in polynomial time [178]. Furthermore, if a problem is convex, an algorithm is guaranteed to converge to a global optimum, provided the problem is feasible. If the problem is infeasible, convex-optimization algorithms can return a "certificate of infeasibility" which provides a mathematical basis for the source of the infeasibility.

Unfortunately, most aerospace optimal control problems are non-convex due to their nonlinear dynamics and complex constraints. Recently, however, there has been a growing interest in solving non-convex optimization problems using convex-optimization-based algorithms, due to their computational efficiency and demonstrated ability to solve a wide range of complex problems [179, 180]. In specific cases, it is possible to perform *lossless* convexification of the non-convex problem, by which the original, non-convex optimal control problem is exactly converted into an equivalent convex optimization problem [61, 181]. Sadly, this only applies to a specific class of optimal control problems. Because equations of motion are enforced as equality constraints in optimal control problem, scenarios that involve highly nonlinear dynamics, such as those in Equation 5.1, cannot be solved using lossless convexification. Instead, a more general convex-optimization framework, known as sequential convex programming (SCP), must be used. This section presents the mathematical principles of convex optimization and SCP, as they are the foundation for the two guidance algorithms presented in Chapters 6 and 7.

5.5.1 Basic Definitions

A parametric optimization problem may generally be stated as

$$\begin{aligned} \min_{\boldsymbol{x} \in \mathcal{X}} \quad & f_0(\boldsymbol{x}) \\ \text{s.t.} \quad & f_i(\boldsymbol{x}) \leq 0 \quad \forall i \in \{1, \ldots, m\}, \\ & h_i(\boldsymbol{x}) = 0 \quad \forall i \in \{1, \ldots, p\} \end{aligned} \tag{P}$$

which is commonly referred to as the nonlinear programming (NLP) form. Here, the optimization variables are $\boldsymbol{x} \in \mathcal{X} \subset \mathbb{R}^n$, where $\mathcal{X}$ is the set in which $\boldsymbol{x}$ is contained. In order for (P) to constitute a *convex* optimization problem, the following criteria must be met [177]:

(1) The set $\mathcal{X}$ must be a *convex* set.

(2) The objective function $f_0(\boldsymbol{x})$ is a *convex* function in $\boldsymbol{x}$.

(3) The inequality constraints $f_i(\boldsymbol{x})$ are *convex* functions in $\boldsymbol{x}$.

(4) The equality constraints $h_i(\boldsymbol{x})$ are *affine* functions of $\boldsymbol{x}$.

A set C is considered convex if the line segment between any two points in C also lies in C. This may be formally stated as

$$\theta \boldsymbol{x}_1 + (1-\theta)\boldsymbol{x}_2 \in C \qquad \forall\, \boldsymbol{x}_1, \boldsymbol{x}_2 \in C \qquad 0 \leq \theta \leq 1 \tag{5.10}$$

More generally, a set is convex if every point in the set can be "seen" by every other point, along an unobstructed straight path between them, where unobstructed means lying in the set [177]. Furthermore, a function $f : \mathbb{R}^n \mapsto \mathbb{R}$ is convex if its domain, $\mathrm{dom}(f)$, is a convex set and

$$f\left(\theta \boldsymbol{x} + (1-\theta)\,\boldsymbol{y}\right) \leq \theta f(\boldsymbol{x}) + (1-\theta)\, f(\boldsymbol{y}) \qquad \forall\, \boldsymbol{x}, \boldsymbol{y} \in \mathrm{dom}(f) \qquad 0 \leq \theta \leq 1 \quad . \tag{5.11}$$

Geometrically, this definition implies that the line segment between the points $(\boldsymbol{x}, f(\boldsymbol{x}))$ and $(\boldsymbol{y}, f(\boldsymbol{y}))$ lies above the graph of f. Similarly, it implies that the function's gradient (tangent line) is always below the function. This is visually represented for a one-dimensional function in Figure 5.6.

Finally, an *affine* function is a function that is linear in the optimization variables such that the equality constraints are expressed in the form $h_i(\boldsymbol{x}) = \boldsymbol{a}_i^T \boldsymbol{x} + b_i$.

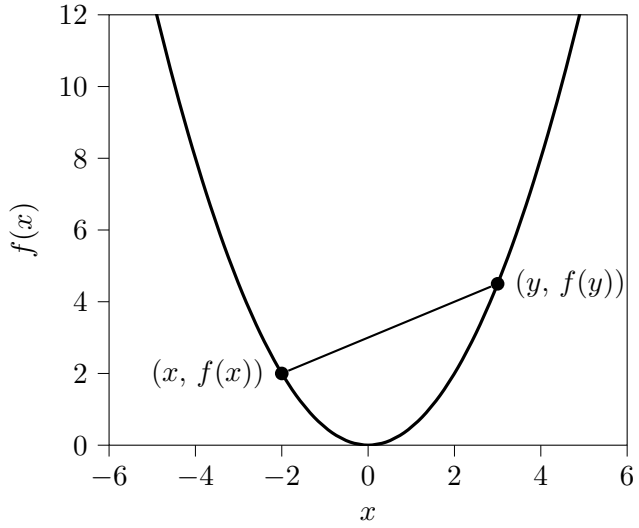

Figure 5.6: Convex function in one dimension.

5.5.2 Sequential Convex Programming

SCP takes an iterative *convexify-and-solve* process to march toward a solution of the original non-convex problem [182]. SCP works by linearizing or convexifying the non-convex parts of the problem — such as nonlinear dynamics and/or constraints — around an initial guess/reference trajectory. Each iteration then solves this convex *subproblem*, typically using efficient convex optimization methods like interior-point algorithms, and updates the reference using the solution to the subproblem. Because of its convex nature, the solution to the subproblem is globally optimal but only valid within the approximations made during the convexification process (i.e. linearization). For example, Figure 5.7 shows an example of using SCP to minimize the non-convex function $f(x) = x^4 - 2x^3 + 0.3x$. The algorithm is initialized for two different initial conditions, and at each iteration k, the convex approximation is formed about the global minimum of the previous *approximation*. Because the convex approximation is only locally optimal, the converged minimum is sensitive to the initial guess x_0. While SCP doesn't guarantee convergence to a *global* optimum of the original problem, its computational efficiency and ability to handle complex systems, like those in aerocapture guidance and control, make it a powerful practical tool.

SCP has seen great success for a variety of applications including, but not limited to, pow-

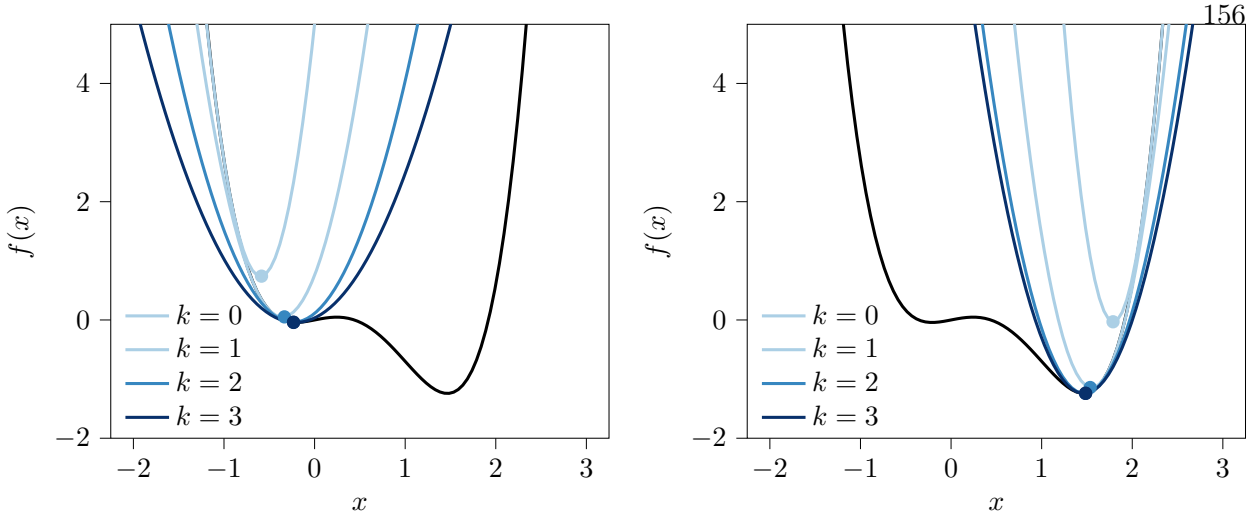

(a) $x_0 = -1$.

(b) $x_0 = 2.25$

Figure 5.7: Sequential convex approximations (blue) of the non-convex function $f(x) = x^4 - 2x^3 + 0.3x$ (black) initialized at $x_0 = -1$ (a) and $x_0 = 2.25$ (b).

ered descent guidance [61, 62, 63], atmospheric entry [64, 65, 66], low-thrust trajectory design [67], aerocapture trajectory optimization [68, 69, 70, 71], and stochastic feedback control policy design [72, 73]. Additionally, SCP is capable of handling complex constraints, such as state-triggered constraints [183, 63], nonlinear equality constraints [184], and reaction control system (RCS) constraints [185] that would be difficult (if not impossible) to enforce using alternative trajectory optimization methods.

Generally, the convex subproblem must be recast into the form of (P) (satisfying criteria for convexity) so that it can be interpreted and solved by an existing convex-optimization algorithm. This can be a tricky task to accomplish by hand, as it requires significant bookkeeping of the problem structure. Fortunately, existing convex-programming interfaces such as `CVX` [186] in Matlab, `CVXPY` [187] in Python, and `Convex.jl` [188] in Julia may be used to automatically parse the optimization problem into a solver-compatible form, allowing the user to express the optimization problem more naturally. All of the convex-optimization implementations in this research are formulated using `Convex.jl`[3] . For further details on the usage of SCP for general aerospace trajectory optimization,

[3] Tutorials on using `Convex.jl` for trajectory optimization can be found on the author's GitHub

the reader is referred to Ref. [182].

5.6 Stochastic Optimization

SCP provides a powerful baseline framework for solving complex trajectory-optimization problems by iteratively approximating nonlinear constraints and dynamics with convex forms, yielding a solution that is *explicitly* defined for the problem's parameters and initial conditions. This means that the optimized trajectory or control sequence is precisely determined for those specific starting parameters; however, any subsequent changes — whether in the environment, vehicle state, or mission requirements — necessitate adjustments to the original solution. In Chapter 6, the control profile is updated in response to perturbations along the trajectory using a predictor-corrector algorithm that resolves a deterministic optimization problem at a given frequency. The iterative nature of a predictor-corrector algorithm is what makes them robust to uncertainty. An alternative approach to addressing uncertainty is to include it directly within the optimization problem.

Stochastic sequential convex programming (SSCP) extends the SCP approach to consider uncertainty directly within the optimization framework. Instead of iteratively solving a deterministic optimization problem, SSCP simultaneously optimizes a control policy that uses both *feedforward* and *feedback* terms by solving a stochastic optimization problem. The optimized feedback control policy allows for the nominal feedforward control profile to be updated in the loop, replacing the need to resolve the optimization problem onboard. In contrast to the deterministic version of SCP, the stochastic optimization process *implicitly* accounts for uncertainty in the system. In this research, uncertainty exists in nearly all models used to simulate an aerocapture trajectory.

Including uncertainty directly within an optimization problem is challenging because constraints involving uncertain variables must be treated probabilistically. For example, consider the scenario where an optimization variable x is constrained to be less than a certain value such that $x_k \leq b \; \forall \; k$. A deterministic solution to the optimization problem may produce a trajectory like the black line in Figure 5.8, where the inequality constraint is active at t_k. However, now let the initial condition of x be normally distributed (a common assumption in control applications) with mean

$\bar{x}$ and covariance P_x such that $x_0 \sim \mathcal{N}(\bar{x}, P_x)$. Propagating the control profile obtained from the deterministic optimization problem would result in a distribution of trajectories, as shown by the shaded region in Figure 5.8, many of which would violate the constraint $x_k \leq b$. So, the question becomes: how can the *distribution* of x be constrained?

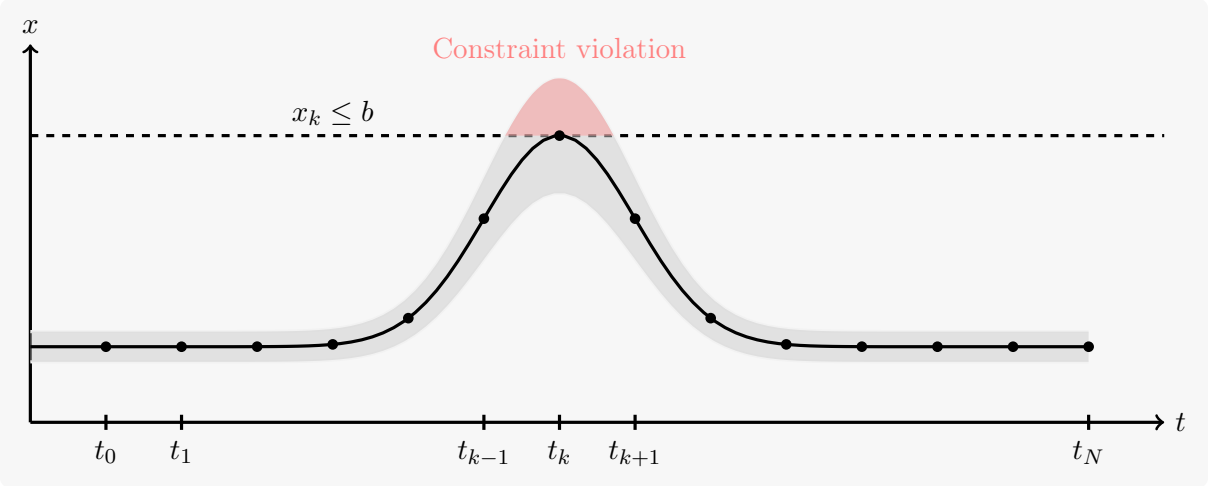

Figure 5.8: Illustration of the effect of uncertainty on the solution to a deterministic trajectory-optimization problem.

5.6.1 Chance Constraints

Constraints enforced on random (uncertain) variables are commonly referred to as *chance* constraints. Chance constraints are stochastic analogues of deterministic inequality constraints. Consider a random variable x with a known probability distribution. In cases where the distribution of x is bounded, classical deterministic inequality constraints (e.g. $x \leq b$) may be used such that $x_{\max} \leq b$, where $x_{\max}$ is the upper bound on x. However, for variables with unbounded distributions, such as Gaussian distributions, no such upper limit exists, as the distribution will never meet the hard constraint. Chance constraints address this by constraining the *probability* of constraint violation. A chance constraint is generically expressed as

$$\Pr\left[\boldsymbol{x} \in \mathcal{S}\right] \geq 1 - \delta \tag{5.12}$$

where $\boldsymbol{x} \in \mathcal{R}^{n_x}$ is a vector of uncertain variables, and $\mathbb{S}$ is the feasible region. The value δ is called the "risk bound" and describes the level of confidence to which the constraint should be satisfied (i.e. $\delta = 10^{-3}$ results in 99.9% constraint satisfaction). Equation 5.12 may be re-written as

$$\Pr\left[\boldsymbol{x} \in \mathcal{S}\right] = \int_{\boldsymbol{x}\in\mathcal{S}} f(\boldsymbol{x})\mathrm{d}\boldsymbol{x} \geq 1 - \delta \tag{5.13}$$

where f is the probability density function for $\boldsymbol{x}$. In general, inclusion of Equation 5.13 within an optimization problem is computationally intractable. However, an equivalent, deterministic form may be obtained if two assumptions are met: 1) the state vectors are normally distributed such that $\boldsymbol{x} \sim \mathcal{N}(\bar{\boldsymbol{x}}, P_x)$, and 2) the feasible regions are affine in the state vectors such that $\mathcal{S} = \left\{\boldsymbol{x} \mid \boldsymbol{a}^T\boldsymbol{x} \leq b\right\}$ (i.e. a half space).

To derive the equivalent, deterministic form of a chance constraint where the state variable is subject to multivariate Gaussian distributions and the feasible region is representable by a half space, recall the definition of the cumulative distribution function for a standard (i.e. zero mean, unit variance), univariate Gaussian random variable γ

$$\Pr\left[\gamma \leq b\right] = \mathrm{cdf}(b) = \int_{-\infty}^{b} f(\gamma)\mathrm{d}\gamma \quad \gamma \sim \mathcal{N}(0,1) \quad . \tag{5.14}$$

Next, notice that $y = \boldsymbol{a}^T\boldsymbol{x}$ is also a univariate Gaussian random variable. It then follows that

$$\mathbb{E}[y] = \bar{y} = \mathbb{E}[\boldsymbol{a}^T\boldsymbol{x}] = \boldsymbol{a}^T\mathbb{E}[\boldsymbol{x}] = \boldsymbol{a}^T\bar{\boldsymbol{x}} \tag{5.15}$$

$$\mathrm{Cov}[y] = \mathbb{E}[(y-\bar{y})(y-\bar{y})^T] = \mathbb{E}[(\boldsymbol{a}^T\boldsymbol{x} - \boldsymbol{a}^T\bar{\boldsymbol{x}})(\boldsymbol{a}^T\boldsymbol{x} - \boldsymbol{a}^T\bar{\boldsymbol{x}})^T] = \boldsymbol{a}^T P_x \boldsymbol{a} \quad . \tag{5.16}$$

Therefore, y is a univariate Gaussian random variable such that $y \sim \mathcal{N}(\boldsymbol{a}^T\bar{\boldsymbol{x}}, \boldsymbol{a}^T P_x \boldsymbol{a})$. Next, recall that any non-standard, univariate Gaussian random variable may be defined in terms of a standard, univariate Gaussian random variable such that

$$y = \boldsymbol{a}^T\bar{\boldsymbol{x}} + \gamma\sqrt{\boldsymbol{a}^T P_x \boldsymbol{a}} \quad \gamma \sim \mathcal{N}(0,1) \quad . \tag{5.17}$$

This allows Equation 5.12 to be written as

$$\Pr\left[\boldsymbol{x} \in \mathcal{S}\right] = \Pr\left[\boldsymbol{a}^T\boldsymbol{x} \leq b\right] = \Pr\left[\boldsymbol{a}^T\bar{\boldsymbol{x}} + \gamma\sqrt{\boldsymbol{a}^T P_x \boldsymbol{a}} \leq b\right] \geq 1 - \delta \tag{5.18}$$

which can be simplified to

$$\Pr\left[\gamma \leq \frac{b - \boldsymbol{a}^T \bar{\boldsymbol{x}}}{\sqrt{\boldsymbol{a}^T P_x \boldsymbol{a}}}\right] \geq 1 - \delta \quad . \tag{5.19}$$

Using the cumulative distribution function, Equation 5.19 can be written as

$$\Pr\left[\gamma \leq \frac{b - \boldsymbol{a}^T \bar{\boldsymbol{x}}}{\sqrt{\boldsymbol{a}^T P_x \boldsymbol{a}}}\right] = \text{cdf}\left(\frac{b - \boldsymbol{a}^T \bar{\boldsymbol{x}}}{\sqrt{\boldsymbol{a}^T P_x \boldsymbol{a}}}\right) \geq 1 - \delta \quad . \tag{5.20}$$

Using the inverse cumulative distribution function, this is simplified to

$$\boldsymbol{a}^T \bar{\boldsymbol{x}} + \text{cdf}^{-1}(1 - \delta)\sqrt{\boldsymbol{a}^T P_x \boldsymbol{a}} \leq b \quad . \tag{5.21}$$

Note that the inverse cumulative distribution function for a univariate, normally-distributed random variable is readily available in most scientific computing languages. Therefore, under the assumptions listed above, an equivalent form of Equation 5.13 is

$$\Pr\left[\boldsymbol{a}^T \boldsymbol{x} \leq b\right], \quad \boldsymbol{x} \sim \mathcal{N}(\bar{\boldsymbol{x}}, P_x) \Leftrightarrow \boldsymbol{a}^T \bar{\boldsymbol{x}} + \text{cdf}^{-1}(1 - \delta)\sqrt{\boldsymbol{a}^T P_x \boldsymbol{a}} \leq b \quad . \tag{5.22}$$

The final obstacle with chance constraints lies in reformulating them into a structure compatible with an SSCP algorithm. Generally, the SSCP method determines sequences of the mean state $\bar{x}_k$ *and* the covariance $P_{x,k}$. Unfortunately, the presence of the square root in Equation 5.22 renders it non-convex with respect to P_x. Chapter 7 addresses this challenge in the development of a stochastic aerocapture guidance algorithm.

5.7 Summary

This chapter outlined the core principles of aerocapture trajectories and the objectives of aerocapture guidance. It began by presenting the equations of motion that govern the simulation of an aerocapture spacecraft's flight, incorporating models for gravitational and aerodynamic forces. Following this, the notion of an aerocapture "entry corridor" was introduced, defining the range of feasible trajectories for a specific problem configuration and offering guidance on the selection of atmospheric-entry conditions. The chapter then provided an explanation of aerocapture guidance, featuring an overview of the state-of-the-art FNPAG algorithm.

The chapter then explored the mathematical foundations of convex optimization and detailed the application of sequential convex programming (SCP) to trajectory optimization. This included fundamental definitions and simple tutorials on the usage of the SCP techniques employed throughout the remainder of this dissertation. The discussion extended the SCP framework from deterministic to stochastic optimization problems with a motivating example that highlighted the benefit of including uncertainty explicitly in the optimization process. Finally, the use of chance constraints to enforce probabilistic constraints on uncertain variables was introduced, and the challenges of including chance constraints within an SCP framework were identified.

Chapter 6

Convex Predictor-Corrector Aerocapture Guidance

"The necessary number of iterations is one more than the number you have currently done. This is true at any point in time."

— Akin's Third Law

6.1 Introduction

This chapter presents the convex predictor-corrector aerocapture guidance (CPAG) algorithm. It seeks to address several of the shortcomings in the current state-of-the-art NPC aerocapture algorithms identified in Chapter 5 by solving a convex optimization problem in the correction phase. Adopting a more advanced optimization technique allows for more complex control profiles to be generated. This removes the need for separate lateral logic to determine the bank-angle sign and allows for system-level path constraints (such as heating and loading) to be considered directly in the correction step. CPAG explicitly enforces constraints on the peak aerodynamic loading and integrated convective heat load, while existing aerocapture guidance algorithms cannot.

The chapter begins by deriving the convex subproblem solved iteratively in CPAG's correction phase, relying on the physical models and mathematical principles outlined in Chapter 5. Next, CPAG's performance is evaluated against the Fully Numerical Predictor-Corrector Aerocapture Guidance (FNPAG) algorithm through a series of Monte Carlo analyses, spanning a variety of planetary destinations and aerocapture vehicle configurations. A broad assessment of CPAG's robustness across multiple scenarios provides deeper insight into its capabilities and highlights conditions that may push its limits. Results are presented for aerocapture at Neptune, Mars, and Earth,

though the algorithm's framework is adaptable to any feasible destination. The chapter concludes with an analysis of an energy-based apoapsis targeting approach, which accounts for non-Keplerian motion between atmospheric exit and apoapsis.

6.2 Background

It is the role of the onboard guidance system to determine the required aerodynamic lift and/or drag force vector orientation during aerocapture, so that the spacecraft can exit the atmosphere and enter its target orbit without consuming a large amount of propellant to perform orbital-correction maneuvers. Current state-of-the-art guidance algorithms employ a fully-numerical, predictor corrector (NPC) approach [52] that has several advantages over heritage guidance systems. These include: the ability to adapt to large trajectory dispersions, no reliance on a pre-planned reference trajectory, and no separate reference trajectory tracking laws. The successful use of NPC guidance, highlighting these benefits, was recently demonstrated during the Artemis 1 skip entry [53]. Despite never being flown, predictor-corrector aerocapture guidance algorithm development has seen substantial interest over the years. Skip entry and aerocapture are relatively similar, and Orion's skip-entry guidance algorithm was also extended to handle aerocapture scenarios [189]. Later, the fully numerical predictor corrector aerocapture guidance (FNPAG) algorithm [55] was introduced, in which a bang-bang bank angle profile was proposed in order to minimize the ΔV required to correct the final orbit. FNPAG has been used for numerous aerocapture studies [56, 57, 58, 59, 174] and is considered to be the state of the art. Further details on FNPAG are found in Chapter 5.

The hypersonic environments encountered during aerocapture produce severe, often highly uncertain, aerodynamic heating. Because the thermal protection system (TPS) must be designed with these uncertainties in mind, the consideration of aerothermodynamic constraints is integral to mission success. If the guidance system can directly consider aerothermodynamic constraints in its decision-making process, a more mass optimal TPS could be designed. Such a capability would be particularly beneficial for highly mass-constrained missions like Ice Giant aerocapture. The benefit of NPC guidance for Ice Giant aerocapture has been demonstrated in recent work

[190, 58, 169], but none of the analyses directly enforce path constrains in the prediction-correction process. An entry vehicle's aerothermodynamic environment is tightly coupled to its trajectory, so the inclusion of aerodynamic loading and heating constraints would be a key characteristic of an *advanced* aerocapture guidance algorithm. However, the optimal aerocapture guidance problem is highly non-convex, making it difficult to solve even without constraints. Because of this, the FNPAG guidance law is derived using only the longitudinal equations of motion and excludes path constraints like aerodynamic loading and heating. While these simplifications make the problem easier to analyze, they prohibit any explicit consideration of out-of-plane targets (e.g. inclination, RAAN) and aerothermodynamic constraints. A method for augmenting the FNPAG control profile to compensate for path constraints is presented in Ref. 52, however, it effectively negates the optimality statement of the control profile originally derived in Ref. 55. The assumptions made by FNPAG do have some benefit, as they produce a one-dimensional, parametric control profile that is assumed by the guidance system. This reduces the computational complexity of the guidance correction phase, as only a single parameter needs to be optimized. Specifically, FNPAG solves a single-variable root-finding problem to obtain either a bank-angle switching time or a constant bank-angle value [55]. While this simplifies the problem to be solved onboard, it makes incorporating features such as finite-rate bank-angle maneuvers and reversals more challenging. In the case of bank-angle modulation, a separate lateral logic algorithm is required to determine the sign of the bank angle (either positive or negative), as the longitudinal channel of the guidance system only computes the bank-angle magnitude [176]. For aerocapture, the lateral logic is used to help target an orbit with a specific inclination (or wedge angle if a specific right ascension of the ascending node is known). If the bank-angle sign has to be reversed, the vehicle must fly an open loop control profile until the reversal is completed which makes it more susceptible to dispersions and accumulation of error. As such, it is desirable for an aerocapture guidance system to be capable of determining the bank-angle magnitude *and* direction simultaneously.

Most aerospace optimal control problems are non-convex due to their nonlinear dynamics and complex constraints. As highlighted in Chapter 5, there has been a growing interest in the

use of convex-optimization [177] based algorithms due to their computational efficiency and ability to solve a wide range of problems [179, 180]. Specifically, sequential convex programming (SCP) takes an iterative *convexify-and-solve* process to march toward a solution of the original non-convex problem [182]. For real-time, onboard guidance, SCP presents an attractive option, as modern interior point methods (IPM) are capable of solving the convex problem in polynomial time [178]. Reynolds et al. [191] have shown that problem-specific structure can be exploited enabling real-time solution capabilities, and the Masten Space Systems Xombie demonstrated the feasibility of SCP on spaceflight hardware [60]. Such demonstrations motivate the use of SCP within other real-time guidance applications such as entry and aerocapture. Ridderhof et al. [74] presented a chance-constrained aerocapture guidance algorithm that optimized a series of feedback control gains using SCP. They demonstrated the effective use of a feedback control policy for a Mars aerocapture scenario under atmospheric density uncertainty, but they did not consider out-of-plane motion or system-level constraints like aerodynamic heating and loading. Recent work by Tracy et al. [192, 193] has investigated the use of SCP within the correction step of an NPC *entry*-guidance algorithm, eliminating the need for a separate lateral guidance law common to existing NPC entry-guidance algorithms. The work presented in [192, 193] did consider system-level constraints in the mathematical formulation of the NPC algorithm, although explicit satisfaction of the constraints was not demonstrated. While entry and aerocapture are different, they both require the accurate guidance of a hypersonic entry vehicle to a terminal target. Therefore, it is of interest to use SCP within an NPC *aerocapture* guidance algorithm.

6.3 Convex-Programming Formulation

Aerocapture trajectories involve highly nonlinear dynamics and both terminal and path constraints. The nonlinear, non-convex optimization aerocapture trajectory optimization problem investigated in this work is expressed as the following free time-of-flight problem

$$\min_{\boldsymbol{x}\in\mathcal{X},u\in\mathcal{U},t_f\geq 0} J(\boldsymbol{x}(t),u(t)) \quad \text{s.t.} \tag{6.1a}$$

$$\dot{\boldsymbol{x}}(t) = \boldsymbol{f}(\boldsymbol{x}(t),u(t)) \tag{6.1b}$$

$$r_a(\boldsymbol{x}(t_f),t_f) = r_{a,\text{target}} \tag{6.1c}$$

$$i(\boldsymbol{x}(t_f),t_f) = i_{\text{target}} \tag{6.1d}$$

$$a(t) \leq a_{\max} \tag{6.1e}$$

$$Q(t_f) \leq Q_{\max} \tag{6.1f}$$

$$u_{\min} \leq u(t) \leq u_{\max} \tag{6.1g}$$

$$\boldsymbol{x}(t_0) = \boldsymbol{x}_0 \quad . \tag{6.1h}$$

Here, $\boldsymbol{x} = [r,\ \theta,\ \phi,\ v,\ \gamma,\ \psi,\ \sigma]$ is the state, the control u is the bank-angle rate, $\boldsymbol{f}(\boldsymbol{x}(t),u(t))$ models the nonlinear dynamics (equations of motion) of the vehicle given in Eq. 5.1, r_a is the apoapsis radius, i is the orbital inclination, $Q(t_f) = \int_{t_0}^{t_f} \dot{q}(\tau)\mathrm{d}\tau$ is the convective heat load obtained by integrating the convective heat flux $\dot{q}(t)$, and $a = \frac{1}{g_0}\sqrt{L^2+D^2}$ is the aerodynamic load factor. Radiative heating is not explicitly considered here because its severity is highly dependent on the entry conditions and atmospheric composition of the destination. However, the optimization framework described herein is easily amenable to its inclusion if radiative heating is expected to be high and an acceptably accurate radiative-heating model is available[1] . The final inclination is expressed as a function of the final state as

$$\cos i_f = \cos\phi_f \cos\psi_f \quad . \tag{6.2}$$

The subscript $(\cdot)_f$ is used to denote a quantity at the final time $t = t_f$, which is taken to be the time at which the vehicle exits the planet's atmosphere. Instead of directly computing the final inclination using an inverse-cosine operation, this work simply uses Eq. 6.2 to enforce the terminal inclination constraint, as the inclination is defined over the range $i \in [0,\pi]$

$$\cos i_f - \cos i_{\text{target}} = 0 \quad . \tag{6.3}$$

[1] At the time of this dissertation, a radiative-heating correlation does not exist for Ice Giant destinations.

In many aerocapture studies [55, 169, 174, 194], the terminal apoapsis target is expressed by directly computing the apoapsis radius as a function of the final state using Equation 5.8. Instead, this work targets a specific apoapsis radius using the energy-like expression

$$\varepsilon(\boldsymbol{x}_f) = \frac{\mu}{r_f}\left[1 - \frac{J_2}{2}\left(\frac{R_e}{r_f}\right)^2\left(3\sin^2\phi_f - 1\right)\right] - \frac{v_f^2}{2} - \frac{\mu}{r_{a,\text{target}}}\left[1 - \frac{J_2}{2}\left(\frac{R_e}{r_{a,\text{target}}}\right)^2\left(3\sin^2\phi_{\text{target}} - 1\right)\right] + \frac{r_f^2 v_f^2 \cos^2\gamma_f}{2r_{a,\text{target}}^2} \quad . \quad (6.4)$$

The form of Eq. 6.4 is derived by equating the specific energy at the atmospheric exit state to the specific energy at the target apoapsis radius. The velocity at the target apoapsis radius is derived by assuming that the specific angular momentum $h = rv\cos\gamma$ is conserved from the atmospheric exit state to the target apoapsis radius. The standard Keplerian gravitational potential μ/r is expanded to account for perturbations due to J_2, which cannot be done with the typical two-body expression in Eq. 5.8. The target latitude ϕ_{target} is an additional constant related to the target inclination and argument of periapsis. As will be shown in Section 6.5, the apoapsis targeting performance is found to be relatively insensitive to the value of ϕ_{target}, so precise selection of this value is not necessary. This work assumes a constant value of $\phi_{\text{target}} = 0$ for all simulations.

The apoapsis target is enforced as a nonlinear equality constraint, which cannot be included directly in a convex optimization setting. A common approach for converting non-convex constraints into an approximate convex form is through local linearization using a first-order Taylor series expansion. This process requires the computation of the constraint's first order sensitivities (i.e. Jacobian) about a reference point. The form of Eq. 5.8 presents a challenge, as the Jacobian of Eq. 5.8 with respect to the final state $\boldsymbol{x}_f = [r_f, \theta_f, \phi_f, v_f, \gamma_f, \psi_f, \sigma_f]$ diverges to infinity when the exit velocity equals the local escape velocity. The specific components of the Jacobian that contain the singularity are those corresponding to the final velocity v_f and flight-path angle γ_f. However, as shown in Fig. 6.1, the magnitude of the Jacobian of Eq. 6.4 is well-defined for all exit-velocity values, whereas, the Jacobian computed using Eq. 5.8 diverges when $v_f = v_{\text{esc}}$.

The increasingly large Jacobian magnitude shown in Fig. 6.1a produces difficulties for numerical optimizers. For example, if an optimizer were given an initial reference trajectory that

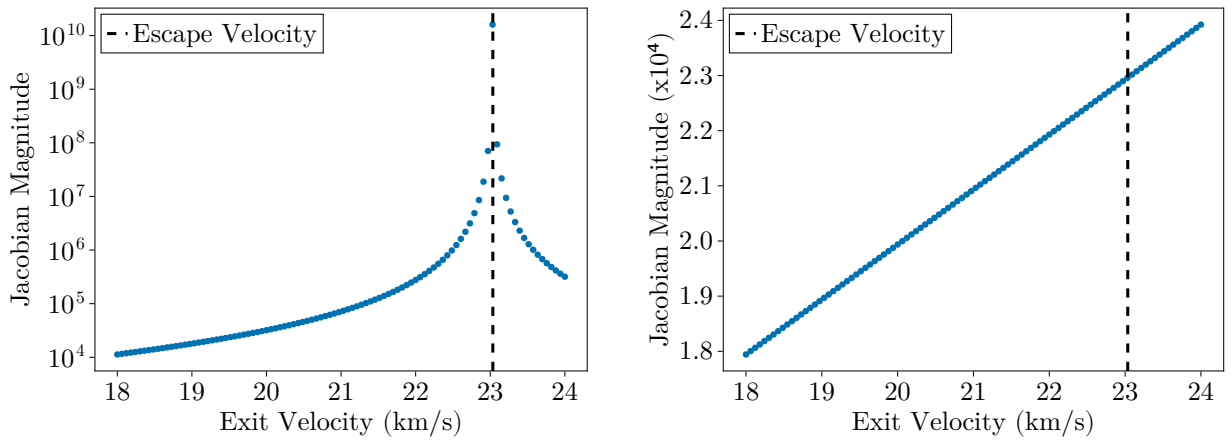

(a) Eq. 5.8.

(b) Eq. 6.4.

Figure 6.1: Jacobian of apoapsis targeting constraint with respect to the final state computed using Eq. 5.8 (a) and Eq. 6.4 (b).

remained hyperbolic (i.e. $v_f > v_{\text{esc}}$), then it would be unlikely to converge on a control profile that would produce a captured trajectory that meets the apoapsis targeting constraint. Physically, this issue is consistent with two-body orbital mechanics in which the apoapsis radius is not defined for hyperbolic orbits. Using an energy-based expression for the terminal apoapsis constraint alleviates this issue, as energy and angular momentum are well-defined for any orbit. For aerocapture, it is assumed that $r_f = r_0$, which is known, so the only unknowns in Equations 5.8 and 6.4 are v_f and γ_f. This is possible because t_0 and t_f correspond to the atmospheric entry and exit times, respectively. While the energy-based expression does not directly solve for the apoapsis radius as a function of the exit state, the values of v_f and γ_f that correspond to the root of Eq. 6.4 are identical to those obtained when setting Eq. 5.8 equal to $r_{a,\text{target}}$ if $J_2 = 0$. When $J_2 \neq 0$, the values of v_f and γ_f determined using Eq. 6.4 will differ in order to account for the effects of J_2 on the exo-atmospheric trajectory from atmospheric exit to apoapsis. This behavior is explored in more depth in the Effects of Planetary Oblateness Section (Sec. 6.5).

In many cases, the ΔV requirements of post-aerocapture orbits produced using an apoapsis targeting approach and a direct ΔV minimization approach are equivalent [54]. This is particularly

the case for highly elliptical target orbits because the periapsis raise maneuver becomes increasingly small as the target apoapsis radius increases. However, in cases where the periapsis raise maneuver dominates the in-plane ΔV requirements, simply targeting an apoapsis radius will not produce ΔV optimal results. A typical expression for computing the required in-plane ΔV is

$$\Delta V_1 + \Delta V_2 = \sqrt{2\mu}\left(\left|\sqrt{\frac{1}{r_a} - \frac{1}{r_a + r_p^*}} - \sqrt{\frac{1}{r_a} - \frac{1}{2a}}\right| + \left|\sqrt{\frac{1}{r_p^*} - \frac{1}{r_a^* + r_p^*}} - \sqrt{\frac{1}{r_p^*} - \frac{1}{r_a + r_p^*}}\right|\right) \quad . \tag{6.5}$$

Here, r_a and r_p are the apoapsis radius and periapsis radius, respectively, determined from the final state, and r_a^* and r_p^* are the corresponding target values. The version of FNPAG in [55] also presents an option to directly minimize Eq. 6.5, however, this work targets an orbit using an apoapsis radius constraint. Equation 6.5 presents two challenges for the proposed SCP framework. First, its Jacobian presents difficulties for numerical optimizers in the same way Eq. 5.8 does. Deriving a singularity-free expression for ΔV minimization is left as an open area of research. Second, it assumes that the vehicle exits the atmosphere on a captured, *elliptical* orbit and is therefore undefined if the predicted trajectory remains hyperbolic. Due to dispersions along the trajectory, it is possible that the predicted trajectory remains hyperbolic, especially if the target orbit is highly elliptical. While evaluating Eq. 6.5 at exit conditions that correspond to a hyperbolic orbit will return a non-zero value, that value has no physical meaning. Therefore, this work introduces a term that encourages the maximization of the resulting periapsis radius in order to minimize the periapsis raise ΔV

$$\frac{1}{r_p(\boldsymbol{x}_f)} = \left[a\left(1 - \sqrt{1 - \frac{v_f^2 r_f^2 \cos^2\gamma_f}{\mu a}}\right)\right]^{-1} \quad . \tag{6.6}$$

The maximization of the resulting periapsis radius will reduce the in-plane ΔV, as the optimization problem constrains the apoapsis radius to match the target. Note that Eq. 6.6 is well-defined for any orbit. Minimizing the periapsis-raise ΔV becomes increasingly important for low target apoapsis radii (e.g. circular orbit), as the periapsis-raise ΔV is typically smaller than the apoapsis-correction ΔV for highly-elliptical target orbits. For this reason, Eq. 6.6 is omitted from the cost function for highly-elliptical target orbits.

The cost function in this work is defined as

$$J(\boldsymbol{x}, u, t) = \alpha_1 \frac{1}{r_p(\boldsymbol{x}(t_f))} + \alpha_2 \int_{t_0}^{t_f} \|u(t)\|_2^2 \, \mathrm{d}t = \alpha_1 \frac{1}{r_p(\boldsymbol{x}(t_f))} + \alpha_2 \int_{t_0}^{t_f} \|\dot{\sigma}(t)\|_2^2 \, \mathrm{d}t \quad . \tag{6.7}$$

Here, the α_i are weighting terms that can be used to scale the cost-function terms accordingly. Additional discussion on cost-function scaling is presented in Section 6.4. The $\mathcal{L}$-2 norm squared of the bank-angle rate is chosen in order to encourage smoothly varying bank-angle profiles. As such, the full nonlinear, non-convex optimization problem to be solved in this work may be stated as

Problem 1. (Nonlinear optimal control problem)*: Determine the optimal* $u(t)$*,* $\boldsymbol{x}(t)$*, and* t_f *that minimize the cost* (6.7) *subject to equality constraints on the dynamics* (6.1b)*, apoapsis radius* (6.4)*, inclination* (6.3)*, and initial condition* (6.1h)*, and inequality constraints on the load factor* (6.1e)*, heat load* (6.1f)*, and control magnitude* (6.1g)*.*

While obtaining a solution to Problem 1 is possible, it cannot be computed rapidly due to the lack of convexity introduced by the nonlinear dynamics and complex path/terminal constraints. Therefore, CPAG replaces Problem 1 with an approximate convex model, called a *convex subproblem*, that is solved sequentially in the prediction-correction process. In general, the sequential convex programming procedure may be summarized as follows. First, a user-supplied initial guess to the trajectory is input. This initial guess is directly supplied by the prediction phase of the algorithm in which the current control plan is integrated forward in time until the current guess of the final time is reached. Any non-convexities in Problem 1 are removed by a local linearization about the reference trajectory [182], while the convex elements are retained. Next, the continuous-time, free final time subproblem is converted into a finite-dimensional, fixed final time convex optimization problem by discretizing the time domain and introducing a time dilation parameter. The convex subproblem can then be solved using a convex optimization algorithm, and the reference control is updated using the solution to the subproblem. The optimized control then becomes the reference for the next prediction, and this process is repeated until convergence. The sequential convex optimization framework described in this work is generally based on the penalized trust region (PTR) [183] algorithm, but it is tailored to be used within a predictor-corrector guidance algorithm framework

instead of a pure trajectory-optimization framework. In the PTR algorithm, the trust-region radii are added to the cost function which allows the optimizer to expand/contract the trust regions accordingly and avoids otherwise complex trust-region logic. Figure 6.2 illustrates the general SCP process, and each of these steps is discussed in further detail for the remainder of this section.

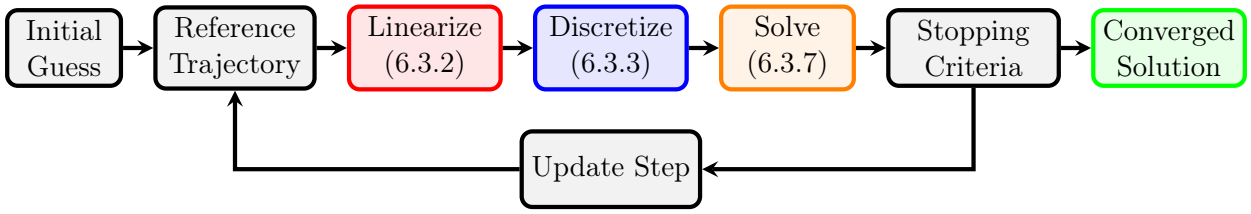

Figure 6.2: Flowchart of the sequential convex programming (SCP) algorithm.

6.3.1 Problem Reformulation

Before converting the nonlinear optimal control problem in Problem 1 into an approximate convex subproblem, a slight reformulation is proposed to reduce the risk of constraint violations for the heat load and load factor constraints. Specifically, the state $\boldsymbol{x}$ is augmented with the heat load and square of the load factor

$$\tilde{\boldsymbol{x}}(t) = [\boldsymbol{x}(t),\ \tilde{a}(t),\ Q(t)] \tag{6.8}$$

where $\tilde{\boldsymbol{x}}$ is referred to as the "augmented" state, and $\tilde{a} = a^2$. Because $a \geq 0$, constraining a^2 is equivalent to constraining a, so the square of the load factor is chosen due to its simpler Jacobian. The augmented state dynamics may then be written as

$$\dot{\tilde{\boldsymbol{x}}}(t) = \begin{bmatrix} \dot{\boldsymbol{x}}(t) \\ \dot{\tilde{a}}(t) \\ \dot{Q}(t) \end{bmatrix} = \begin{bmatrix} \boldsymbol{f}(\boldsymbol{x}(t), u(t)) \\ \dfrac{\partial \tilde{a}(\boldsymbol{x}(t))}{\partial t} \\ \dot{q}(\boldsymbol{x}(t)) \end{bmatrix} = \boldsymbol{F}(\boldsymbol{x}(t), u(t)) \tag{6.9}$$

where $\dot{q}(\boldsymbol{x}(t))$ is the instantaneous heat flux, and the partial derivative $\dfrac{\partial \tilde{a}}{\partial t}$ may be obtained using Eqs. 5.1, 5.3, and 5.4. The form of Eqs. 6.8 and 6.9 is chosen for two reasons. First, computing the heat load requires integrating the heat flux along the entire trajectory. While convex approximations of integral quantities exist (e.g. trapezoidal integration [182]), the linearization and discretization

procedure discussed in the following sections allows for a more exact enforcement of the heat load constraint in the convex subproblem. Because the reference trajectory, about which the convex subproblem is formed, is produced by numerically integrating the nonlinear equations of motion, a more accurate description of how the heat load changes along the corrected trajectory is obtained. Finally, expressing the path constraint variables as an additional state variable allows the proposed framework to be amenable to continuous time constraint satisfaction through the selection of an appropriate penalty function [195]. Following this reformulation, Eq. 6.1 may be equivalently stated as

$$\min_{\tilde{\boldsymbol{x}}\in\mathcal{X},u\in\mathcal{U}} J(\tilde{\boldsymbol{x}}(t),u(t)) \quad \text{s.t.} \tag{6.10a}$$

$$\dot{\tilde{\boldsymbol{x}}}(t) = \boldsymbol{F}(\tilde{\boldsymbol{x}}(t),u(t)) \tag{6.10b}$$

$$\boldsymbol{g}(\tilde{\boldsymbol{x}}(t_0),\tilde{\boldsymbol{x}}(t_f)) = 0 \tag{6.10c}$$

$$\boldsymbol{h}(\tilde{\boldsymbol{x}}(t),u(t)) \leq 0 \tag{6.10d}$$

where $\boldsymbol{g}\in\mathbb{R}^3$ enforces the apoapsis (Eq. 6.4), inclination (Eq. 6.2), and initial condition equality constraints (Eq. 6.1h), and $\boldsymbol{h}\in\mathbb{R}^3$ enforces the load factor (Eq. 6.1e), heat load (Eq. 6.1f), and control inequality constraints (Eq. 6.1g). For a more concise presentation, the tilde notation denoting the augmented state is dropped for the remainder of the chapter.

6.3.2 Linearization

In a predictor-corrector algorithm, the time of flight is not known *a priori*, so the optimizer needs to solve for it. To include the time of flight t_f in the optimization problem, the time domain is scaled to be in the non-dimensional time interval $\tau\in[0,1]$, where τ is defined such that $\tau = t/t_f = t/p$. Using the chain rule, the augmented-state dynamics with respect to the non-dimensional time become

$$\frac{\mathrm{d}\boldsymbol{F}}{\mathrm{d}\tau} = \overset{\circ}{\boldsymbol{x}}(\tau) = \frac{\mathrm{d}\boldsymbol{F}}{\mathrm{d}t}\frac{\mathrm{d}t}{\mathrm{d}\tau} = t_f\boldsymbol{F}(\boldsymbol{x}(t),u(t)) = p\boldsymbol{F}(\boldsymbol{x}(t),u(t)) \tag{6.11}$$

where $\overset{\circ}{(\cdot)}$ denotes the derivative with respect to the non-dimensional time. Equation 6.11 is often referred to as *time dilation* and is used to convert a free final time problem into an equivalent fixed

final time problem [183].

The next step in converting Eq. 6.10 into a convex form is to linearize the nonlinear dynamics using a first-order Taylor-series expansion

$$\mathring{\boldsymbol{x}}(\tau) \approx A(\tau)\boldsymbol{x}(\tau) + B(\tau)u(\tau) + \Sigma(\tau)p + \boldsymbol{z}(\tau) \tag{6.12}$$

where $\boldsymbol{x} \in \mathbb{R}^{n+2}$ is the augmented state vector, $u \in \mathbb{R}$ is the bank-angle rate, $p \in \mathbb{R}^p$ represents any problem-dependent parameters, A is the state zero-input dynamics matrix, B is the control matrix, Σ is the parameter matrix, and $\boldsymbol{z}$ is a drift term that is produced by the linearization of the nonlinear equations of motion. In this work, the parameter p is taken to be the final time of flight $t_f \in \mathbb{R}_{++}$. Using Eq. 6.11, the matrices in Eq. 6.12 are computed using a first-order Taylor series expansion about a reference trajectory $(\bar{\boldsymbol{x}}(\tau), \bar{u}(\tau))$ and time of flight $\bar{p}$

$$A(\tau) = \bar{p}\left.\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{x}}\right|_{\bar{\boldsymbol{x}},\bar{u},\bar{p}} \tag{6.13}$$

$$B(\tau) = \bar{p}\left.\frac{\partial \boldsymbol{F}}{\partial u}\right|_{\bar{\boldsymbol{x}},\bar{u},\bar{p}} \tag{6.14}$$

$$\Sigma(\tau) = \left.\frac{\partial \boldsymbol{F}}{\partial p}\right|_{\bar{\boldsymbol{x}},\bar{u},\bar{p}} = \boldsymbol{F}(\bar{\boldsymbol{x}}(\tau), \bar{u}(\tau)) \tag{6.15}$$

$$\boldsymbol{z}(\tau) = -A(\tau)\bar{\boldsymbol{x}}(\tau) - B(\tau)\bar{u}(\tau) \quad . \tag{6.16}$$

Next, the non-convex terminal equality constraints are linearized about the reference. Using Eq. 6.2, the terminal inclination constraint is approximated as

$$\cos\left(i(\boldsymbol{x}_f)\right) \approx \left.\frac{\partial \cos(i)}{\partial \boldsymbol{x}}\right|_{\text{ref}} \cdot (\boldsymbol{x}_f - \bar{\boldsymbol{x}}_f) + \cos\left(i(\bar{\boldsymbol{x}}_f)\right) = \cos\left(i_{\text{target}}\right) \tag{6.17}$$

where the subscript $(\cdot)_{\text{ref}}$ is used to denote the quantity being evaluated at the reference $(\bar{\boldsymbol{x}}, \bar{u}, \bar{p})$. Similarly, using Eq. 6.4, the terminal apoapsis constraint is approximated as

$$\varepsilon(\boldsymbol{x}_f) \approx \left.\frac{\partial \varepsilon}{\partial \boldsymbol{x}}\right|_{\text{ref}} \cdot (\boldsymbol{x}_f - \bar{\boldsymbol{x}}_f) + \varepsilon(\bar{\boldsymbol{x}}_f) = 0 \quad . \tag{6.18}$$

Note that these constraints depend only on the state $\boldsymbol{x}$, so no additional partial derivatives with respect to the control u or problem parameters p are required. Furthermore, because the heat load and load factor are included in the augmented state definition, the Jacobian of each is automatically

captured in Eq. 6.12. Further details on each of the Jacobians described in this section are found in Appendix B.1.

6.3.3 Discretization

To convert the continuous-time, infinite-dimensional optimal control problem into an equivalent parameter optimization problem that can be solved by convex optimization algorithms, the continuous expressions from the previous section must be converted into discrete-time equivalents. To do so, a set of N temporal nodes is chosen such that

$$0 = \tau_1 < \tau_2 < \ldots < \tau_{N-1} < \tau_N = 1 \qquad . \tag{6.19}$$

The spacing of the temporal nodes τ_k is up to the user, and this work elects to distribute them uniformly across the interval $[0, 1]$. Recall that by introducing the time dilation parameter $p = t_f$, the free-time-of-flight problem can always be discretized over the domain $\tau \in [0, 1]$. While the values of τ_k remain fixed, an adaptive mesh procedure could be implemented using a technique known as "generalized time dilation" [195] which expresses the time-dilation factor as a continuous-time control input instead of a fixed parameter. At each node, let $\Phi(\tau_{k+1}, \tau_k) : \mathbb{R} \mapsto \mathbb{R}^{(n+2)\times(n+2)}$ denote the state transition matrix (STM) for the system in Eq. 6.12. From linear systems theory, the STM evolves according to

$$\overset{\circ}{\Phi}(\tau, \tau_0) = A(\tau)\Phi(\tau, \tau_0), \qquad \Phi(\tau_0, \tau_0) = I_{(n+2)\times(n+2)} \qquad . \tag{6.20}$$

The control is discretized using a first-order hold [196, 197, 183], which assumes a linear variation between temporal nodes k and $k+1$

$$u(\tau) = \sigma_k^-(\tau)u(\tau_k) + \sigma_k^+(\tau)u(\tau_{k+1}) \qquad \forall\ \tau \in [\tau_k, \tau_{k+1}) \tag{6.21}$$

where $\sigma_k^-(\tau)$ and $\sigma_k^+(\tau)$ are defined as

$$\sigma_k^-(\tau) = \frac{\tau_{k+1} - \tau}{\tau_{k+1} - \tau_k} \qquad \sigma_k^+(\tau) = \frac{\tau - \tau_k}{\tau_{k+1} - \tau_k} \qquad . \tag{6.22}$$

Utilizing the STM, the continuous-time, linear dynamics of the augmented state can now be transformed into an equivalent discrete-time, linear form such that

$$\boldsymbol{x}_{k+1} = A_k \boldsymbol{x}_k + B_k^- u_k + B_k^+ u_{k+1} + \Sigma_k p + \boldsymbol{z}_k \quad k = 1, \ldots, N-1 \tag{6.23}$$

where

$$\begin{aligned}
A_k &= \Phi(\tau_{k+1}, \tau_k) && \text{(6.24a)} \\
B_k^- &= A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k) \sigma_k^-(\xi) B(\xi) \mathrm{d}\xi && \text{(6.24b)} \\
B_k^+ &= A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k) \sigma_k^+(\xi) B(\xi) \mathrm{d}\xi && \text{(6.24c)} \\
\Sigma_k &= A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k) \Sigma(\xi) \mathrm{d}\xi && \text{(6.24d)} \\
\boldsymbol{z}_k &= A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k) \boldsymbol{z}(\xi) \mathrm{d}\xi \quad . && \text{(6.24e)}
\end{aligned}$$

Here ξ is a dummy variable introduced for the integration of Eq. 6.24 over the interval $\tau \in [\tau_k, \tau_{k+1}]$. For more details on the implementation aspects of this discretization procedure, the reader is referred to Appendix B.2.

The process of linearization and discretization may lead to a solution that is not dynamically feasible, meaning that the solution does not obey the full nonlinear dynamics. This is typically mitigated by introducing unconstrained slack variables called "virtual control" to the discrete, linear dynamics to relax the requirement for the solution to be dynamically feasible [198]. These slack variables are then heavily penalized in the subproblem cost function to encourage the converged solution to be dynamically feasible. This approach is particularly useful for problems in which the initial reference trajectory is not dynamically feasible [182].

Instead of relaxing the dynamics constraints using virtual control, this work takes an approach made possible by the nature of predictor-corrector guidance algorithms. The prediction phase of the guidance algorithm produces a dynamically feasible trajectory $\bar{\boldsymbol{x}}(t)$ that acts as the reference about which the problem is linearized. Note that the "reference" trajectory described here is not the same as a pre-planned reference trajectory used in some guidance algorithms, but is rather

a reference for the local linearization procedure. Instead of solving for a sequence of full state $\boldsymbol{x}_k$ and control u_k vectors and final time t_f at each iteration, the proposed algorithm solves for corrections $\delta\boldsymbol{x}_k = \boldsymbol{x}_k - \bar{\boldsymbol{x}}_k$, $\delta u_k = u_k - \bar{u}_k$, and $\delta p = p - \bar{p}$ about the reference values. Because the reference is dynamically feasible, a trivial solution of all zeros for $\delta\boldsymbol{x}_k$, δu_k, and δp always exists, which eliminates concerns of the optimizer not converging on a dynamically feasible trajectory [192]. This is a crucial characteristic of a robust on-board guidance algorithm. Therefore, the constraints describing the dynamics in Eq. 6.23 are reformulated as

$$\delta\boldsymbol{x}_{k+1} = A_k\delta\boldsymbol{x}_k + B_k^-\delta u_k + B_k^+\delta u_{k+1} + \Sigma_k\delta p \quad \forall\ k = 1,\ldots,N-1 \quad . \tag{6.25}$$

Next, the apoapsis and inclination targeting constraints are approximated as

$$R\cdot\delta\boldsymbol{x}_N + \cos\left(i(\bar{\boldsymbol{x}}_N)\right) - \cos\left(i_{\text{target}}\right) = 0 \tag{6.26}$$

$$H\cdot\delta\boldsymbol{x}_N + \varepsilon(\bar{\boldsymbol{x}}_N) = 0 \tag{6.27}$$

where $R = \left.\dfrac{\partial\cos(i)}{\partial\boldsymbol{x}}\right|_{k=N}$, and $H = \left.\dfrac{\partial\varepsilon}{\partial\boldsymbol{x}}\right|_{k=N}$. Similarly, the periapsis minimization term in Eq. 6.6 is approximated via

$$\frac{1}{r_p(\boldsymbol{x}_f)} \approx P\cdot\delta\boldsymbol{x}_N + \frac{1}{r_p(\bar{\boldsymbol{x}}_N)} \tag{6.28}$$

where $P = \left.\dfrac{\partial(1/r_p)}{\partial\boldsymbol{x}}\right|_{k=N}$. Because the load factor and heat load are included in the augmented state, the constraints on the load factor and heat load are expressed simply as

$$\boldsymbol{E}_{n+1}\cdot(\bar{\boldsymbol{x}}_k + \delta\boldsymbol{x}_k) \quad\leq\quad a_{\max} \quad \forall\ k = 1,\ldots,N \tag{6.29}$$

$$\boldsymbol{E}_{n+2}\cdot(\bar{\boldsymbol{x}}_N + \delta\boldsymbol{x}_N) \quad\leq\quad Q_{\max} \tag{6.30}$$

where $\boldsymbol{E}_i \in \mathbb{R}^{n+2}$ is a vector to index the i-th entry in the state vector. While the use of a dynamically feasible reference trajectory alleviates concerns of dynamic infeasibility in each subproblem, it does not guarantee that the load factor and heat load path constraints are satisfied. Therefore, non-negative slack variables are introduced such that

$$\boldsymbol{E}_{n+1}\cdot(\bar{\boldsymbol{x}}_k + \delta\boldsymbol{x}_k) - a_{\max} \leq \nu_a \quad \forall\ k = 1,\ldots,N \tag{6.31}$$

$$\boldsymbol{E}_{n+2} \cdot (\bar{\boldsymbol{x}}_N + \delta \boldsymbol{x}_N) - Q_{\max} \leq \nu_Q \tag{6.32}$$

and ν_a and ν_Q are appended to the cost function. Finally, the initial condition and bank-angle-rate-magnitude constraints may be expressed as

$$\|\bar{u}_k + \delta u_k\|_2 \leq u_{\max} \quad \forall\ k = 1, \ldots, N \tag{6.33}$$

$$\delta \boldsymbol{x}_1 = 0 \quad . \tag{6.34}$$

6.3.4 Improved Balancing of In-Plane and Out-of-Plane Accuracy

The current treatment of the optimal control problem only penalizes deviations in the apoapsis radius and inclination at the final node $t_k = t_f$. This is sufficient for the apoapsis-targeting problem, however, it presents a challenge if high-accuracy final apoapsis *and* final inclination are desired. In early implementations of CPAG, achieving a balance between in-plane and out-of-plane accuracy proved challenging, often requiring a trade-off where in-plane accuracy was compromised to ensure sufficiently precise out-of-plane metrics [71]. When only penalizing the final inclination error, CPAG produces bank-angle profiles that contain a single switch in the bank-angle sign from positive to negative. This, in turn, produces an inclination profile like the black line in Figure 6.3 that initially deviates from the desired value and then asymptotically approaches the target at the end of the trajectory. While optimal from a control-effort perspective, it provides no margin for dispersions along the trajectory. Physically, the need for balancing the two performance metrics is a product of the coupling of in-plane and out-of-plane motion inherent to bank-angle modulation.

A more robust formulation is to penalize not just the final inclination error, but also the inclination error at intermediate times $t_k < t_f$. While requiring more control effort, it produces a lower *average* inclination error along the entire trajectory, which improves robustness. To meet the more stringent intermediate inclination-error requirements, the bank-angle profiles produced under this formulation contain multiple switching events, and the resultant inclination profiles tend to look more like the gray line in Figure 6.3. Additional switching events are easily handled by CPAG because it solves for the control magnified *and* direction simultaneously, so the switching is

explicitly accounted for in both the prediction and correction phases.

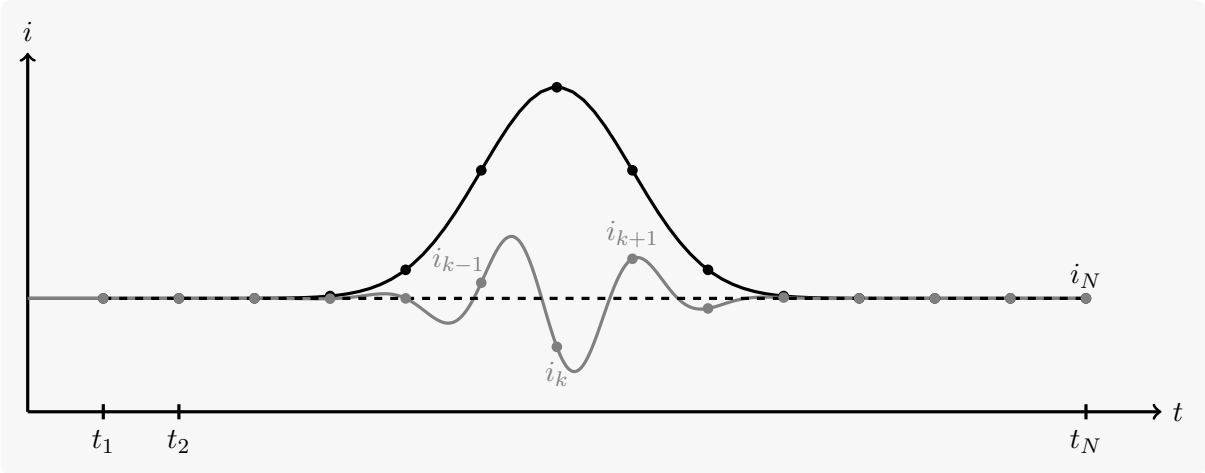

Figure 6.3: Comparison of inclination-targeting formulations.

The numerical implementation of the intermediate inclination error penalization is achieved by constraining the final $N - N_i - 1$ points in the subproblem

$$R_k \cdot \delta \boldsymbol{x}_k + \cos\left(i(\bar{\boldsymbol{x}}_k)\right) - \cos\left(i_{\text{target}}\right) \leq \nu_{\Delta i,k} \quad \forall k = N_i, \ldots, N-1 \quad . \tag{6.35}$$

Note that this constraint is only applied until the $N - 1$ node because the terminal constraint in Equation 6.26 is applied at the final node. From numerical experimentation, a value of $N_i = 0.8N$ is found to give good results, as constraining every node proved to be too restrictive. The cost function is then augmented with the slack variables $\nu_{\Delta,i}$ to encourage small inclination errors at each constrained node

$$J = J_0 + \lambda \left\| \boldsymbol{\nu}_{\Delta i} \right\|_2 \tag{6.36}$$

where $\lambda \in \mathbb{R}_+$ is a scalar weight and J_0 is the original cost. As the value of λ increases, the inclination errors at intermediate time steps are penalized more, and more bank-angle reversals are produced. Figure 6.4 shows an example comparing the inclination profiles produced with and without the intermediate inclination-error penalty. When minimizing the inclination error only at the final time step, large errors are observed along the trajectory for $t = 200$-350 seconds. In contrast, the use of Equation 6.35 keeps the inclination error within 0.5 degrees for the entire trajectory, which provides

CPAG with more flexibility to adjust the entire inclination profile to meet the terminal inclination constraint. The tighter inclination-errors require produce greater control costs, often times resulting in control saturation. Therefore, when using Eq. 6.35, the control magnitude is limited to $0.5u_{\max}$ for the first call to CPAG, which ensures sufficient control-authority margin is maintained. After the first call, this constraint is removed to allow CPAG to use more control when needed.

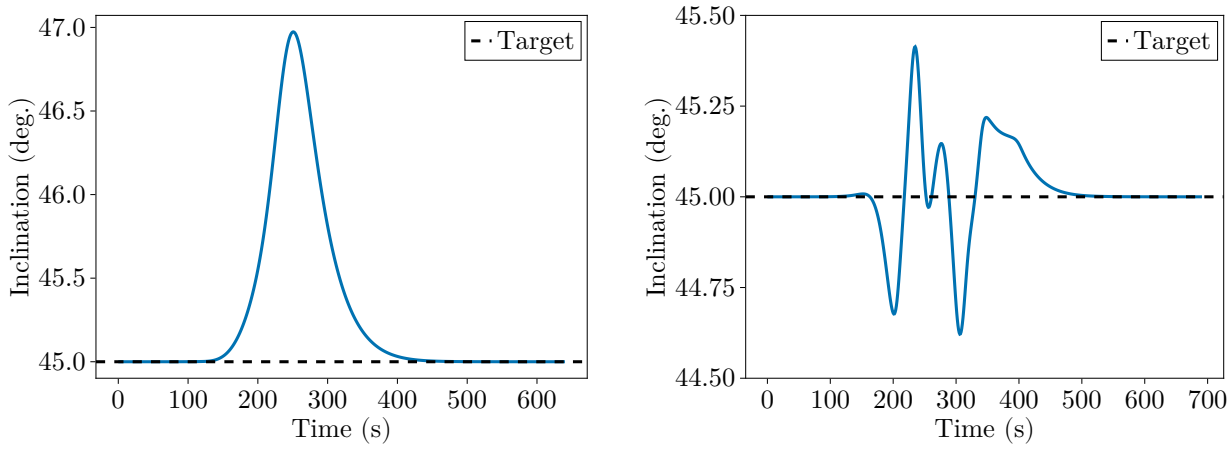

(a) Without intermediate penalty.

(b) With intermediate penalty.

Figure 6.4: Comparison of inclination profiles produced without (a) and with (b) penalties on intermediate inclination errors.

This work assumes the target inclination is the inclination at atmospheric entry. In this situation, a bank-angle profile that is full lift-up or full lift-down along the entire trajectory would produce no intermediate inclination error. For this reason, when the value of λ is large enough to produce multiple bank reversals, CPAG tends to fly the end of the trajectory full lift-up. In practice, this saturates the control and leaves no margin in the longitudinal channel[2] . In an aerocapture guidance algorithm, a more robust bank-angle profile would terminate somewhere near the midpoint of lift-up and lift-down [55]. A naive approach of encouraging this behavior would be to constrain the final $N - N_{\text{bank}}$ nodes to be equal to 90 degrees. This is suboptimal because it may be more efficient to have the final bank angle be -90 degrees. Instead, the optimizer should be able to decide

[2] This is a known pitfall of bank-angle modulation.

which direction the final $N - N_{\text{bank}}$ bank-angle values should be

$$\sigma_k \in \{-90, 90\} \quad k = N - N_{\text{bank}}, \ldots, N \quad . \tag{6.37}$$

Such a problem typically requires either a mixed-integer optimization algorithm [199], or a workaround such as homotopic paths combined with smooth approximations as done in [200]. To avoid both of these additional complexities, the constraint in Eq. 6.37 may be rewritten treating σ_k as a continuous variable subject to the following nonlinear equality constraint [184]

$$(\sigma_k - 90)(\sigma_k + 90) = 0 \quad \sigma_k \in \mathbb{R} \quad . \tag{6.38}$$

Unfortunately, this is not affine in σ_k and cannot be explicitly forced as an equality constraint in CPAG's convex subproblem. However, it is convex when relaxed into epigraphic form [201] and can be solved using the augmented convex-concave decomposition (ACCD) [184]. The ACCD decomposes Eq. 6.38 into the following four inequality constraints

$$-\varepsilon_1 \leq 0 \tag{6.39a}$$

$$-\varepsilon_2 \leq 0 \tag{6.39b}$$

$$\sigma_k^2 - (a+b)\,\sigma_k + ab - \varepsilon_1 \leq 0 \tag{6.39c}$$

$$(a + b - 2\bar{\sigma}_k)\,\sigma_k + \bar{\sigma}_k^2 - ab - \varepsilon_2 \leq 0 \tag{6.39d}$$

where $\bar{\sigma}_k$ is the bank angle along the reference trajectory $\bar{\boldsymbol{x}}_k$, $a = 90$ degrees, and $b = -90$ degrees. The slack variables ε_1 and ε_2 are appended to the cost function to encourage constraint satisfaction at convergence. Equation 6.39 is only included in the subproblem during the first call to CPAG, as constraining the final bank angle later in the trajectory results in less robust apoapsis targeting.

Figure 6.5 compares the bank-angle profiles produced by CPAG with and without Eq. 6.39 for values of $\lambda = 10^3$ and $N_{\text{bank}} = 25$. Without the ACCD, CPAG flies the final portion of its trajectory nearly full lift-up, which leaves no margin for apoapsis undershoot cases. As shown in Figure 6.5b, the ACCD is able to alleviate this concern in a robust manner, as CPAG is free to choose a final bank-angle of ± 90 degrees.

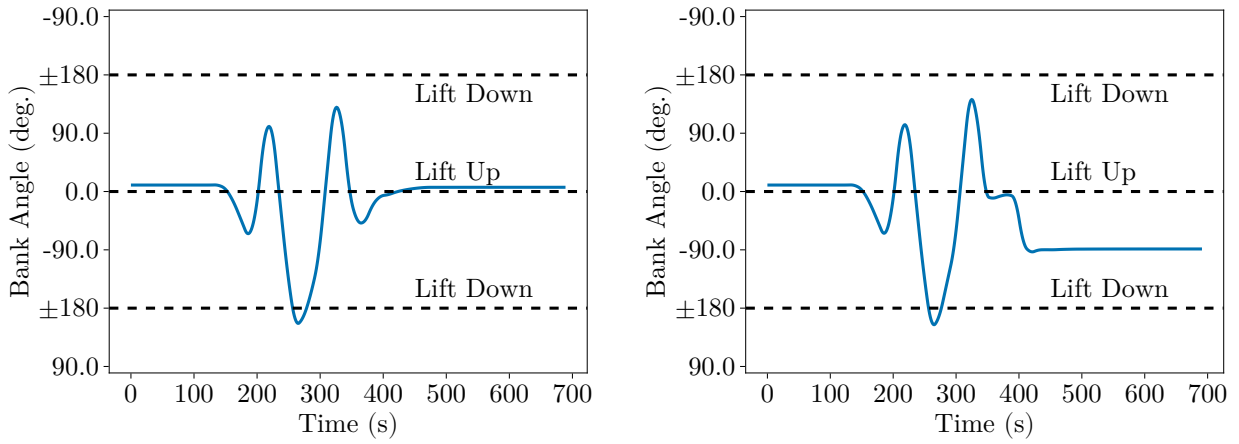

(a) Without ACCD.

(b) With ACCD.

Figure 6.5: Comparison of bank-angle profiles produced without (a) and with (b) ACCD decision-making on the final bank-angle value.

6.3.5 Scaling of the Optimization Variables

In trajectory optimization problems, the magnitude of the optimization variables can vary significantly. For example, in Eq. 5.1, the radius $r \sim \mathcal{O}(10^6)$, whereas, the bank angle rate $\dot{\sigma} \sim \mathcal{O}(10^0)$. This large discrepancy in magnitude can lead to significant buildup of numerical error from iteration to iteration of the optimization process, which is undesirable for an onboard guidance algorithm. One of two approaches may be taken in the scaling of the optimal control problem. The first scales the continuous-time equations of motion and constraints as done in [202, 203], while the second scales the discretized, parameter-optimization problem [204, 191]. This work adopts the latter technique and chooses to scale the discretized problem, as it allows for more precise scaling within the SCP framework. To this end, the following affine transformation is introduced

$$\delta \boldsymbol{x}_k = S_x \delta \hat{\boldsymbol{x}}_k + \boldsymbol{c}_x \tag{6.40a}$$

$$\delta u_k = S_u \delta \hat{u}_k + \boldsymbol{c}_u \tag{6.40b}$$

$$\delta p = S_p \delta \hat{p} + \boldsymbol{c}_p \tag{6.40c}$$

where S_x, S_u, S_p are diagonal, square matrices, $\boldsymbol{c}_x$, $\boldsymbol{c}_u$, $\boldsymbol{c}_p$ are vectors of commensurate size that center the variables, and the hat notation $\hat{(\cdot)}$ denotes the scaled variable. The specific form of the

i-th diagonal entry of S and i-th entry of $\boldsymbol{c}$ for some quantity a is

$$S_{a,ii} = \frac{a_{i,\text{upper}} - a_{i,\text{lower}}}{2} \tag{6.41a}$$

$$c_{a,i} = \frac{a_{i,\text{upper}} + a_{i,\text{lower}}}{2} \quad . \tag{6.41b}$$

Here $a_{i,\text{upper}}$ and $a_{i,\text{upper}}$ are upper and lower bounds on the i-th component of quantity a that will produce the scaled variables to lie in the range $\hat{a}_i \in [-1, 1]$. The form of Eq. 6.41 is a general approach for scaling all the optimization variables to be of similar order of magnitude. However, because the optimization variables in this work are corrections about the reference values, the lower and upper values of each optimization variable are taken to be equal and opposite in magnitude. This particular choice yields $\boldsymbol{c}_a = 0$ for the state, control, and parameter scaling. The scaling values used in this work are summarized in Table 6.1.

Table 6.1: Values used to scale optimization variables in CPAG.

Variable	δr	$\delta\theta$	$\delta\phi$	δv	$\delta\gamma$	$\delta\psi$	$\delta\sigma$	δQ	δu	δt_f
$S_{a,ii}$	1 km	0.2 deg.	0.2 deg.	50 m/s	0.2 deg.	0.2 deg.	15 deg.	0.5 kJ/cm^2	0.1 deg/s	5 s

6.3.6 Complete Convex Subproblem

Following the previous sections, Problem 1 is successfully converted into an approximate, first-order, discrete-time problem that can be solved using a variety of convex-optimization algorithm tools. The first-order approximation is only valid in a neighborhood around the reference solution, so the solution must be kept close enough to the reference trajectory. This is accomplished by using the following trust region constraint

$$\|\delta \boldsymbol{x}_k\|_2^2 + \|\delta u_k\|_2^2 \leq \eta_k \quad \forall\ k = 1, \ldots, N \tag{6.42a}$$

$$\|\delta p\|_2^2 \leq \eta_p \tag{6.42b}$$

where $\eta_k \in \mathbb{R}_+$, and $\eta_p \in \mathbb{R}_+$ are the trust region radii. To aid in convergence, the trust region radii are added as solution variables, and the cost function is augmented with

$$J_{tr} = \boldsymbol{w}_{tr}^T \boldsymbol{\eta} + w_{tr,p}\eta_p \tag{6.43}$$

where $\boldsymbol{w}_{tr} \in \mathbb{R}^N_{++}$ and $w_{tr,p} \in \mathbb{R}_{++}$ are trust region weights. The use of a PTR lets the trust region radii be determined by the optimizer instead of a separate trust-region adaptation law. Finally, the hard constraints on the apoapsis radius (Eq. 6.27) and final inclination (Eq. 6.26) may produce feasibility issues. Therefore, for a more numerically robust implementation in a predictor-corrector guidance algorithm, the constraints are expressed using an exact penalty term in the cost function. With that, the complete cost function may be stated as

$$\begin{aligned} J = \sum_{k=1}^{N} \alpha_1(\bar{u}_k + \delta u_k)^2 + \alpha_2 \left\| R \cdot \delta \boldsymbol{x}_N + \cos\left(i(\bar{\boldsymbol{x}}_N)\right) - \cos\left(i_{\text{target}}\right) \right\| + \\ \alpha_3 \left\| H \cdot \delta \boldsymbol{x}_N + \varepsilon(\bar{\boldsymbol{x}}_N) \right\| + \alpha_4 \left\| P \cdot \delta \boldsymbol{x}_N + \frac{1}{r_p(\bar{\boldsymbol{x}}_N)} \right\| + \alpha_5 \left\| \boldsymbol{\nu} \right\|_1 + J_{tr} + \lambda \left\| \boldsymbol{\nu}_{\Delta i} \right\|_2 \end{aligned} \tag{6.44}$$

where $\boldsymbol{\nu} = [\nu_a,\ \nu_Q]^T$, the $\alpha_i \in \mathbb{R}_+$ are constants used to weight each term in the cost appropriately, and $\lambda \in \mathbb{R}_+$ is used to penalize the intermediate inclination errors. The values of α_i should be chosen such that the terminal accuracy (i.e. apoapsis, inclination, etc.) terms are penalized the most, and the trust region radii do not get too large. More detail on the selection of these weights is provided in Section 6.4. Recall that the slack variables for the ACCD are appended to the cost function only during the first call to CPAG. Finally, the solution variables are scaled by substituting $\delta \boldsymbol{x}_k \leftarrow S_x \delta \hat{\boldsymbol{x}}_k$, $\delta u_k \leftarrow S_u \delta \hat{u}_k$, and $\delta t_f \leftarrow S_p \delta \hat{t}_f$ wherever they appear, so that the decision variables become the scaled state, control, and final time corrections. The full convex subproblem to solve at each guidance cycle is then

Problem 2. (CPAG convex subproblem): *Determine the values of* $\delta \hat{\boldsymbol{x}}_k$, $\delta \hat{u}_k$, $\delta \hat{t}_f$, *and* $\boldsymbol{\nu}$ *that minimize the cost* (6.44) *subject to equality constraints on the linearized dynamics* (6.25) *and initial condition* (6.34)*, and inequality constraints on the load factor* (6.31)*, heat load* (6.32)*, control magnitude* (6.33)*, intermediate inclination error* (6.35)*, and trust regions* (6.42).

6.3.7 Predictor-Corrector Formulation

The CPAG algorithm is summarized in Algorithm 1. At time t_n when the guidance algorithm is called, the current state $\boldsymbol{x}(t_n) = \boldsymbol{x}_0$ and control profile $u = \bar{u}$ are used to integrate the trajectory forward in time until the final time is reached. Next, the convex subproblem described in Problem 2 is solved to obtain a correction δu to update the nominal control profile such that $u_k = \bar{u}_k + \delta u_k \; \forall \; k$, as shown in Fig. 6.6. The loop in Algorithm 1 repeats until the change in the trajectory $\|\delta\hat{\boldsymbol{x}}\|_2$ is below some threshold or a maximum number of iterations is reached. The optimized control profile determined at $t_n > 0$ is flown until a time of $t_{n+1} = t_n + \Delta t$, where Δt is some specified time interval, and the process is repeated. When CPAG is called for the first time, the initial value of t_f is determined by numerically integrating the remainder of the trajectory until the atmospheric exit radius is reached. The initial guess for the control is assumed to be $u(t) = 0$. While a more specific choice for the initial control is allowable, assuming an initial value of zero is acceptable even if it produces a trajectory that remains hyperbolic due to the improved numerical robustness of the energy-based apoapsis radius expression in Eq. 6.4.

Algorithm 1: CPAG Algorithm

Input: Current state $\boldsymbol{x}_0$, control $u = \bar{u}$, and final time t_f
for n = 1: max iterations **do**
  Predict remaining trajectory $\bar{\boldsymbol{x}}$ using $\boldsymbol{x}_0$ and u
  Compute A_k, B_k^-, B_k^+, Σ_k, $\boldsymbol{z}_k$ about $\bar{\boldsymbol{x}}$ using Eq. 6.24
  Compute R, H, and P using Eqs. 6.26, 6.27, and 6.28
  Solve convex subproblem for $\delta\hat{\boldsymbol{x}}_k$, $\delta\hat{u}_k$, and $\delta\hat{t}_f$
  $u_k \mathrel{+}= S_u\delta\hat{u}_k$, Update control profile
  if $\|\delta\hat{\boldsymbol{x}}\|_2 \leq$ tolerance **then**
    Break
  end
end
Output: u

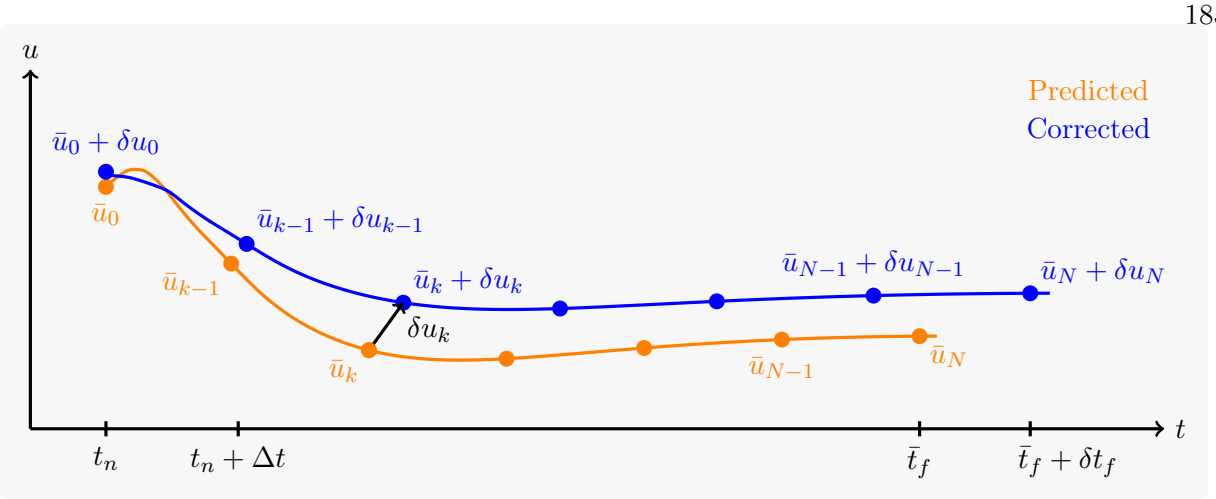

Figure 6.6: Schematic of convex predictor-corrector process performed at $t = t_n$.

6.4 Numerical Monte Carlo Results

In this section, CPAG is applied to aerocapture scenarios at Neptune, Mars, and Earth. CPAG is implemented in Julia [205] using the `Convex.jl` [188] toolbox and `Clarabel` [206] solver. All simulations are carried out on a MacBook Pro laptop with 32 GB of RAM and an M2 Pro CPU. The performance of CPAG is compared to the state-of-the-art algorithm FNPAG[3] by performing Monte Carlo analyses that disperse aerocapture-relevant parameters, including aerodynamic coefficients, initial conditions, and the atmospheric density profile. For each destination, the cost function scaling parameters α_i are set according to the values in Table 6.2. The cost-function terms with the highest priority, specifically the apoapsis target and the constraint satisfaction, are weighted more heavily. On the other hand, the trust region and control weights are relatively small to allow for large enough changes to the trajectory and control.

The standard version of FNPAG [55] uses Keplerian expressions for targeting the apoapsis radius and computing the in-plane ΔV (e.g. Eqs. 5.8 and 6.5), which assume the exo-atmospheric portion of the trajectory is not affected by J_2. For a more consistent comparison of CPAG and FNPAG, the Monte Carlo analyses in Sections 6.4.1–6.4.3 maintain this assumption such that J_2 in Equation 6.4 is set to zero. Section 6.5 relaxes this assumption, where the use of Eq. 6.4 to

[3] FNPAG is also implemented in Julia.

incorporate the effects of J_2 within the apoapsis targeting is explored in more detail.

Table 6.2: Values used to scale cost function in CPAG.

Term	Value
Control, α_1	1×10^{-2}
Inclination, α_2	5.0
Apoapsis, α_3	1,000
Periapsis, α_4	0.0 (Neptune and Mars), 1.0 (Earth)
Constraints, α_5	1,000
Trust region, w_{tr}, $w_{tr,p}$	1×10^{-3}
Interm. Inclination, λ	1,000

6.4.1 Neptune Monte-Carlo Results

The entry state is derived from an existing Neptune aerocapture systems analysis [13], and the initial conditions $\boldsymbol{x}_0 = [r_0, \phi_0, \theta_0, v_0, \gamma_0, \psi_0, \sigma_0]$, expressed in the inertial frame, are summarized in Table 6.3. The initial flight path angle is chosen to be in the middle of the nominal aerocapture corridor, and the initial bank-angle value is chosen to be consistent with the value of σ_0 used by FNPAG. Atmospheric properties are modeled using NeptuneGRAM [121], and all dispersed density profiles use a density perturbation scale of 1.0, which corresponds to the 3σ uncertainties computed within GRAM.

Table 6.3: Nominal inertial initial conditions for Neptune mission.

h_0 (km)	ϕ_0 (deg.)	θ_0 (deg.)	v_0 (km/s)	γ_0 (deg.)	ψ_0 (deg.)	σ_0 (deg.)
1000	0	0	29.0	-11.64	45	10

A $100,000 \times 3,000$ km orbit with a final inclination of 45 degrees is targeted. The entry vehicle is modeled as an MSL-style, 70-degree sphere cone with a base diameter of 5 meters and a nominal lift-to-drag ratio and ballistic coefficient of $L/D = 0.3$ and $\beta = 110$ kg/m^2, respectively. Once the sensed acceleration exceeds a limit of 0.1 g, the guidance system is initialized and is called at a frequency of 1 Hz. When CPAG is first called, Problem 2 is discretized using $N = 50$ nodes,

and each subsequent call to CPAG only uses nodes that are defined for $t_k \geq t_0$, where t_0 is the current time. Control and problem-specific constraints for the Neptune aerocapture scenario are summarized in Table 6.4. The bank-angle rate and acceleration are limited to ± 20 deg/s and ± 5 deg/s^2, respectively, which is consistent with previous work [174]. A load factor limit of $a_{\max} = 10$ g is assumed, which is within the structural capabilities of atmospheric entry capsules. The heat load constraint is set to $Q_{\max} = 98$ kJ/cm^2. As will be shown later, this specific value is chosen to be approximately equal to the 50-th percentile of the heat-load distribution produced by FNPAG. The Sutton-Graves correlation [158] for Neptune is used to model the stagnation-point convective heat flux. It is assumed that the guidance system has direct access to surface heating measurements, however, a more sophisticated estimation algorithm, such as that presented in [207], could be used to estimate the aerothermal environments if more limited measurements are available. Following Section 6.3.4, the intermediate inclination errors are penalized with a weight of $\lambda = 1000$, and the final $N_{\text{bank}} = 25$ nodes are constrained using the ACCD.

Table 6.4: Control and problem-specific constraints for Neptune aerocapture.

Parameter	Max. Value
Bank-angle rate	20 deg/s
Bank-angle accel.	5 deg/s^2
Load factor $a_{\max}$	10 g
Heat load $Q_{\max}$	98 kJ/cm^2

The performance of CPAG is first assessed in a nominal scenario (i.e. no dispersions). In this configuration, the guidance simulation and the truth simulation use the same density, aerodynamic, and initial condition. Evaluating a guidance algorithm in a nominal scenario allows for easy diagnosis of the algorithm's fundamental settings, as it should result in excellent aerocapture performance. This work quantifies the performance of the guidance system using the final apoapsis targeting error and total ΔV required to correct the post-aerocapture orbit. For a trajectory that exits the atmosphere on an elliptical orbit, a three-burn concept of operations is assumed. The first burn is applied at apoapsis and is used to raise the trajectory periapsis out of the atmosphere to the desired

altitude, and the second burn is applied at periapsis to correct for any apoapsis targeting errors. The final burn is applied at the corrected apoapsis and is used to correct orbit's inclination

$$\Delta V_3 = 2v_a^* \sin(\Delta i/2) \tag{6.45}$$

where v_a^* is the apoapsis velocity of the target orbit, and Δi is the error in the orbit's inclination. Using Equations 6.5 and 6.45, the total ΔV is $\Delta V_{\text{tot}} = \Delta V_1 + \Delta V_2 + \Delta V_3$.

Figure 6.7 shows the resulting trajectory profile and inclination in the nominal scenario, and Fig. 6.8 shows the corresponding bank-angle rate and bank-angle profiles. For this case, the in-plane orbital corrections require a $\Delta V_{1,2} = 190.0$ m/s. The in-plane cost is dominated by the periapsis raise maneuver, as CPAG produces a nominal apoapsis error of $\Delta r_a = 100$ m, and the target orbit is highly elliptical (eccentricity of 0.94). The out-of-plane corrections require a $\Delta V_3 = 0.0$ m/s, as the final inclination error is $\Delta i = 0$ degrees.

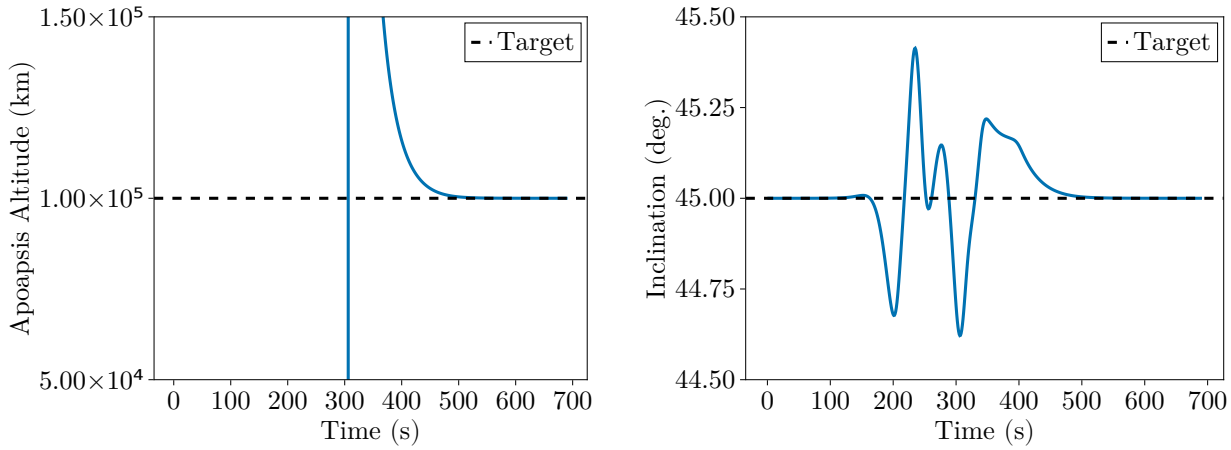

(a) Apoapsis altitude.

(b) Inclination.

Figure 6.7: Time history of apoapsis altitude (a) and inclination (b) produced using CPAG in a nominal scenario.

As shown in Fig. 6.8, once CPAG is initialized at $t \approx 120$ seconds, it begins modulating the bank angle rate in order to hit the target apoapsis radius and inclination. To maintain a sufficiently small inclination error along the entire trajectory, several bank-angle reversals are observed. Recall that CPAG solves for an entire bank-angle rate profile (i.e. magnitude and direction) at each

time step. This allows for the bank-angle reversals illustrated in Fig. 6.8b to be accounted for explicitly within *both* the prediction *and* the correction phase, as the algorithm is trying to meet both apoapsis and inclination targets simultaneously. This characteristic is significantly different from existing aerocapture guidance algorithms that use discrete lateral logic algorithms to determine bank-angle reversals. However, the bank-angle profile in Fig. 6.8b is not a "bang-bang" profile, as the corresponding bank-angle rates are within the limit of 20 deg/s. In [55], it is shown that the ΔV optimal bank-angle profile is bang-bang, suggesting that CPAG may be losing some optimality in its performance. However, as shown in the following results, the ΔV produced by CPAG is close to the ΔV produced by FNPAG, which does fly bang-bang bank-angle profiles. Therefore, the optimality lost by not flying a pure bang-bang profile is sufficiently small.

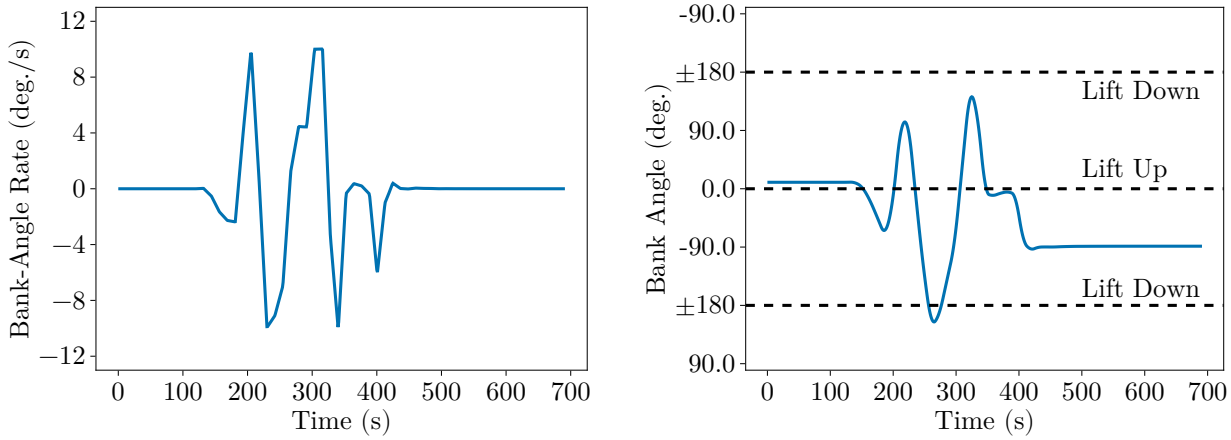

(a) Bank-angle rate.

(b) Bank angle.

Figure 6.8: Bank-angle rate profile determined by CPAG (a) and resulting bank-angle profile (b) in a nominal scenario.

As described in Section 6.3, CPAG uses a penalized trust region (PTR) technique to keep the solution near the reference and encourage convergence to a feasible solution. As the algorithm converges, the trust region radii shrink. This characteristic is demonstrated in Fig. 6.9 which shows the evolution of the trust region radii at each discrete node when CPAG is called for the first time in the nominal scenario. At the first point ($k = 1$), the trust region radius is driven to zero due to

the initial-condition constraint. For later values of k, the trust region radii adapt as needed in order to shape the trajectory and control profiles, and once the algorithm is in the vicinity of the optimal solution, they begin to shrink quickly. This behavior indicates that the algorithm is functioning as intended and showcases the flexibility of the PTR formulation.

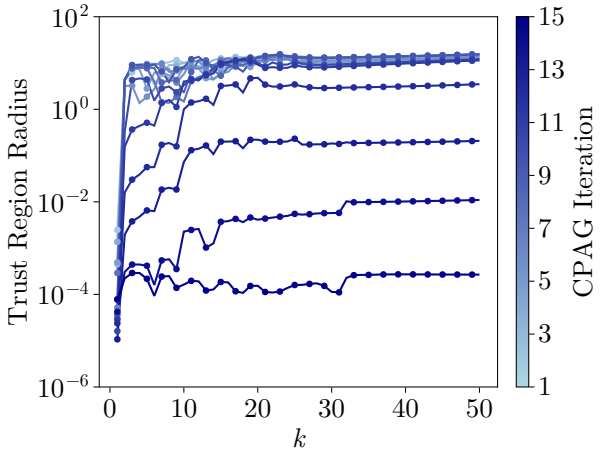

Figure 6.9: Convergence of trust region radii during first call to CPAG.

To assess the performance of CPAG in off-nominal scenarios, a 1000-case Monte Carlo analysis is performed in which the vehicle aerodynamic coefficients, vehicle mass, entry velocity, entry flight-path angle, and atmospheric density profile are dispersed according to the values in Table 6.5. In each of the Monte Carlo runs, the predictor step of CPAG utilizes the nominal aerodynamic model and nominal density profile from NeptuneGRAM, while the truth simulation uses the dispersed aerodynamics and density profile. The GRAM profiles are implemented as lookup tables in both the guidance and truth simulation. Two first-order fading-memory filters, as described in Ref. 55, are used to scale the nominal lift and drag accelerations using sensed acceleration data accessible from the vehicle's inertial measurement unit (IMU) to help account for the differences in the aerodynamic/density models used onboard and those used by the truth simulation. The navigation state is assumed to be perfect for all simulations. Uncertainties in the initial velocity and flight-path angle are chosen to be consistent with recent interplanetary navigation solutions for Ice Giant aerocapture missions [116] and recent aerocapture guidance analyses [208, 128, 209].

The aerodynamic dispersions are chosen to produce an uncertainty in the lift-to-drag ratio of approximately ± 0.05, and an uncertainty in the ballistic coefficient of approximately ± 10.0 kg/m^2. Atmospheric density dispersions are applied using NeptuneGRAM with a density perturbation scale of 1.0, which returns the 3σ GRAM density distribution. Samples for all parameters, except for density, are generated from normal distributions using the nominal and 3σ values in Table 6.5.

Table 6.5: Dispersions applied in Monte Carlo analysis.

Quantity	Nominal Value	3σ Uncertainty	Distribution
Lift coefficient, C_L	0.495	0.065	Normal
Drag coefficient, C_D	1.65	0.05	Normal
Mass, m	2200 kg	150 kg	Normal
Entry velocity, v_0	29.0 km/s	100 m/s	Normal
Entry FPA, γ_0	-11.64 deg	0.2 deg	Normal
Density	Mean GRAM	Dispersed GRAM	From GRAM

For consistent comparison, a 1000-case Monte Carlo analysis with identical dispersions is run using FNPAG. Similar to the settings used for CPAG, FNPAG is initialized once the sensed acceleration reaches 0.1 g, after which it is called at a frequency of 1 Hz. The predictor step of FNPAG also utilizes the nominal aerodynamic model and density profile from NeptuneGRAM, while the truth simulation uses the dispersed aerodynamics and density profile. Two first-order fading-memory filters, identical to those used by CPAG, are used to scale the nominal lift and drag accelerations. Finite-rate bank maneuvers are modeled with a maximum bank-angle rate of 20 deg/s, and a maximum bank-angle acceleration of 5 deg/s^2. FNPAG has two tuning parameters σ_0 and σ_f that dictate the bang-bang bank-angle profile flown during its first phase. This work assumes values of $\sigma_0 = 10$ degrees and $\sigma_f = 90$ degrees for all trajectories in the Monte Carlo analysis. These values are chosen to provide sufficient control margin and crossrange control authority [55]. A choice of $\sigma_f = 90$ degrees is made because it is found to produce a good balance of robustness and performance, as larger values tend to result in fewer trajectories capturing. The bank-angle sign is determined using predictive logic that manages inclination error [176] in which a maximum of five bank reversals is assumed. Note that this specific lateral logic was developed separately from the

original FNPAG algorithm, which used a simpler inclination-based dead band to determine bank reversals. Similar to CPAG, FNPAG assumes that the navigation state is perfect. However, unlike CPAG, FNPAG has no way of modifying its commands to account for the heat-load constraint.

For all 1000 cases, CPAG produces a capture rate of 100%, while FNPAG had two cases that remained hyperbolic producing a capture rate of 99.8%. This is potentially because FNPAG uses the standard apoapsis radius expression in Eq. 5.8 which is not well-behaved for hyperbolic orbits. The two algorithms produce similar apoapsis targeting accuracy, and Fig. 6.10 shows the apoapsis accuracy and required in-plane ΔV. Each algorithm produces similar trends in the resultant periapsis and apoapsis radii, with more apoapsis undershoot occurring for lower periapsis radii and more apoapsis overshoot at higher periapsis radii. Generally, a lower periapsis radius indicates that too much energy has been depleted during the atmospheric flight, so the resulting apoapsis radius is below the target value, whereas, a higher periapsis radius indicates the opposite result. The

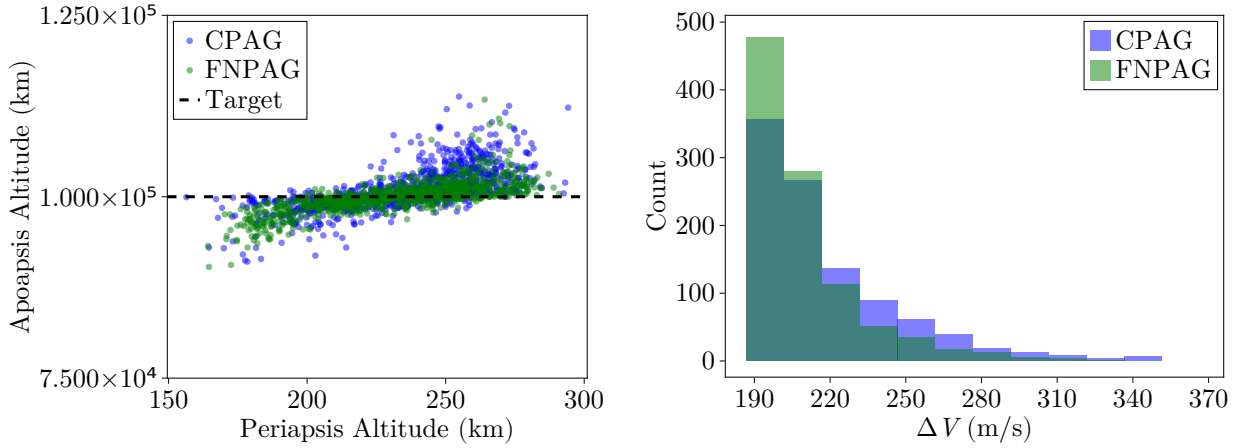

(a) Apoapsis accuracy.

(b) In-plane ΔV.

Figure 6.10: Comparison of apoapsis targeting accuracy (a) and required in-plane ΔV (b) for Neptune aerocapture.

average error in apoapsis radius is 1.69% and 1.54% using CPAG and FNPAG, respectively. The slightly larger apoapsis error produced by CPAG is due to the relative weighting of the apoapsis and inclination terms in the convex subproblem. This relative weighting is a trade that may be

tailored to a specific aerocapture scenario to place a greater emphasis on either the in-plane or out-of-plane accuracy (or a balance between the two). Figure 6.10b shows the in-plane ΔV required to correct the apoapsis and periapsis radii for both algorithms. The distribution of in-plane ΔV for each algorithm is relatively similar, with CPAG having slightly more cases in the higher ΔV bins. The control law for FNPAG is derived from optimal control theory, so the similarity in the two distributions demonstrates that the sequential convex optimization routine employed by CPAG is not sacrificing significant performance by solving a sequence of approximate convex subproblems.

The in-plane ΔV may be separated into two components: the periapsis raise, and the apoapsis correction. Figure 6.11 displays histograms for each of these components for both algorithms. The magnitude of the periapsis raise burn required for both algorithms is similar, with mean ΔV requirements of 188 m/s and 190 m/s for CPAG and FNPAG, respectively. As shown in Fig. 6.10a the similarity in the periapsis raise ΔV magnitude is attributed to the similar range of periapsis radii produced by CPAG and FNPAG. The magnitude of the apoapsis correction burn required when using CPAG is slightly higher than that required when using FNPAG, with mean apoapsis correction ΔV requirements of 30.3 m/s and 29.9 m/s for CPAG and FNPAG, respectively.

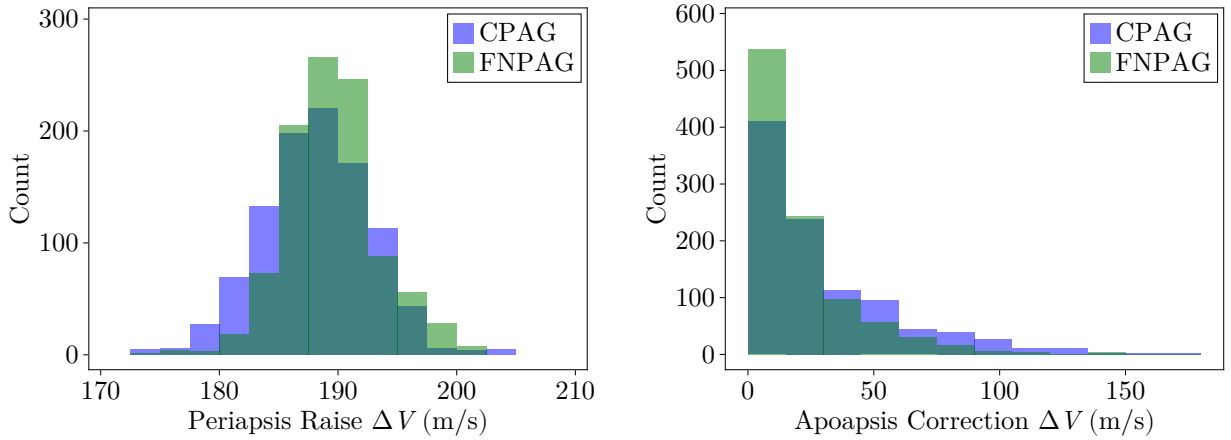

(a) Periapsis raise ΔV.

(b) Apoapsis correction ΔV.

Figure 6.11: Comparison of periapsis raise ΔV (a) and apoapsis correction ΔV (b) for Neptune aerocapture.

In addition to the in-plane ΔV requirements, consideration must be also be given to the ΔV necessary to correct the final orbit inclination. The out-of-plane ΔV_3 requirements for each algorithm are shown in Fig. 6.12a, and the corresponding total ΔV requirements are shown in Fig. 6.12b. A noticeable difference between the two algorithms is observed for the out-of-plane performance, with 50-th percentile requirements of 0.44 m/s and 4.5 m/s for CPAG and FNPAG, respectively. This difference is attributed to the difference in the architecture of the two algorithms. Specifically, FNPAG separates the longitudinal (i.e. bank-angle magnitude) and lateral (i.e. bank-angle direction) guidance channels, whereas CPAG solves for both the lateral and longitudinal controls simultaneously. By separating the two channels, the bank reversal is not accounted for within the prediction phase and the desire to target a specific inclination is not considered in the correction phase. As a result, the FNPAG correction phase only seeks to minimize in-plane (i.e. apoapsis) errors, while CPAG is able to consider both in-plane and out-of-plane targets. During the discrete bank-angle reversals performed by FNPAG, the vehicle must fly an open-loop control profile until the reversal is complete, during which it is susceptible to the accumulation of additional error.

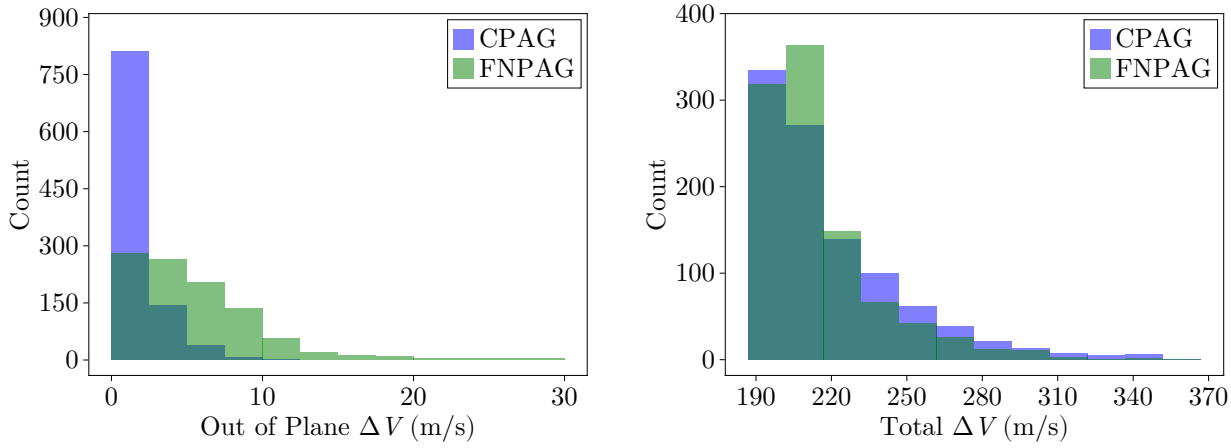

(a) ΔV_3

(b) Total ΔV.

Figure 6.12: Required ΔV for inclination change (a) and total orbit correction (b) for Neptune aerocapture.

Figure 6.13 compares the integrated stagnation-point heat load and load factor for both algorithms. Recall that CPAG is able to directly enforce the satisfaction of both the heat-load and load-factor constraints in its correction phase, while FNPAG has no way of doing so. Therefore, in order to stress CPAG more, the magnitude of the heat-load constraint used in CPAG is approximately chosen to be the 50-th percentile of the heat load distribution produced by FNPAG. Figure 6.13a demonstrates that CPAG is able to satisfy the heat-load constraint of 98 kJ/cm^2 in 100% of the cases, while FNPAG only produces a constraint satisfaction rate of 47.4%, by design. The addition of the heat-load constraint also explains why the in-plane ΔV requirements for CPAG are slightly higher than for FNPAG, indicating that CPAG is sacrificing some orbit targeting accuracy in order to satisfy the heat-load constraint. Finally, as shown in Fig. 6.13b, both algorithms satisfy the load-factor constraint of 10 g. The initial phase of an aerocapture trajectory is typically flown with a near lift-up bank angle, so the peak aerodynamic loading is typically not that high. Overall, CPAG produces less severe aerothermodynamic (heating and loading) environments compared to FNPAG due to the explicit consideration of these quantities in the correction process.

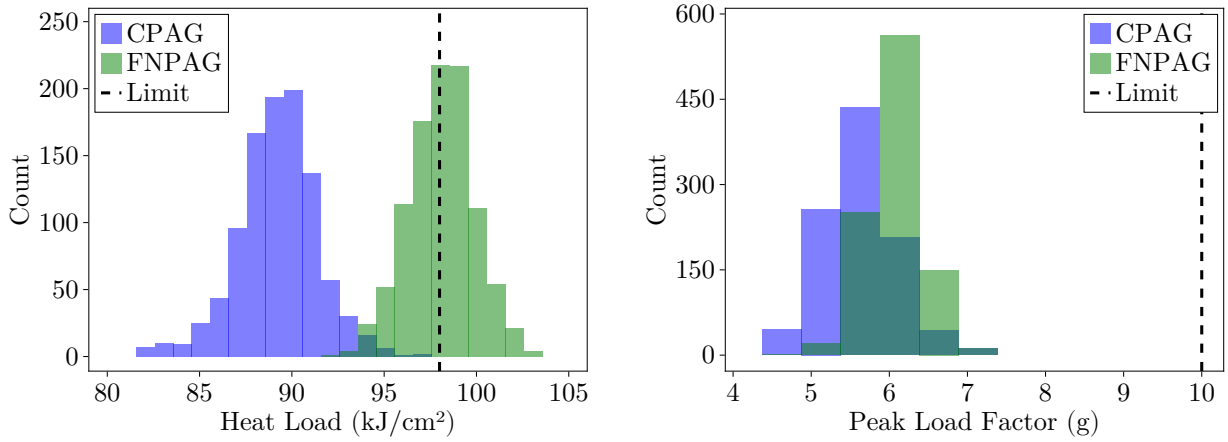

(a) Heat Load.

(b) Maximum load factor.

Figure 6.13: Integrated heat load (a) and maximum load factor (b) for both CPAG and FNPAG for Neptune aerocapture.

Table 6.6 presents an overall comparison for the statistics of all quantities of interest obtained

using both CPAG and FNPAG. With respect to the ΔV requirements, both algorithms produce similar results. Note that the 2 cases that did not capture for FNPAG produce an inflated standard deviation in the in-plane and total ΔV requirements. CPAG satisfies the system-level constraints in 100% of cases, whereas, FNPAG fails to do so. Table 6.6 also compares the computational time of each algorithm, where the reported times correspond to the amount of time spent on each predictor-corrector iteration. Figure 6.14 illustrates a representative time history of the computational time required by CPAG. Generally, at later iterations, less time is required as the number of points N in the discretized problem has decreased which leads to faster solve times. In general, FNPAG requires slightly less time than CPAG, although the solve times for both are sufficiently small.

Table 6.6: Comparison of Neptune aerocapture statistics for CPAG and FNPAG.

CPAG (Capture Rate = 100.0%)	Mean	Std. Dev.	1st Perc.	50th Perc.	99th Perc.
In-Plane ΔV (m/s)	218.4	28.94	188.9	209.1	326.0
Out-of-Plane ΔV (m/s)	1.235	1.686	7.900×10^{-4}	4.385×10^{-1}	7.010
Total ΔV (m/s)	219.6	29.14	189.2	210.8	328.5
Apoapsis Error (%)	1.688	1.744	1.136×10^{-2}	1.119	7.557
Max. Load Factor (g)	5.626	0.4686	4.615	5.605	6.983
Heat Load (kJ/cm^2)	89.24	2.160	83.06	89.32	94.27
Computational Time (ms)	18.8	4.75	12.9	17.5	34.0
FNPAG (Capture Rate = 99.8%)	Mean	Std. Dev.	1st Perc.	50th Perc.	99th Perc.
In-Plane ΔV (m/s)	219.6	216.8	187.9	202.6	306.7
Out-of-Plane ΔV (m/s)	5.413	4.499	4.777×10^{-2}	4.487	23.30
Total ΔV (m/s)	225.1	217.0	190.5	208.3	313.4
Apoapsis Error (%)	1.543	9.503	1.612×10^{-2}	7.443×10^{-1}	6.293
Max. Load Factor (g)	6.081	3.463×10^{-1}	5.211	6.096	6.918
Heat Load (kJ/cm^2)	98.03	2.236	93.94	98.10	102.2
Computational Time (ms)	5.10	1.76	4.22	5.10	9.46

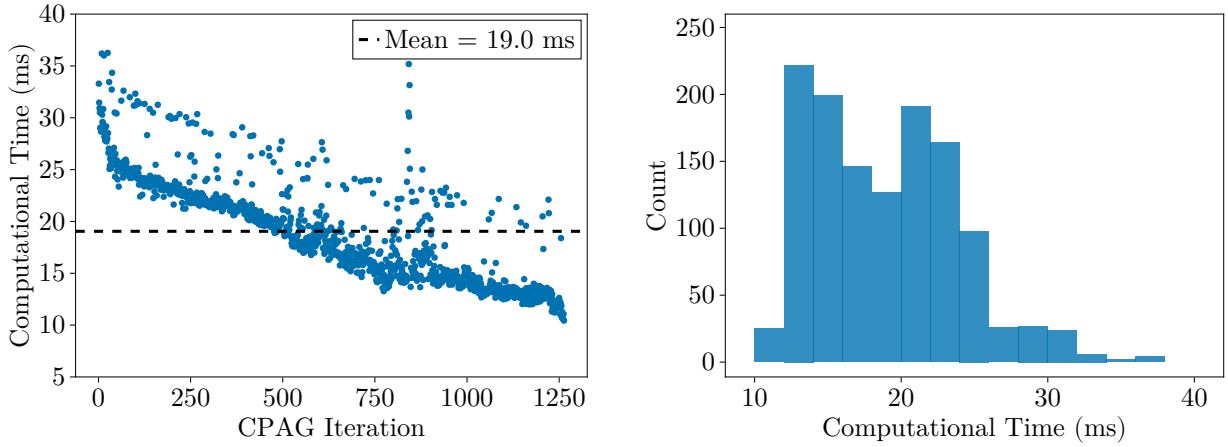

(a) Time history.

(b) Histogram.

Figure 6.14: Time history (a) and histogram (b) of computational time required during each CPAG iteration.

6.4.2 Mars Monte-Carlo Results

The entry state for Mars is derived from Refs. [210] and [208], which investigated the use of aerocapture for human-scale missions to Mars, and the initial conditions are summarized in Table 6.7. All atmospheric properties are modeled using MarsGRAM [211], where the dispersed density profiles produced use a density perturbation scale of 1.0. The target orbit for the Mars aerocapture demonstration is a highly elliptical, 1-sol orbit, with an apoapsis altitude of 33,793 km, a periapsis altitude of 250 km, and an inclination of 90 degrees (i.e. polar). In contrast to the low L/D blunt capsule studied for the Neptune aerocapture scenario, the Mars aerocapture scenario considers a mid L/D geometry with nominal mass and aerodynamic properties summarized in Table 6.8. A higher L/D vehicle has more control authority than a blunter vehicle, and comparing CPAG and FNPAG across a variety of vehicle geometries provides a more comprehensive comparison of the two algorithms.

CPAG and FNPAG are both initialized once the sensed acceleration exceeds a limit of 0.1 g, after which each guidance system is called at a frequency of 1 Hz. When CPAG is first called, Problem 2 is discretized using $N = 50$ nodes, and each subsequent call to CPAG only uses nodes that

Table 6.7: Nominal inertial initial conditions for Mars mission [210, 208].

h_0 (km)	ϕ_0 (deg.)	θ_0 (deg.)	v_0 (km/s)	γ_0 (deg.)	ψ_0 (deg.)	σ_0 (deg.)
150	0	0	6.2	-10.1	90	10

are defined for $t_k \geq t_0$, where t_0 is the current time. Control and problem-specific constraints for the Mars aerocapture scenario are summarized in Table 6.9. The load factor and heat load constraints in CPAG are set to $a_{\max} = 3$ g (to be consistent with human-rated mission requirements) and 6.35 kJ/cm^2, respectively, and both CPAG and FNPAG limit the bank-angle rate to a magnitude of 20 deg/s and the bank-angle acceleration to 5 deg/s^2. For both algorithms, the prediction step utilizes the nominal density profile from MarsGRAM, while the truth simulation uses a dispersed density profile, and the predicted lift and drag accelerations are scaled using first-order fading-memory filters. The stagnation-point heat flux is modeled using the Sutton-Graves correlation for the Martian atmosphere. A 1000-case Monte Carlo analysis is run for both CPAG and FNPAG using the 3σ aerodynamic and initial condition uncertainties listed in Table 6.5, and the atmospheric density is dispersed using MarsGRAM. A comparison of the performance of each algorithm is presented in Figures 6.15–6.18.

Overall, CPAG and FNPAG produce similar distributions of in-plane ΔV requirements. Due to the highly elliptical 1-sol target orbit, the periapsis raise burn is much smaller than the apoapsis correction burn. As a result, the decreased apoapsis targeting accuracy observed for CPAG produces slightly higher overall in-plane ΔV requirements. Because the inclination burn is applied at apoapsis, the ΔV required to correct the final inclination is significantly smaller than those shown for Neptune. As a result, both algorithms have small out-of-plane ΔV requirements.

Table 6.8: Nominal mass and aerodynamic properties for mid L/D vehicle [210].

Mass (kg)	L/D	β (kg/m^2)
62,000	0.55	380

Table 6.9: Control and problem-specific constraints for Mars aerocapture.

Parameter	Max. Value
Bank-angle rate	20 deg/s
Bank-angle accel.	5 deg/s^2
Load factor $a_{\max}$	3 g
Heat load $Q_{\max}$	6.35 kJ/cm^2

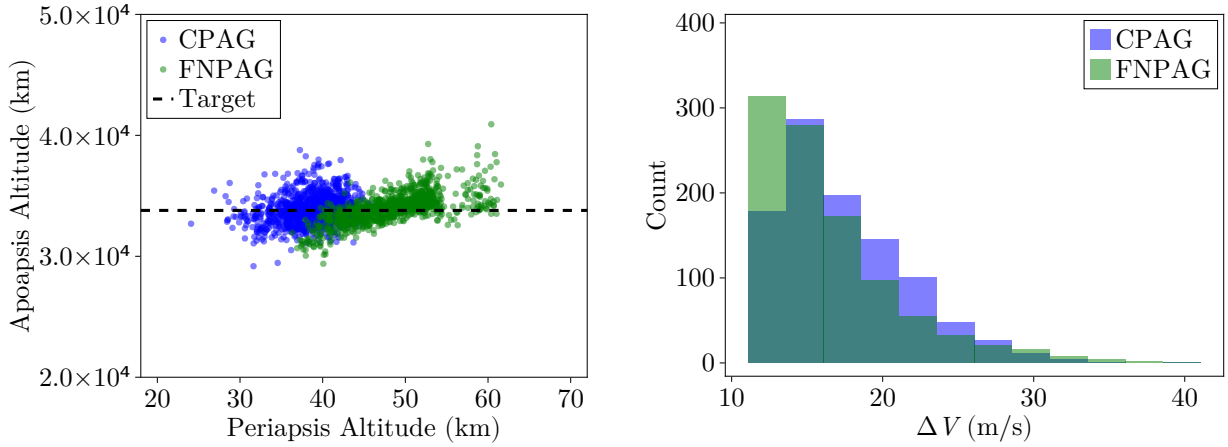

(a) Apoapsis accuracy

(b) In-plane ΔV.

Figure 6.15: Comparison of apoapsis targeting accuracy (a) and in-plane ΔV (b) for Mars aerocapture.

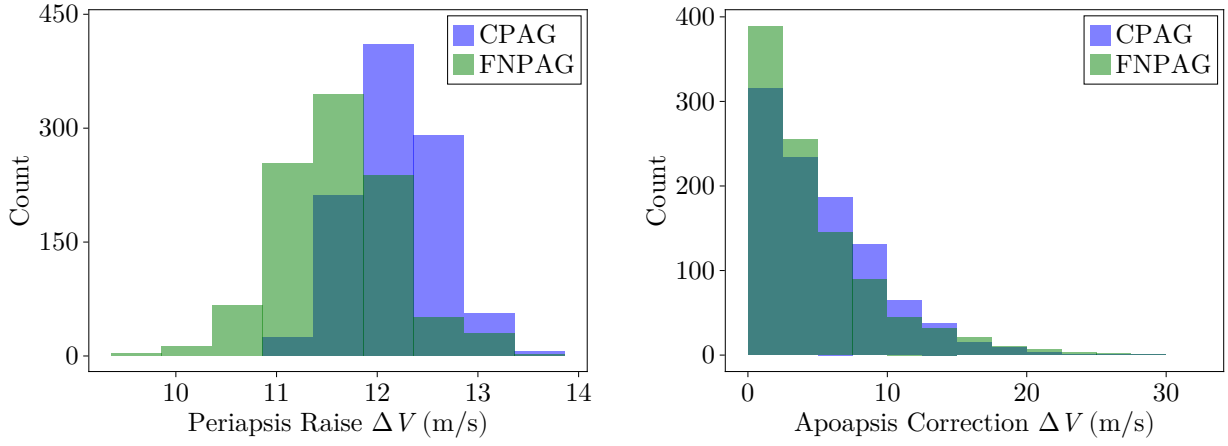

(a) Periapsis raise ΔV.

(b) Apoapsis correction ΔV.

Figure 6.16: Comparison of periapsis raise ΔV (a) and apoapsis correction ΔV (b) for Mars aerocapture.

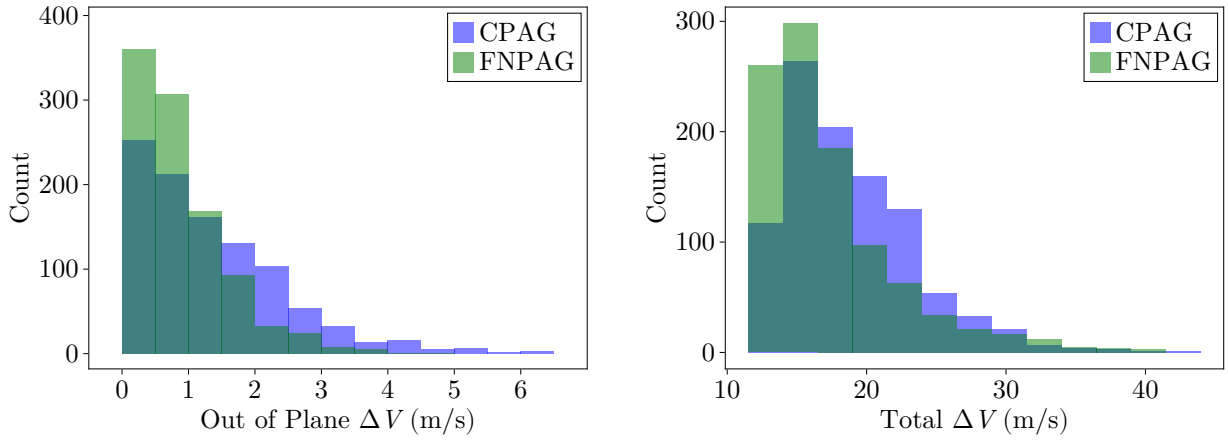

(a) ΔV_3 (b) Total ΔV.

Figure 6.17: Required ΔV for inclination change (a) and total orbit correction (b) for Mars aerocapture.

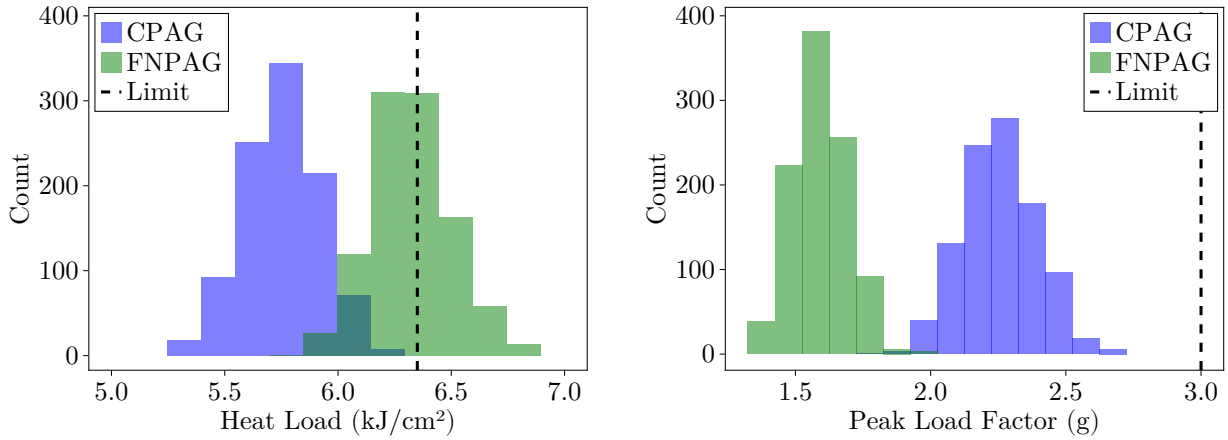

(a) Heat load. (b) Load factor.

Figure 6.18: Comparison of stagnation-point heat lead (a) and load factor (b) for Mars aerocapture.

Both FNPAG and CPAG produce load-factor distributions that are below the limit of 3 g. On the other hand, the heat load constraint of 6.35 kJ/cm^2 is satisfied in 100% of the cases using CPAG but only 59.4% of the cases using FNPAG. The heat-load constraint causes CPAG to produce lower periapsis radii, which explains the difference in the periapsis altitude distributions in Fig. 6.15a. It is interesting that lower periapsis radii are observed for the Mars aerocapture case, while higher periapsis radii are observed for the Neptune aerocapture case. This is likely explained by the different vehicle geometries assumed for each scenario. Due to the lower entry velocity and highly elliptical target orbit, the aerothermodynamic environments are less extreme relative to the Neptune results. It may be possible to increase the apoapsis targeting accuracy of CPAG if the heat-load constraint is relaxed. An overall comparison of the statistics for the aerocapture quantities of interest obtained using both CPAG and FNPAG is presented in Table 6.10.

Table 6.10: Comparison of Mars aerocapture statistics for CPAG and FNPAG.

CPAG (Capture Rate = 100.0%)	Mean	Std. Dev.	1st Perc.	50th Perc.	99th Perc.
In-Plane ΔV (m/s)	17.50	4.280	12.29	16.51	30.27
Out-of-Plane ΔV (m/s)	1.383	1.267	4.380×10^{-3}	1.118	5.769
Total ΔV (m/s)	18.88	4.739	12.48	17.91	33.36
Apoapsis Error (%)	2.585	2.120	4.623×10^{-2}	2.101	10.03
Max. Load Factor (g)	2.255	1.406×10^{-1}	1.947	2.253	2.585
Heat Load (kJ/cm^2)	5.755	1.705×10^{-1}	5.347	5.765	6.124
Computational Time (ms)	17.8	4.13	11.2	17.9	32.5
FNPAG (Capture Rate = 100.0%)	Mean	Std. Dev.	1st Perc.	50th Perc.	99th Perc.
In-Plane ΔV (m/s)	16.51	4.667	11.65	15.08	32.80
Out-of-Plane ΔV (m/s)	8.669×10^{-1}	7.205×10^{-1}	1.148×10^{-2}	6.8967×10^{-1}	3.229
Total ΔV (m/s)	17.37	4.948	11.932	15.96	34.88
Apoapsis Error (%)	2.390	2.306	3.622×10^{-1}	1.713	10.28
Max. Load Factor (g)	1.592	1.005×10^{-1}	1.3752	1.586	1.822
Heat Load (kJ/cm^2)	6.327	1.756×10^{-1}	5.916	6.312	6.764
Computational Time (ms)	5.08	1.23	4.18	5.64	10.26

6.4.3 Earth Monte-Carlo Results

The entry state for Earth is taken from the original FNPAG paper [55], and the initial conditions are summarized in Table 6.11. All atmospheric properties are modeled using EarthGRAM

[212], where the dispersed density profiles produced use a density perturbation scale of 1.0. The target orbit for the Earth aerocapture scenario (representative of a lunar return mission) is a 200 km circular orbit with an inclination of zero degrees, and the entry vehicle is chosen to be consistent with the Orion MPCV. The nominal mass and aerodynamic properties, taken from Ref. [55], are summarized in Table 6.12.

Table 6.11: Nominal inertial initial conditions for Earth mission [55].

h_0 (km)	ϕ_0 (deg.)	θ_0 (deg.)	v_0 (km/s)	γ_0 (deg.)	ψ_0 (deg.)	σ_0 (deg.)
122	0	0	11.02	-5.74	0	10

Table 6.12: Nominal mass and aerodynamic properties for Orion MPCV [55].

Mass (kg)	L/D	β (kg/m^2)
8983.4	0.27	323

CPAG and FNPAG are both initialized once the sensed acceleration exceeds a limit of 0.1 g, after which each guidance system is called at a frequency of 1 Hz. When CPAG is first called, Problem 2 is discretized using $N = 50$ nodes, and each subsequent call to CPAG only uses nodes that are defined for $t_k \geq t_0$, where t_0 is the current time. Control and problem-specific constraints for the Earth aerocapture scenario are summarized in Table 6.13. The load factor and heat load constraints in CPAG are set to $a_{\max} = 6$ g and 13.5 kJ/cm^2, respectively, and both CPAG and FNPAG limit the bank-angle rate to a magnitude of 20 deg/s and the bank-angle acceleration to 5 deg/s^2. For both algorithms, the prediction step utilizes the nominal density profile from EarthGRAM, while the truth simulation uses a dispersed density profile, and the predicted lift and drag accelerations are scaled using first-order fading-memory filters. The stagnation-point heat flux is modeled using the Sutton-Graves correlation for the Earth atmosphere. A 1000-case Monte Carlo analysis is run for both CPAG and FNPAG using the 3σ aerodynamic and initial condition uncertainties listed in Table 6.5, and the atmospheric density is dispersed using EarthGRAM. A comparison of the

performance of each algorithm is presented in Figures 6.19–6.22.

Table 6.13: Control and problem-specific constraints for Earth aerocapture.

Parameter	Max. Value
Bank-angle rate	20 deg/s
Bank-angle accel.	5 deg/s^2
Load factor $a_{\max}$	6 g
Heat load $Q_{\max}$	13.5 kJ/cm^2

While CPAG produces slightly worse overall apoapsis targeting, it actually results in lower overall in-plane ΔV requirements. From Fig. 6.19a, CPAG tends to have higher periapsis radii compared to FNPAG. Because the target orbit is circular, the periapsis raise burn dominates the overall in-plane ΔV requirements. Compared to FNPAG, the slightly higher periapsis radii for CPAG result in lower overall periapsis raise ΔV values, while the apoapsis error requires a relatively larger apoapsis correction burn. However, the magnitude of the apoapsis correction burn is significantly less than the periapsis raise burn, explaining why the overall in-plane ΔV requirements for CPAG are lower than that for FNPAG. Similar to the previous two cases, the ΔV required to correct the final inclination is low for CPAG. Unlike the previous two cases, there are a few outlier cases for CPAG where the inclination correction burn is high. The out-of-plane ΔV for FNPAG is larger due to the low-altitude circular target orbit, as the lower-altitude orbit has a higher velocity than the target orbits for Neptune and Mars, so any inclination targeting errors will be magnified. As a result, the total ΔV requirements for CPAG are lower than those for FNPAG. The FNPAG results could likely be improved by considering a larger value of σ_d, as the Earth scenario has little risk of any trajectories remaining hyperbolic.

From Fig. 6.22a, the heat-load distributions for CPAG and FNPAG are much more similar than those for Neptune and Mars. The heat load constraint of 13.5 kJ/cm^2 is satisfied for 100% of the cases using CPAG, but only in 89.0% of the cases when using FNPAG. A larger similarity in the heat-load profiles is observed likely because the lower target apoapsis altitude constrains the trajectory space more than if the target orbit were highly elliptical. The load factor constraint of 6

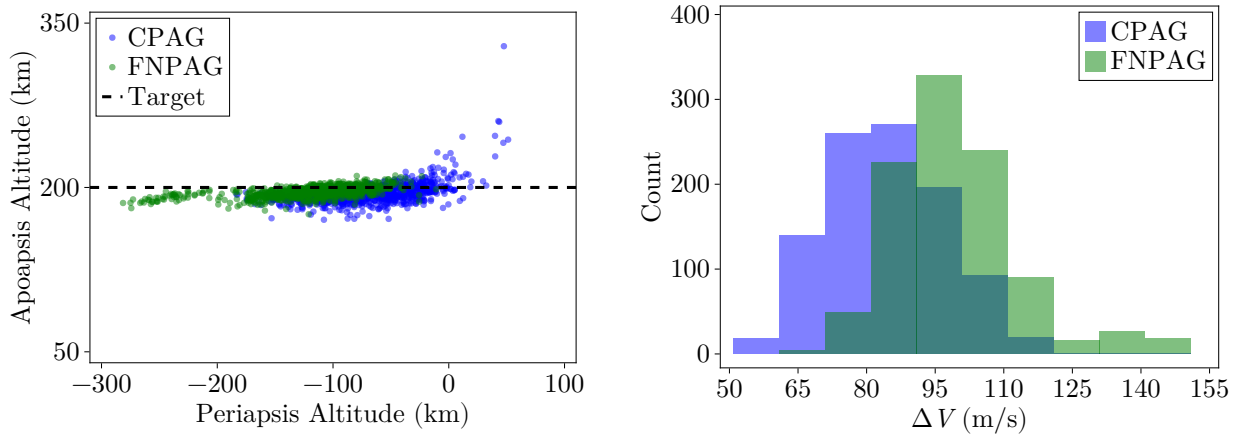

(a) Apoapsis accuracy

(b) In-plane ΔV.

Figure 6.19: Comparison of apoapsis targeting accuracy (a) and in-plane ΔV (b) for Earth aerocapture.

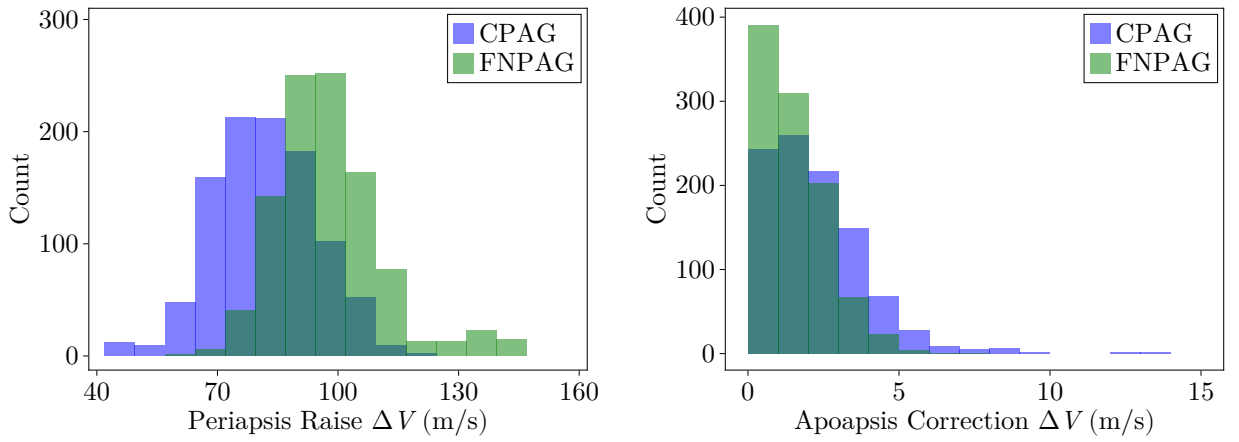

(a) Periapsis raise ΔV.

(b) Apoapsis correction ΔV.

Figure 6.20: Comparison of periapsis raise ΔV (a) and apoapsis correction ΔV (b) for Earth aerocapture.

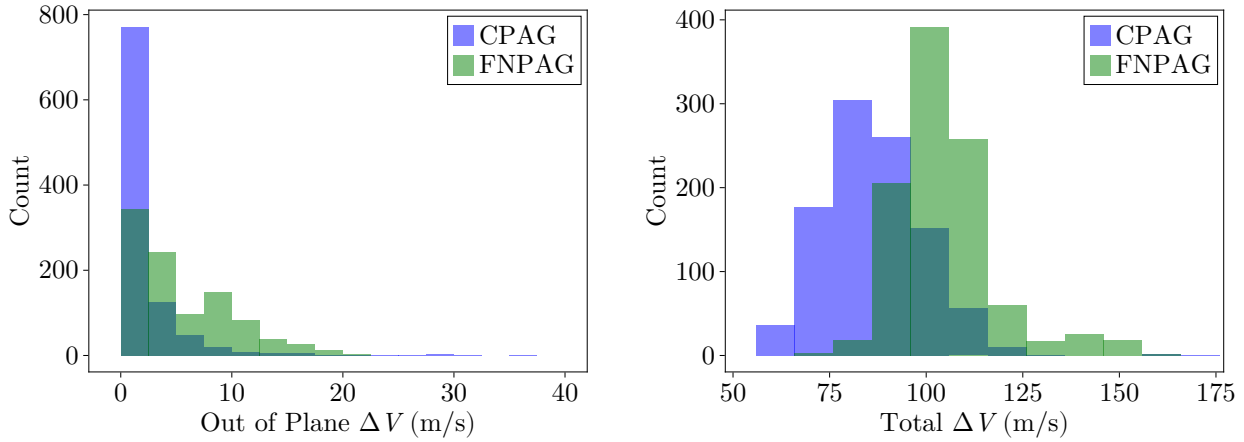

(a) ΔV_3

(b) Total ΔV.

Figure 6.21: Required ΔV for inclination change (a) and total orbit correction (b) for Earth aerocapture.

g is satisfied for 100% of the cases using both algorithms. An overall comparison of the statistics for the aerocapture quantities of interest obtained using both CPAG and FNPAG is presented in Table 6.14.

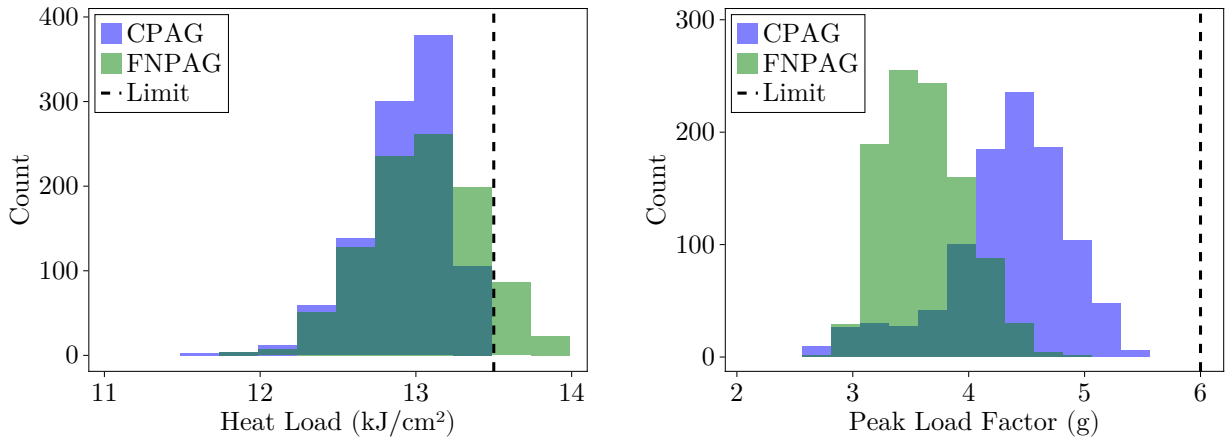

(a) Heat load.

(b) Load factor.

Figure 6.22: Comparison of stagnation-point heat lead (a) and load factor (b) for Earth aerocapture.

Table 6.14: Comparison of Earth aerocapture statistics for CPAG and FNPAG.

CPAG (Capture Rate = 100.0%)	Mean	Std. Dev.	1st Perc.	50th Perc.	99th Perc.
In-Plane ΔV (m/s)	84.56	13.08	59.39	83.58	116.6
Out-of-Plane ΔV (m/s)	1.960	4.507	2.197×10^{-3}	3.112×10^{-1}	23.01
Total ΔV (m/s)	86.52	13.55	61.39	85.23	118.1
Apoapsis Error (%)	1.414×10^{-1}	3.284×10^{-1}	3.167×10^{-3}	1.024×10^{-1}	6.637×10^{-1}
Max. Load Factor (g)	4.332	5.300×10^{-1}	2.842	4.407	5.297
Heat Load (kJ/cm^2)	12.93	2.805×10^{-1}	12.15	12.97	13.42
Computational Time (ms)	19.5	4.55	10.1	20.1	34.4
FNPAG (Capture Rate = 100.0%)	**Mean**	**Std. Dev.**	**1st Perc.**	**50th Perc.**	**99th Perc.**
In-Plane ΔV (m/s)	99.35	14.00	74.36	97.42	145.4
Out-of-Plane ΔV (m/s)	5.415	4.472	1.009×10^{-1}	3.770	18.39
Total ΔV (m/s)	104.7	12.99	84.41	102.7	148.7
Apoapsis Error (%)	7.738×10^{-2}	5.612×10^{-2}	1.933×10^{-3}	6.780×10^{-2}	2.360×10^{-1}
Max. Load Factor (g)	3.623	3.524×10^{-1}	2.978	3.592	4.493
Heat Load (kJ/cm^2)	13.05	3.657×10^{-1}	12.22	13.05	13.84
Computational Time (ms)	5.77	1.91	4.00	5.81	11.2

6.5 Effects of Planetary Oblateness

The effects of J_2 on the exo-atmospheric phase of an aerocapture trajectory and resulting ΔV requirements have yet to be explored in the literature. The apoapsis radii and in-plane ΔV in the previous section are computed using the Keplerian expressions in Eqs. 5.8 and 6.5, respectively. However, these assume that the effects of J_2 after the vehicle exits the planet's atmosphere are negligible. Equation 6.4, on the other hand, presents an alternative expression for the apoapsis radius as a function of the atmospheric exit condition that relaxes the assumption of Keplerian motion in the exo-atmospheric phase. Because the magnitude of the J_2 perturbations depends on the latitude, the user must specify a target latitude ϕ_{target} in Eq. 6.4 which may or may not be known. This section assesses the sensitivity of Eq. 6.4 to the choice of ϕ_{target} and shows how using Eq. 6.4 to target an apoapsis radius instead of the Keplerian expression in Eq. 5.8 can result in more accurate ΔV predictions.

To demonstrate how Eq. 6.4 depends on ϕ_{target}, the Neptune aerocapture scenario in Section 6.4.1 is used. For target inclinations $i_{\text{target}} \in [0, 90]$ degrees, and values of $\phi_{\text{target}} \in [0, 90]$ degrees, CPAG is used to guide the entry vehicle to the target orbit using Eq. 6.4. A value of $J_2 = 3.411 \times 10^{-3}$ is used in both the equations of motion and Eq. 6.4, and CPAG is assumed to have perfect knowledge of its environment (i.e. no dispersions). For each pair of i_{target} and ϕ_{target}, the ΔV required to correct the final orbit is computed using two methods. The first is the standard Keplerian expression in Eq. 6.5, and the second uses a numerical single-shooting approach in order to account for the effects of J_2 on the exo-atmospheric trajectory. Figure 6.23 compares the ΔV computed using each technique assuming an initial latitude of $\phi_0 = 0$ degrees.

For each approach, the computed ΔV shows almost no dependence on the selection of ϕ_{target}. For all pairs of i_{target} and ϕ_{target}, the numerically computed ΔV varies by less than 1%, while there is a noticeable variation in the Keplerian ΔV with the target inclination. This indicates that, under certain conditions, using the Keplerian expression in Eq. 6.5 to compute the required in-plane ΔV can produce substantial errors. Recall that the previous Monte Carlo analysis in Sections 6.4.1–6.4.3

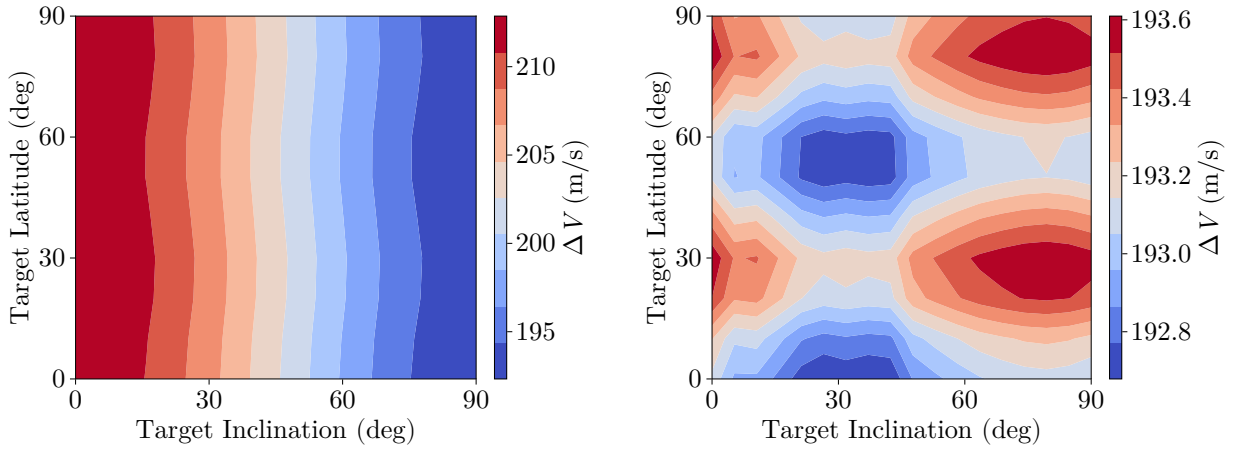

(a) Keplerian (Eq. 6.5)

(b) Numerical.

Figure 6.23: ΔV computed using Keplerian expression (a) and numerical single-shooting approach (b) for Neptune aerocapture for $\phi_0 = 0$ degrees.

assumed $J_2 = 0$ in Eq. 6.4. As such, if the distributions of in-plane ΔV were to be recomputed using the numerical integration approach in this section, they would be shifted toward larger ΔV values, as the Keplerian expressions used to compute the ΔV are inaccurate. The conditions under which Eq. 6.5 approaches the numerically computed ΔV relate to the specific energy at the atmospheric exit condition. Specifically, the local "Keplerian" energy at the atmospheric exit condition must be close to the "true" energy at the resulting apoapsis radius. Here, the "Keplerian" and "true" energy are defined using Eqs. 6.46 and 6.47, respectively

$$\varepsilon_{\text{Keplerian}} = \frac{\mu}{r} - \frac{v^2}{2} \tag{6.46}$$

$$\varepsilon_{\text{true}} = \frac{\mu}{r}\left[1 - \frac{J_2}{2}\left(\frac{R_e}{r}\right)^2\left(3\sin^2\phi - 1\right)\right] - \frac{v^2}{2} \quad . \tag{6.47}$$

Note these are the negative of the standard specific energy in orbital mechanics, which is chosen for consistency with Eq. 6.4. Figure 6.24 shows the time history of the true energy after atmospheric exit for two different cases. For each case, the solid line corresponds to the true energy computed using Eq. 6.47, and the dashed line corresponds to the Keplerian energy computed at atmospheric exit. The first case, denoted by the red lines, corresponds to a scenario in which the Keplerian energy at atmospheric exit ($t = 0$ min.) is close to the true energy at apoapsis ($t \approx 400$ min.). When

this happens, the ΔV computed using Eq. 6.5 and the ΔV computed numerically agree well. The second, case, denoted by the blue lines, corresponds to a scenario in which the Keplerian energy at atmospheric exit is *not* close to the true energy at apoapsis. Here, the Keplerian expression overpredicts the ΔV by approximately 10%.

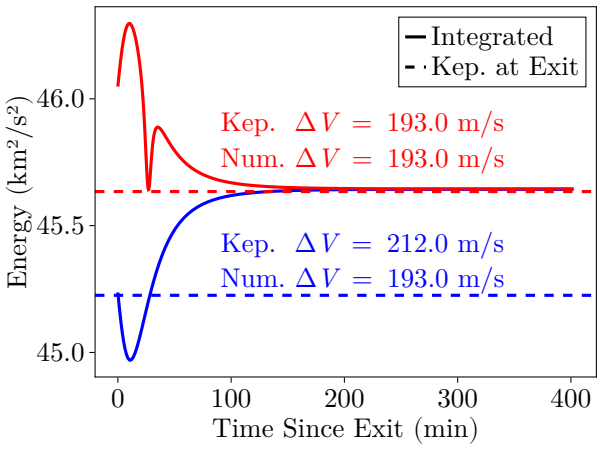

Figure 6.24: Comparison of Keplerian and true energy time histories during exo-atmospheric phase of trajectory.

Because J_2 perturbations depend strongly on the latitude, the same analysis is repeated for initial latitudes in the range $\phi_0 \in [0, 80]$ degrees. Figure 6.25 shows the Keplerian ΔV computed using Eq. 6.5 as a function of the target inclination. Here, depending on the initial latitude, the use of Keplerian expressions can result in overpredictions of ΔV by up to 25%. On the other hand, the numerically computed ΔV remains relatively constant across all scenarios, demonstrating that Eq. 6.4 is able to account for the effects of J_2 on the exo-atmospheric phase of the aerocapture trajectory.

The analysis presented here indicates that existing Keplerian expressions (e.g. Eqs. 5.8 and 6.5) are inaccurate when considering higher-order gravitational effects. Therefore, it is recommended that the energy-based expression in Eq. 6.4 be used for targeting an apoapsis radius in aerocapture guidance algorithms. The form of Eq. 6.4 is general and may be used by any NPC aerocapture algorithm, not just CPAG. Additionally, Eq. 6.4 could be extended to even higher-order gravitational

perturbations if the spherical-harmonic coefficients are known and are sufficiently large.

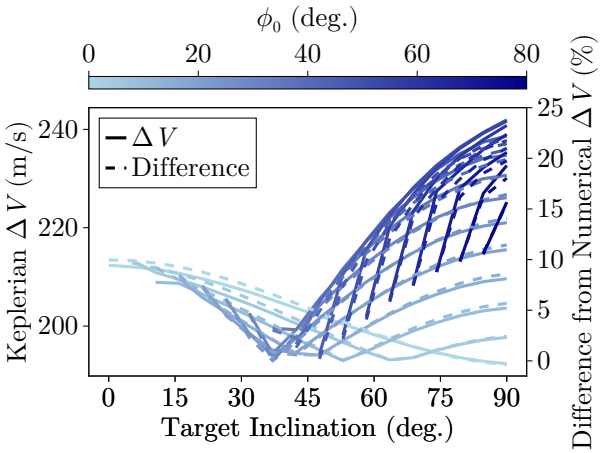

Figure 6.25: ΔV computed using standard Keplerian expression for a variety of initial latitudes and target inclinations.

6.6 Summary

This chapter presented CPAG, an advanced predictor-corrector aerocapture guidance algorithm based in convex optimization. It was shown that replacing the standard, simple one-dimensional optimization problem typically solved in the correction phase of a fully numerical predictor-corrector (NPC) algorithm with the solution to a convex optimization problem allowed the algorithm to produce more complex control profiles. This eliminated the need for a separate lateral logic algorithm to determine the bank-angle direction. CPAG also easily incorporated system-level constraints, such as aerodynamic loading and integrated heat load, into its correction phase, whereas, existing NPC aerocapture guidance algorithms either required separate logic or were unable to do so. The increase in control-profile complexity came at a slight reduction in performance and increase in computational time. An improved balancing between in- and out-of-plane accuracy was achieved by penalizing intermediate inclination errors and using the augmented convex-concave decomposition (ACCD).

Numerical results demonstrated that CPAG was capable of guiding both low and mid lift-to-drag ratio vehicles to their desired orbit with a high level of accuracy in the presence of aerodynamic,

delivery, and atmospheric uncertainty. Monte Carlo analyses were performed to compare the performance of CPAG to the current state-of-the art NPC aerocapture guidance algorithm FNPAG for aerocapture scenarios at Neptune, Mars, and Earth. For Neptune, and Mars, the total ΔV requirements obtained using CPAG were slightly greater than those obtained using FNPAG. At Earth, CPAG produced lower overall total ΔV requirements. A significant improvement in inclination targeting was observed using CPAG, which was attributed to the direct consideration of the final inclination within CPAG's correction phase. CPAG satisfied the imposed heat-load and load-factor constraints for 100% of the cases at all three destinations, whereas, FNPAG produced significantly lower heat-load constraint satisfaction rates. The ability to improve aerothermal constraint satisfaction allows for a more precise sizing of the thermal protection system (TPS). CPAG can easily adapt to the aerothermodynamic constraints, while FNPAG is unable to do so. Overall, the proposed convex-optimization framework was shown to be sufficiently robust to the relevant sources of uncertainty in aerocapture missions. It incorporated system-level constraints in a straightforward way and eliminated the need for a separate lateral logic algorithm to determine bank-angle sign.

Finally, an energy-based expression for targeting an apoapsis radius that can account for perturbations from non-spherical gravitational fields was presented. It was shown that assuming Keplerian motion during the exo-atmospheric phase of the aerocapture trajectory can produce substantial overpredictions in the ΔV required to correct the final orbit. The use of the energy-based expression for targeting an apoapsis radius, on the other hand, was able to relax this assumption and improve the accuracy in the predicted ΔV requirements by almost 25%. For missions where the expected ΔV requirements are large, a 25% improvement in the predictive accuracy would produce a drastic decrease in onboard fuel requirements, allowing for more mass to be allocated to science objectives.

Chapter 7

Stochastic Aerocapture Guidance

"It is not the unknown that should concern us, but the rigidity of our response to it."
— *Dune*

7.1 Introduction

This chapter takes a stochastic optimal control approach to the aerocapture guidance problem, in which uncertainty in the atmospheric density, vehicle aerodynamics, and aerothermodynamic environment are explicitly considered in the design of a feedback control policy. This is substantially different from the approach of CPAG in Chapter 6, which repeatedly solves a deterministic optimal control problem. Instead, a sequence of feedforward controls and feedback gain matrices are optimized offline in a stochastic optimal control setting, which are then implemented as lookup tables in the guidance system. In support of this explicit uncertainty approach, the theory describing how multiple interacting and correlated noise sources affect the system is developed.

7.2 Background

Generally, the approach for obtaining a solution to a trajectory optimization problem may be categorized into one of two methods: indirect methods and direct methods [213]. In the indirect approach, the optimal control profile is indirectly solved for by computing the Hamiltonian and associated adjoint variables using Pontryagin's minimum principle. Vinh et al. has extensively employed this approach for the optimization of lifting entry trajectories [214, 215], however, their

analyses never included any aerothermodynamic or vehicle-scale constraints. Furthermore, solving an optimal control problem with an indirect method is often more difficult, particularly in the presence of inequality constraints, due to the sensitivity of the adjoint dynamics to their initial conditions. Direct methods introduce a more practical approach to solving trajectory optimization problems by reducing the infinite-dimensional optimal control problem to a finite-dimensional parametric optimization problem. Both of these approaches, however, are only made more challenging when uncertainty in the optimization variables and constraints is introduced.

In a stochastic setting, constraints must be treated probabilistically, and a common approach is to use stochastic analogs to deterministic path/terminal constraints called chance constraints [216]. The introduction of uncertainty brings with it a substantial increase in computational complexity, as evaluating chance constraints requires the propagation of multi-variate uncertainty distributions through the nonlinear aerocapture dynamics. Significant effort has been directed toward the development of solution techniques for nonlinear, stochastic optimal control problems in the presence of chance constraints. These include the use of sequential convex programming (SCP), as in the work of Oguri et al. [217, 218] and Ridderhof et al. [219, 74], and the use of higher-order state-transition tensors, as in the case of Boone et al. [220]. In any solution procedure, a model for the relevant sources of uncertainty is required. Determining a high-fidelity estimate of the aerothermodynamic uncertainty distributions, even just at a single flight condition, requires significant computational resources, and the level of uncertainty is not constant along an aerocapture trajectory [45]. It is computationally intractable to model the aerothermodynamic uncertainty at every point along the aerocapture trajectory, but it is also necessary to incorporate the spatial variation of the uncertainty along a trajectory. Ridderhof et al. [74] present an approach for incorporating spatially varying atmospheric density uncertainty in chance-constrained optimal control problems using Gaussian random fields (also known as Gaussian processes). As such, this work utilizes Gaussian-process (GP) based surrogate models to efficiently map the uncertainty at key trajectory points to the entire trajectory domain [140]. The GP model is trained on data obtained from analyses performed using a high-fidelity computational fluid dynamics (CFD) code to account

for the various hypersonic thermochemical nonequilibrium phenomena relevant to aerocapture trajectories.

Overall, this chapter seeks to directly consider the aerothermodynamic uncertainty present in Ice Giant aerocapture scenarios within a stochastic trajectory optimization problem. Gaussian-process based surrogate models trained on data from high-fidelity models are used to efficiently incorporate the uncertainty in the aerothermodynamic heating environments along the aerocapture trajectory, and the aerothermodynamic constraints are enforced using chance constraints. A sequential convex programming approach (SCP) [177] is developed due to its wide use in the aerospace trajectory optimization field [179, 68, 182, 180, 71, 70] and demonstrated ability to efficiently solve chance-constrained trajectory optimization problems [221, 217, 74]. A Neptune aerocapture mission is used for numerical demonstration, however, the proposed methodology could be applied to any aerocapture or planetary entry mission in which significant aerothermodynamic uncertainty exists.

The chapter is organized as follows. Section 7.3 discusses the development of stochastic aerothermodynamic and atmospheric density models using Gaussian processes. A nonlinear, stochastic aerocapture trajectory optimization problem is formulated in Section 7.4, and a solution methodology using sequential convex programming is provided in Section 7.5. A 3,000 case Monte Carlo analysis is performed in Section 7.6, and conclusions are given in Section 7.7.

7.3 Stochastic Aerothermodynamic and Atmosphere Models

This section provides details on the models used to approximate the uncertainty present in aerocapture trajectories. These include the aerodynamics and aerothermodynamics of the spacecraft and the atmospheric density profile.

7.3.1 Stochastic Aerothermodynamic Model for Neptune Aerocapture

The aerothermodynamic constraint considered in this chapter is the stagnation-point convective heat load. Previous work [70] modeled the stagnation-point convective heat flux using a GP as a function of altitude and velocity. However, because this work also considers uncertainty in

the atmospheric density, it is necessary to capture the coupling between the atmospheric density uncertainty and the heat-flux uncertainty. To this end, let the nominal stagnation-point convective heat flux be modeled using a generalized Sutton-Graves [158] form

$$\dot{\bar{q}} = C\rho^a v^b \tag{7.1}$$

where ρ is the freestream density in kg/m^3, v is the freestream velocity in km/s, and C, a, and b are constants. Using the training data in Chapter 4, a least-squares fit of Equation 7.1 is performed. The fit parameters are summarized in Table 7.1, which produce heat-flux values with units of W/cm^2.

Table 7.1: Generalized Sutton-Graves coefficients for Neptune atmosphere.

C	a	b
11.84	0.4897	2.779

The uncertain stagnation-point convective heat flux is given by

$$\dot{q} = \dot{\bar{q}}\,(1 + \delta\dot{q}) \tag{7.2}$$

where $\delta\dot{q}$ is a zero-mean GP taking values as a function of the atmospheric density, ρ, and freestream velocity, v. Based on the uncertainty intervals computed in Chapter 3, and the GP's fit in Chapter 4, the heat-flux variation covariance function is defined using a squared-exponential kernel with variance $\sigma_q^2 = 1.5 \times 10^{-3}$, and lengthscales $l_\rho = 10$ kg/m^3 and $l_v = 1$ km/s. These values produce heat-flux uncertainties of $3\sigma \sim 12\%$, which is consistent with the highest convective-heating uncertainties observed in Chapter 3. The methodology presented in this Chapter does not require that the data used to construct the aerothermodynamic model be derived from high-fidelity CFD simulations. A standard Sutton-Graves correlation [158], could be used just as easily.

Similarly, the lift and drag coefficients are given by

$$C_L = \bar{C}_L\,(1 + \delta_{C_L}) \tag{7.3}$$

$$C_D = \bar{C}_D\,(1 + \delta_{C_D}) \tag{7.4}$$

where δ_{C_L} and δ_{C_D} are zero-mean GP's taking values as a function of the altitude $h = r - R_e$, where R_e is the planet radius, and velocity. The aerodynamic coefficients are not modeled as a function of density because they are relatively invariant at high hypersonic Mach numbers. While each aerodynamic model could be extended to depend on density, doing so would introduce additional complexity, and the coupling between density and aerodynamic coefficients is less critical than that between density and heat flux. Both the lift- and drag-coefficient variations are modeled using a squared-exponential kernel with a variance of $\sigma^2_{C_{L,D}} = 1 \times 10^{-3}$ and lengthscales of $l_h = 100$ km and $l_v = 5$ km/s, which provide aerodynamic uncertainties consistent with the Monte-Carlo analysis performed in Chapter 6.

7.3.2 Stochastic Atmosphere Modeling

Following [74], the atmospheric density is given by

$$\rho = \bar{\rho}\,(1 + \delta\rho) \tag{7.5}$$

where $\bar{\rho}$ is the nominal NeptuneGRAM [121] profile. The density variation $\delta\rho$ is a zero-mean GP taking values as a function of altitude. Based on the NeptuneGRAM model, the density variation covariance function is defined as [74]

$$k_\rho(h_1, h_2) = \exp\left(-\frac{|h_1 - h_2|}{H_{\text{scale}}}\right) \times \begin{cases} b(\min(h_1, h_2)) & \min(h_1, h_2) \leq h_{\text{trans}} \\ \sigma^2_{\rho,\max} & \min(h_1, h_2) > h_{\text{trans}} \end{cases} \tag{7.6}$$

where H_{scale} is the scale height, and

$$b(h) = \sigma^2_{\rho,\max} \exp\left(\frac{h - h_{\text{trans}}}{c_{\text{scale}}}\right) \tag{7.7}$$

The constants h_{trans} and c_{scale} determine the scale of the exponential variance model, and $\sigma^2_{\rho,\max}$ is the maximum density variance. The specific values used in this work are $H_{\text{scale}} = 50$ km, $\sigma_{\rho,\max} = 1 \times 10^{-2}$, $h_{\text{trans}} = 60$ km, and $c_{\text{scale}} = 30$ km. Samples of $\delta\rho$ are compared to density variations computed using NeptuneGRAM in Figure 7.1. It is acknowledged that the GP-based atmosphere model underpredicts the density variability from NeptuneGRAM at high altitudes.

However, the density in this region is low, so some underprediction is acceptable, as it does not significantly affect the trajectory. This is corroborated by the results at the end of the chapter, which demonstrate sufficient robustness to density variability.

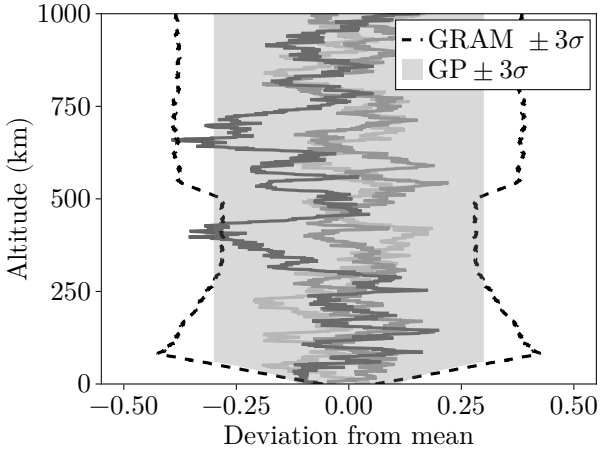

Figure 7.1: Samples from the density variation process with the 3σ confidence interval shaded.

7.4 Stochastic Optimal Control Problem Formulation

This section presents the physical models used in the development of the stochastic aerocapture trajectory optimization problem. The equations of motion governing the atmospheric flight mechanics are provided in Section 7.4.1, and the nonlinear, stochastic aerocapture trajectory optimization problem under aerothermodynamic uncertainty is formulated in Sections 7.4.2 and 7.4.3.

7.4.1 Aerocapture Equations of Motion

The equations of motion over an oblate-spheroid, rotating planet are given in Equation 5.1 of Chapter 5. A gravitational model up to J_2 harmonics is used, and the control is taken to be the bank angle rate $\dot{\sigma}$.

7.4.2 Variable Time of Flight Formulation

The heat-load is sensitive to the duration of the trajectory, so it is of interest to allow the time of flight to vary as needed. A free time of flight problem may be cast as a fixed time of flight problem using a transformation called *generalized time dilation* [195]. In this approach, a strictly increasing, continuously differentiable mapping $t : [0, 1] \mapsto \mathbb{R}_+$, with boundary conditions $t(0) = t_0$ and $t(1) = t_f$, is defined. The derivative of this mapping is defined as

$$s(\tau) = \frac{\mathrm{d}t(\tau)}{\mathrm{d}\tau} = \overset{\circ}{t}(\tau) \tag{7.8}$$

for $\tau \in [0, 1]$, where $\overset{\circ}{(\cdot)}$ denotes the derivative with respect to the non-dimensional time. The term $s(\tau)$ is referred to as the *time dilation factor* and is treated as an additional control input. Next, the state $\boldsymbol{x} = [r, \theta, \phi, v, \gamma, \psi, \sigma]$ is augmented with the heat load and physical time t such that

$$\tilde{\boldsymbol{x}}(\tau) = [\boldsymbol{x}(t(\tau)),\ Q(t(\tau)),\ t(\tau)] \tag{7.9}$$

where $\tilde{\boldsymbol{x}}$ is the "augmented" state. The augmented state dynamics may then be written as

$$\overset{\circ}{\tilde{\boldsymbol{x}}}(\tau) = \begin{bmatrix} \dot{\boldsymbol{x}}(t) \\ \dot{Q}(t) \\ 1 \end{bmatrix} \frac{\mathrm{d}t(\tau)}{\mathrm{d}\tau} = \begin{bmatrix} \boldsymbol{f}(\boldsymbol{x}(t), u(t)) \\ \dot{q}(\boldsymbol{x}(t)) \\ 1 \end{bmatrix} \frac{\mathrm{d}t(\tau)}{\mathrm{d}\tau} = \boldsymbol{F}(\tilde{\boldsymbol{x}}(\tau), \tilde{\boldsymbol{u}}(\tau)) \tag{7.10}$$

where $\dot{q}(\boldsymbol{x}(t))$ is the instantaneous heat flux, and $\tilde{\boldsymbol{u}} = [u(t(\tau)), s(\tau)]$ is the original control augmented with the time dilation factor. The form of Equations 7.9 and 7.10 is chosen for two reasons. First, computing the heat load requires integrating the heat flux along the entire trajectory. While convex approximations of integral quantities exist (e.g. trapezoidal integration [182]), the linearization and discretization procedure discussed in the following sections allows for a more exact enforcement of the heat load constraint in the convex subproblem. Because the reference trajectory, about which the convex subproblem is formed, is produced by numerically integrating the nonlinear equations of motion, a more accurate description of how the heat load changes along the updated trajectory is obtained. Finally, expressing the path constraint variables as an additional state variable allows the proposed framework to be amenable to continuous time constraint satisfaction [222, 195].

Next, the inclusion of the aerothermodynamic uncertainty is considered. Following the work of [74], the nonlinear, stochastic state equations of motion are defined as

$$\overset{\circ}{\tilde{\boldsymbol{x}}}(\tau) = \boldsymbol{F}(\tilde{\boldsymbol{x}}(\tau), \tilde{\boldsymbol{u}}(\tau), \boldsymbol{\Psi}(\tilde{\boldsymbol{x}}(\tau))) \tag{7.11}$$

where $\boldsymbol{\Psi}$ is a vector of the Gaussian processes used to model the uncertainties in the system. For the uncertainties considered in this work, this produced $\boldsymbol{\Psi} = [\delta q, \delta C_L, \delta C_D, \delta\rho]^T$. For a more concise presentation, the tilde notation denoting the augmented state is dropped for the remainder of the chapter.

7.4.3 Nonlinear Optimization Problem Formulation

The initial state is assumed to be Gaussian-distributed about the nominal value such that $\boldsymbol{x}_0 \sim \mathcal{N}(\bar{\boldsymbol{x}}_0, P_0)$, where $\bar{\boldsymbol{x}}_0 = \mathbb{E}[\boldsymbol{x}_0]$ is the nominal initial state and $P_0 = \mathbb{E}[(\boldsymbol{x}_0 - \bar{\boldsymbol{x}}_0)(\boldsymbol{x}_0 - \bar{\boldsymbol{x}}_0)^T]$ is the initial state covariance. Note that, even if the aerothermodynamic environment were not modeled as a stochastic process, the assumption of a random initial condition inherently makes the entire augmented state dynamics stochastic. To help keep the state close to the nominal trajectory, the control is assumed to follow a state-dependent feedback control policy defined at time τ

$$\boldsymbol{u}(\tau) = \bar{\boldsymbol{u}}(\tau) + K(\tau)(\boldsymbol{x}(\tau) - \bar{\boldsymbol{x}}(\tau)) \tag{7.12}$$

where $\bar{\boldsymbol{u}}(\tau) \in \mathbb{R}^m$ is the nominal feedforward control, and $K(\tau) \in \mathbb{R}^{m\times(n+2)}$ are feedback gains to be determined in the numerical optimization process. The control in Equation 7.12 is inherently stochastic due to its dependence on the state $\boldsymbol{x}$. It is emphasized that the form of Equation 7.12 is general and does not assume any parameterization of the control (e.g. zero-order hold, first-order hold, etc.).

The constraints considered in this work consist of both path and terminal constraints. Because the state, and therefore control, are not deterministic, any constraint on the control or state must be treated stochastically. Constraints on uncertain (random) variables are typically referred to as chance constraints, which probabilistically constrain unbounded random variables; for instance,

$$\mathbb{P}[\xi \leq a] \geq 1 - \delta \tag{7.13}$$

where $\delta \in (0,1)$ is the maximum allowable probabilistic bound for the violation of $\xi \leq a$. For example, a value of $\delta = 10^{-3}$ would correspond to 99.9% confidence. A series of path chance constraints are enforced on the control to satisfy bank-angle rate magnitude constraints such that

$$\mathbb{P}[\|\dot{\sigma}_k\|_2 \leq \dot{\sigma}_{\max}] \geq 1 - \delta_{\dot{\sigma}} \quad \forall\ k \tag{7.14}$$

$$\mathbb{P}[\|s_k\|_2 \leq s_{\max}] \geq 1 - \delta_s \quad \forall\ k \tag{7.15}$$

where $\dot{\sigma}_{\max}$ and $s_{\max}$ are the upper bounds on the magnitude of the bank-angle rate and time-dilation factor, respectively.

Overall, it is desired that the vehicle be on a trajectory that allows it to reach a specified apoapsis radius when it exits the planet's atmosphere. Therefore, a terminal constraint is formulated to ensure that the nominal trajectory hits the target apoapsis radius

$$r_a(\bar{\boldsymbol{x}}_f) - r_{a,\text{target}} = 0 \quad . \tag{7.16}$$

The subscript $(\cdot)_f$ is used to denote a quantity at the final time $\tau = \tau_f$, which is taken to be the time at which the vehicle exits the planet's atmosphere. In most work [55, 169, 174, 194], the terminal apoapsis target is expressed by directly computing the apoapsis radius as a function of the final state

$$r_a(\boldsymbol{x}_f) = a\left(1 + \sqrt{1 - \frac{v_f^2 r_f^2 \cos^2\gamma_f}{\mu a}}\right) \tag{7.17}$$

where a is the semi-major axis

$$a = \frac{\mu}{2\mu/r_f - v_f^2} \quad . \tag{7.18}$$

Instead, this work targets a specific apoapsis radius using the energy-like expression

$$\varepsilon(\bar{\boldsymbol{x}}_f) = \frac{\mu}{r_f} - \frac{v_f^2}{2} - \frac{\mu}{r_{a,\text{target}}} + \frac{r_f^2 v_f^2 \cos^2\gamma_f}{2r_{a,\text{target}}^2} = 0 \tag{7.19}$$

which is obtained through the conservation of specific energy and angular momentum from the atmospheric exit condition to the target apoapsis radius. This particular form of the terminal apoapsis constraint is chosen because it produces the same solution as Equation 7.17 and has more

robust numerical properties [71]. A constraint is placed on the nominal initial condition such that

$$\bar{\boldsymbol{x}}(\tau_0) \quad = \quad \bar{\boldsymbol{x}}_0 \tag{7.20}$$

where $\bar{\boldsymbol{x}}_0$ is the nominal initial state.

The cost function used in this work considers minimizing the following expression

$$J = \mathrm{Var}\left[r_a(\boldsymbol{x}(\tau_f))\right] + \int_0^1 \|\bar{\boldsymbol{u}}(\tau)\|_2^2 \,\mathrm{d}\tau + \int_0^1 \left(\boldsymbol{u}(\tau) - \bar{\boldsymbol{u}}(\tau)\right)^T R_u \left(\boldsymbol{u}(\tau) - \bar{\boldsymbol{u}}(\tau)\right) \mathrm{d}\tau + \int_0^1 \left(\boldsymbol{x}(\tau) - \bar{\boldsymbol{x}}(\tau)\right)^T Q_x \left(\boldsymbol{x}(\tau) - \bar{\boldsymbol{x}}(\tau)\right) \mathrm{d}\tau \tag{7.21}$$

where $\mathrm{Var}\left[r_a(\boldsymbol{x}(\tau_f))\right]$ is the apoapsis radius variance, and $Q_x \in \mathbb{S}_+^{(n+2)\times(n+2)}$ and $R_u \in \mathbb{S}_+^{m\times m}$ are positive semi-definite matrices that penalize the deviation in the state and control, respectively. Note that the standard cost function seeking to directly minimize $\int \|\boldsymbol{u}(\tau)\|_2^2 \,\mathrm{d}\tau$ is not applicable due to the stochastic nature of the control in Equation 7.12.

Finally, it must be ensured that the heat load accumulated during the atmospheric flight does not exceed the capabilities of the vehicle's thermal protection system. A property of Gaussian processes is that the outputs are jointly normally distributed, which makes a GP-based aerothermodynamic model amenable to the proposed chance constraint framework. Specifically, it is desired that the final heat load Q_N not exceed some threshold

$$\mathbb{P}[Q(\tau_f) \leq Q_{\max}] \geq 1 - \delta_Q \tag{7.22}$$

where $Q_{\max}$ is the imposed heat-load limit. As such, the nonlinear stochastic optimal control problem to be solved in this work is

Problem 3. *Find the optimal history of $\bar{\mathbf{u}}(\tau)$, $K(\tau)$, and $\bar{\boldsymbol{x}}(\tau)$ that minimize the cost Eq. 7.21 subject to equality constraints on the dynamics Eq. 7.10, terminal apoapsis radius Eq. 7.19, and initial condition Eq. 7.20, and inequality constraints on the control magnitude Eq. 7.14 and 7.15, and terminal heat-load Eq. 7.22.*

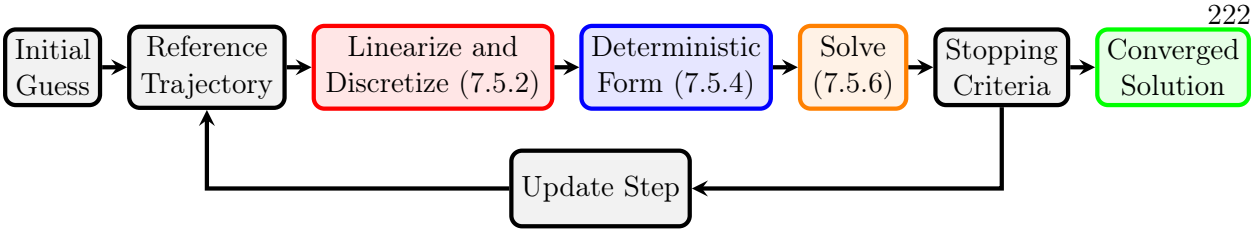

Figure 7.2: Flowchart of the stochastic sequential convex programming (SCP) algorithm.

7.5 Solution via Sequential Convex Programming

Obtaining a solution to Problem 3 is exceptionally difficult due to its stochastic nature and the lack of convexity introduced by the non-linear dynamics and complex path/terminal constraints. Therefore, this work proposes to replace Problem 3 with an approximate convex model, called a *convex subproblem*, that is solved sequentially. A graphical representation of this process is shown in Figure 7.2. In general, the sequential convex programming procedure may be summarized as follows. First, the nonconvexities in Problem 3 are removed by a local linearization about the reference trajectory [182], while the convex elements are retained. The continuous-time, free final time subproblem is then converted into finite-dimensional, fixed final time convex optimization problem by discretizing the time domain. Additional consideration must also be given to the conversion of the stochastic chance constraints into equivalent deterministic forms. Then, the convex subproblem can be solved by a variety of existing convex solvers, and the reference trajectory and control are updated with the solution to the subproblem.

7.5.1 Special Considerations for Interacting Gaussian Processes

In [74], only a single GP was used in the development of the stochastic optimal control problem. Because this work models multiple sources of uncertainty, it is necessary to consider the situation where two of the GP's interact with one another. For example, the drag force is given as

$$D = \frac{1}{2}\rho v^2 C_D A_{\text{ref}} \tag{7.23}$$

where A_{ref} is the vehicle reference area. Substituting the uncertain density and drag coefficient yields

$$D = \frac{1}{2}\bar{\rho}\left(1+\delta\rho\right)v^2\bar{C}_D\left(1+\delta C_D\right)A_{\text{ref}} = \frac{1}{2}\bar{\rho}v^2\bar{C}_D\left(1+\delta\rho+\delta C_D+\delta\rho\delta C_D\right) \quad . \tag{7.24}$$

In general, the interaction term cannot be represented using a GP because the product of two GP's is no longer Gaussian. However, because the solution procedure outlined in the following sections only requires the statistics of the GP's evaluated along the reference trajectory, the product of the two GP's can be treated as Gaussian in *distribution* space. This is analogous to the fact that the product of two Gaussian random *variables* is not Gaussian, but the product of two Gaussian *distributions* is[1] . Here, it is assumed that the GP's are statistically independent. As shown in the following result, the mean and covariance of the product of two GP's only depends on the mean and covariance functions of the GP's in the interaction term.

Proposition 1. *The product of two Gaussian processes $\Psi_1(t)$ and $\Psi_2(t)$ with mean functions $\mu_1(t)$ and $\mu_2(t)$, respectively, has the mean function*

$$\mathbb{E}\left[\Psi_1(t)\Psi_2(t)\right] = \mu_1(t)\mu_2(t) \quad . \tag{7.25}$$

Proof. The GP's are assumed to be statistically independent which yields

$$\mathbb{E}\left[\Psi_1\Psi_2\right] = \mathbb{E}\left[\Psi_1\right]\mathbb{E}\left[\Psi_2\right] = \mu_1\mu_2 \tag{7.26}$$

□

Proposition 2. *The product of two Gaussian processes $\Psi_1(t)$ and $\Psi_2(t)$ with mean functions $\mu_1(t)$ and $\mu_2(t)$, respectively, and covariance functions $k_1(t,\tau)$ and $k_2(t,\tau)$, respectively, has the covariance function*

$$k_3(t,\tau) = k_1(t,\tau)k_2(t,\tau) + \mu_1(t)\mu_1(\tau)k_2(t,\tau) + \mu_2(t)\mu_2(\tau)k_1(t,\tau) \quad . \tag{7.27}$$

Proof. From the definition of the covariance

$$\text{Cov}\left[\Psi_1(t)\Psi_2(t), \Psi_1(\tau)\Psi_2(\tau)\right] = \mathbb{E}\left[\left(\Psi_1(t)\Psi_2(t) - \mu_1(t)\mu_2(t)\right)\left(\Psi_1(\tau)\Psi_2(\tau) - \mu_1(\tau)\mu_2(\tau)\right)\right] \tag{7.28}$$

[1] See `https://mathworld.wolfram.com/NormalProductDistribution.html`.

where the result of Proposition 1 has been used. The above term is simplified to yield

$$\mathbb{E}\left[\Psi_1(t)\Psi_1(\tau)\right]\mathbb{E}\left[\Psi_2(t)\Psi_2(\tau)\right] - \mu_1(t)\mu_1(\tau)\mu_2(t)\mu_2(\tau) \quad . \tag{7.29}$$

By definition, the following holds

$$\mathbb{E}\left[\Psi_1(t)\Psi_1(\tau)\right] = k_1(t,\tau) + \mu_1(t)\mu_1(\tau) \tag{7.30}$$

$$\mathbb{E}\left[\Psi_2(t)\Psi_2(\tau)\right] = k_2(t,\tau) + \mu_2(t)\mu_2(\tau) \quad . \tag{7.31}$$

Substituting this into the above expression produces

$$\left[k_1(t,\tau) + \mu_1(t)\mu_1(\tau)\right]\left[k_2(t,\tau) + \mu_2(t)\mu_2(\tau)\right] - \mu_1(t)\mu_1(\tau)\mu_2(t)\mu_2(\tau) \tag{7.32}$$

which gives the desired result upon simplification. □

To this end, whenever the product of two GP's is encountered in the equations of motion, it is replaced with a separate GP that has mean and covariance functions defined by Equations 7.25 and 7.27. This result is similar to the well-known fact that the product of two Gaussian probability density functions yields a separate Gaussian probability density function.

7.5.2 Linearization and Discretization of Dynamics

The first step in converting Problem 3 into a convex form is to linearize the nonlinear dynamics using a first-order Taylor-series expansion around a nominal trajectory $\hat{\boldsymbol{x}}(\tau)$, control $\hat{\boldsymbol{u}}(\tau)$, and mean disturbance $\hat{\boldsymbol{\mu}}(\tau)$

$$\mathring{\boldsymbol{x}}(\tau) \approx A(\tau)\boldsymbol{x}(\tau) + B(\tau)\boldsymbol{u}(\tau) + G(\tau)\hat{\boldsymbol{\Psi}}(\tau) + \boldsymbol{z}(\tau) \tag{7.33}$$

where $\boldsymbol{x} \in \mathbb{R}^{n+2}$ is the augmented state vector, $\boldsymbol{u} \in \mathbb{R}^m$ is the control, A is the state zero-input dynamics matrix, B is the control matrix, G is a matrix mapping the GP into the state, $\hat{\boldsymbol{\Psi}}$ is the Gaussian process vector evaluated along the nominal trajectory (i.e. $\hat{\boldsymbol{\Psi}} = \boldsymbol{\Psi}(\hat{\boldsymbol{x}})$), and $\boldsymbol{z}$ is a drift term that is produced by the linearization of the nonlinear equations of motion. As described in [74], the nominal trajectory serves as a mapping between time and space, which is used to reduce

the spatially varying Gaussian process to a temporal process. This allows perturbations from the GP to be approximated by a sequence of jointly Gaussian random vectors, the statistics of which are computed *about* the nominal trajectory. Perturbations from the GRF are then approximated by a sequence of jointly Gaussian random vectors, which produces a more tractable subproblem in which the system is only subject to temporal disturbances. The matrices in Equation 7.33 are defined as

$$A(\tau) = \left.\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{x}}\right|_{\hat{\boldsymbol{x}},\hat{\boldsymbol{u}},\hat{\boldsymbol{\mu}}} \tag{7.34a}$$

$$B(\tau) = \left.\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{u}}\right|_{\hat{\boldsymbol{x}},\hat{\boldsymbol{u}},\hat{\boldsymbol{\mu}}} \tag{7.34b}$$

$$G(\tau) = \left.\frac{\partial \boldsymbol{F}}{\partial \boldsymbol{\Psi}}\right|_{\hat{\boldsymbol{x}},\hat{\boldsymbol{u}},\hat{\boldsymbol{\mu}}} \tag{7.34c}$$

$$\boldsymbol{z}(\tau) = \boldsymbol{F}(\hat{\boldsymbol{x}}(\tau),\hat{\boldsymbol{u}}(\tau),\hat{\boldsymbol{\mu}}(\tau)) - A(\tau)\hat{\boldsymbol{x}}(\tau) - B(\tau)\hat{\boldsymbol{u}}(\tau) - G(\tau)\hat{\boldsymbol{\mu}}(\tau) \tag{7.34d}$$

where $\hat{\boldsymbol{\mu}}$ is a vector of the mean functions of the Gaussian processes evaluated along the nominal trajectory (i.e. $\hat{\boldsymbol{\mu}} = \boldsymbol{\mu}(\hat{\boldsymbol{x}})$). Further details on these Jacobians are found in Appendix B.1.

The convective heat flux is modeled using the generalized Sutton-Graves expression in Equation 7.1, which raises the atmospheric density to the power ρ^a. This introduces substantial complexity to the problem, as it involves raising a Gaussian process to the power of a. Instead, before linearizing about the reference, the density dependence in Equation 7.1 is approximated using the tangent-line as a global overestimator

$$\rho^a \leq a\hat{\rho}^{a-1}\rho + (1-a)\hat{\rho}^a \quad \forall \rho, \hat{\rho} > 0 \tag{7.35}$$

which holds for any $a \leq 1$. In Equation 7.35, $\hat{\rho}$ is the density evaluated along the nominal trajectory $\hat{\boldsymbol{x}}$. As a result, the approximate heat-load dynamics are a linear function on the density GP $\delta\rho$. While this introduces additional conservatism into the optimization framework, it is relatively small as demonstrated by the numerical results in Section 7.6.

To convert the continuous-time, infinite-dimensional optimal control problem into an equivalent parameter optimization problem that can be solved by convex optimization algorithms, the

continuous expressions from the previous section must be converted into discrete-time equivalents. To do so, a set of N temporal nodes is chosen such that

$$0 = \tau_1 < \tau_2 < \ldots < \tau_{N-1} < \tau_N = 1 \qquad . \tag{7.36}$$

Recall that by introducing the time dilation factors, the free-time-of-flight problem can always be discretized over the domain $\tau \in [0, 1]$. At each node, let $\Phi(\tau_{k+1}, \tau_k) : \mathbb{R} \mapsto \mathbb{R}^{n \times n}$ denote the state transition matrix (STM) for the system in Equation 6.12. From linear systems theory, the STM evolves according to

$$\overset{\circ}{\Phi}(\tau, \tau_0) = A(\tau)\Phi(\tau, \tau_0), \qquad \Phi(\tau_0, \tau_0) = I_{n \times n} \qquad . \tag{7.37}$$

The control is discretized using a first-order hold [196, 197, 223], which assumes that the control varies linearly between temporal nodes k and $k + 1$

$$\boldsymbol{u}(\tau) = \lambda_k^-(\tau)\boldsymbol{u}(\tau_k) + \lambda_k^+(\tau)\boldsymbol{u}(\tau_{k+1}) \qquad \forall\ \tau \in [\tau_k, \tau_{k+1}) \tag{7.38}$$

where $\lambda_k^-(\tau)$ and $\lambda_k^+(\tau)$ are defined as

$$\lambda_k^-(\tau) = \frac{\tau_{k+1} - \tau}{\tau_{k+1} - \tau_k} \qquad \lambda_k^+(\tau) = \frac{\tau - \tau_k}{\tau_{k+1} - \tau_k} \qquad . \tag{7.39}$$

Utilizing the STM, the continuous-time, linear dynamics can now be transformed into an equivalent discrete-time, linear form such that

$$\boldsymbol{x}_{k+1} = A_k\boldsymbol{x}_k + B_k^-\boldsymbol{u}_k + B_k^+\boldsymbol{u}_{k+1} + \boldsymbol{w}_k + \boldsymbol{z}_k \quad k = 1, \ldots, N-1 \tag{7.40}$$

where the discrete matrices and vectors are defined via

$$A_k = \Phi(\tau_{k+1}, \tau_k) \tag{7.41a}$$

$$B_k^- = A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k)\lambda_k^-(\xi)B(\xi)\mathrm{d}\xi \tag{7.41b}$$

$$B_k^+ = A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k)\lambda_k^+(\xi)B(\xi)\mathrm{d}\xi \tag{7.41c}$$

$$\boldsymbol{w}_k = A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k)G(\xi)\hat{\boldsymbol{\Psi}}(\xi)\mathrm{d}\xi \tag{7.41d}$$

$$\boldsymbol{z}_k = A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k)\boldsymbol{z}(\xi)\mathrm{d}\xi \quad . \tag{7.41e}$$

Here ξ is a dummy variable introduced for the integration of Equation 7.41 over the interval $\tau \in [\tau_k, \tau_{k+1}]$. Following [74], the Gaussian disturbance term $\boldsymbol{w}_k$ has mean and covariance

$$\mathbb{E}[\boldsymbol{w}_k] = A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k)G(\xi)\hat{\boldsymbol{\mu}}(\xi)\mathrm{d}\xi \tag{7.42}$$

$$\mathrm{Cov}(\boldsymbol{w}_k, \boldsymbol{w}_\ell) = A_k \left(\int_{\tau_k}^{\tau_{k+1}} \int_{\tau_\ell}^{\tau_{\ell+1}} \Phi^{-1}(\xi_1, \tau_k)G(\xi_1)\hat{K}(\xi_1, \xi_2)G^T(\xi_2)\Phi^{-T}(\xi_2, \tau_\ell)\mathrm{d}\xi_2\mathrm{d}\xi_1 \right) A_k^T \tag{7.43}$$

where $\hat{K}(t, \tau) = \mathrm{diag}\left[k_1(t, \tau), k_2(t, \tau), \ldots\right]$ is a diagonal matrix with elements corresponding to the covariance functions of the Gaussian processes evaluated along the nominal trajectory, and ξ_1 and ξ_2 are additional integration dummy variables.

The dynamics of the mean state evolve according to

$$\mathbb{E}[\boldsymbol{x}_{k+1}] = \bar{\boldsymbol{x}}_{k+1} = A_k\bar{\boldsymbol{x}}_k + B_k^-\bar{\boldsymbol{u}}_k + B_k^+\bar{\boldsymbol{u}}_{k+1} + \boldsymbol{z}_k + \bar{\boldsymbol{w}}_k \tag{7.44}$$

where $\bar{\boldsymbol{w}}_k$ is computed using Equation 7.42. To express the discrete-time dynamics of the covariance, the change in the state around the mean is first written as

$$\boldsymbol{x}_{k+1} - \bar{\boldsymbol{x}}_{k+1} = \delta\boldsymbol{x}_{k+1} = A_k\delta\boldsymbol{x}_k + B_k^-\delta\boldsymbol{u}_k + B_k^+\delta\boldsymbol{u}_{k+1} + \delta\boldsymbol{w}_k \tag{7.45}$$

where $\delta\boldsymbol{w}_k = \boldsymbol{w}_k - \bar{\boldsymbol{w}}_k$. Using the feedback-control policy in Equation 7.12 and recognizing that $\delta\boldsymbol{u}_k = \boldsymbol{u}_k - \bar{\boldsymbol{u}}_k = K_k\delta\boldsymbol{x}_k$, Equation 7.45 may be written as

$$\delta\boldsymbol{x}_{k+1} = (I - B_k^+K_{k+1})^{-1}\left[(A_k + B_k^-K_k)\delta\boldsymbol{x}_k + \delta\boldsymbol{w}_k\right] \quad . \tag{7.46}$$

Because the initial state is assumed to be Gaussian-distributed with mean $\bar{\boldsymbol{x}}_0$ and covariance P_0, it then follows that the covariance evolves through the discrete-time linear dynamics as

$$P_{k+1} = (I - B_k^+K_{k+1})^{-1}\left[(A_k + B_k^-K_k)P_k(A_k + B_k^-K_k)^T + \Sigma_k\right](I - B_k^+K_{k+1})^{-T} \tag{7.47}$$

where $\Sigma_k = \mathrm{Cov}(\boldsymbol{w}_k, \boldsymbol{w}_k)$ is the Gaussian process covariance matrix from Equation 7.43 where the integration domain is $[\tau_k, \tau_{k+1}] \times [\tau_k, \tau_{k+1}]$. Simplifying Equation 7.47 yields

$$P_{k+1} - P_{k+1}K_{k+1}^T B_k^{+^T} - B_k^+ K_{k+1}P_{k+1} + B_k^+ K_{k+1}P_{k+1}K_{k+1}^T B_k^{+^T} = A_k P_k A_k^T + A_k P_k K_k^T B_k^{-^T} + B_k^- K_k P_k A_k + B_k^- K_k P_k K_k^T B_k^{-^T} + \Sigma_k \quad . \quad (7.48)$$

Equation 7.48 is non-convex in P_k and K_k. To transform the covariance dynamics into an equivalent, convex form, the following transformations are introduced [224]

$$U_k = K_k P_k \quad (7.49)$$

$$Y_k = U_k P_k^{-1} U_k^T \quad . \quad (7.50)$$

This produces the following form of the covariance dynamics

$$P_{k+1} = A_k P_k A_k^T + A_k U_k^T B_k^{-^T} + B_k^- U_k A_k^T + B_k^- Y_k B_k^{-^T} + U_{k+1}^T B_k^{+^T} + B_k^+ U_{k+1} - B_k^+ Y_{k+1} B_k^{+^T} + \Sigma_k \quad (7.51)$$

which is now convex in the decision variables $P_k \succeq 0$, U_k, and $Y_k \succeq 0$. According to [224], the following linear matrix inequality (LMI) must be imposed in order for Equation 7.50 to be satisfied

$$\begin{bmatrix} P_k & U_k^T \\ U_k & Y_k \end{bmatrix} \succeq 0 \quad (7.52)$$

which is convex and is obtained by taking the Schur compliment of Equation 7.50. In physical terms, when solved to optimality, Y_k is the covariance of the control. After solving for the optimal sequence of P_k, U_k, and Y_k, the feedback gain matrices are obtained by $K_k = U_k P_k^{-1}$.

7.5.3 Treatment of Correlated Noise

The form of the covariance dynamics in Equation 7.51 is only valid for systems subject to Brownian motion disturbances [224]. Clearly, this is not the case when the noise is modeled using a GP, as the disturbances at each time step are correlated to all other time steps. Unfortunately, including this correlation using the block-covariance form of Equation 7.51 yields a result that is non-convex. To see this, let

$$\tilde{B}_k = \left(I - B_k^+ K_{k+1}\right)^{-1} \quad (7.53)$$

$$\tilde{A}_k \quad = \quad A_k + B_k^- K_k \quad . \tag{7.54}$$

The cross-covariance between the state $\boldsymbol{x}_k$ and disturbance $\boldsymbol{w}_k$ is thus given by

$$\text{Cov}(\boldsymbol{x}_k, \boldsymbol{w}_k) = \mathbb{E}\left[\tilde{B}_k \tilde{A}_k \delta\boldsymbol{x}_k \delta\boldsymbol{w}_k^T \tilde{B}_k^T\right] \quad . \tag{7.55}$$

Substituting Equation 7.45 yields the following result

$$\text{Cov}(\boldsymbol{x}_k, \boldsymbol{w}_k) = \tilde{B}_k \left[\sum_{i=1}^{k-1} \left(\prod_{j=k}^{j=i+1} \tilde{A}_j \tilde{B}_{j-1}\right) \Sigma_{i,k}\right] \tilde{B}_k^T = M_k \tag{7.56}$$

which assumes that the initial state $\boldsymbol{x}_{k=1}$ is uncorrelated from the disturbances $\boldsymbol{w}_k$. The covariance matrix $\Sigma_{i,k} = \text{Cov}\left(\boldsymbol{w}_i, \boldsymbol{w}_k\right)$ is computed using Equation 7.43. This form is clearly non-convex due to the multiplicative terms involving K_k.

Failing to account for the cross-covariance terms in the covariance dynamics leads to lossy propagation of the covariance, which is not desirable for an aerocapture guidance algorithm. To maintain the convex form of 7.51, but still incorporate the cross-covariance terms, this work proposes to relax M_k into the covariance dynamics using the solution of the previous subproblem. Therefore, after the first iteration of the SCP procedure, Equation 7.51 becomes

$$P_{k+1} = A_k P_k A_k^T + A_k U_k^T B_k^{-^T} + B_k^- U_k A_k^T + B_k^- Y_k B_k^{-^T} + U_{k+1}^T B_k^{+^T} + \\ B_k^+ U_{k+1} - B_k^+ Y_{k+1} B_k^{+^T} + \Sigma_k + M_k + M_k^T \tag{7.57}$$

where the additional M_k^T accounts for the symmetric nature of the GP cross-covariance matrix.

7.5.4 Convex Constraint Formulation

Because chance constraints are probabilistic by nature, equivalent deterministic forms must be derived. First, a convex approximation of the terminal constraint on the apoapsis radius in Equation 7.19 is obtained using a first-order Taylor series expansion about the reference such that

$$\varepsilon(\bar{\boldsymbol{x}}_f) \approx \left.\frac{\partial \varepsilon}{\partial \boldsymbol{x}}\right|_{\text{ref}} \cdot (\bar{\boldsymbol{x}}_N - \hat{\boldsymbol{x}}_N) + \varepsilon(\hat{\boldsymbol{x}}_N) = 0 \qquad . \tag{7.58}$$

The initial condition constraint in Equation 7.20 is simply

$$\bar{\boldsymbol{x}}_{k=1} \quad = \quad \bar{\boldsymbol{x}}_0 \quad . \tag{7.59}$$

Equivalent forms of the chance constraints on the control magnitude in Equations 7.14 and 7.15 are obtained using the results of Ref. [219]. Letting $\boldsymbol{y} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$ be an m-dimensional random vector with $m = 1$ or $m = 2$, then an equivalent form of Equation 7.14 or 7.15 is

$$\|\boldsymbol{\mu}\|_2 + \sqrt{\lambda_{\max}(\Sigma)}\sqrt{2\log\frac{1}{\delta}} \leq \rho \Leftrightarrow \mathbb{P}\left[\|\boldsymbol{y}\|_2 \leq \rho\right] \geq 1 - \delta \tag{7.60}$$

where $\rho > 0$, and $0 < \delta < 1$. A proof of Equation 7.60 may be found in Ref. [219]. Since the control is normally distributed with mean $\bar{\boldsymbol{u}}_k$ and covariance $P_{u_k} = Y_k$, the chance constraints on the bank-angle rate and time dilation factor magnitude in Equations 7.14 and 7.15 are expressed as

$$\|\bar{\dot{\sigma}}_k\|_2 + \sqrt{\lambda_{\max}(Y_k)}\sqrt{2\log\frac{1}{\delta_{\dot{\sigma}}}} \quad \leq \quad \dot{\sigma}_{\max} \quad \forall\ k = 1, \ldots, N \tag{7.61}$$

$$\|\bar{s}_k\|_2 + \sqrt{\lambda_{\max}(Y_k)}\sqrt{2\log\frac{1}{\delta_s}} \quad \leq \quad s_{\max} \quad \forall\ k = 1, \ldots, N \quad . \tag{7.62}$$

Unfortunately, the square root in Equations 7.61 and 7.62 makes them non-convex. Therefore, this work takes the approach of [225] by introducing the variables $\rho_{u,k}$, $\chi_{u,k}$, $\rho_{s,k}$, $\chi_{s,k} \geq 0$ such that

$$\rho_{u,k} I_m - Y_k \quad \succeq \quad 0 \qquad \forall\ k = 1, \ldots, N \tag{7.63}$$

$$\rho_{s,k} I_m - Y_k \quad \succeq \quad 0 \qquad \forall\ k = 1, \ldots, N \tag{7.64}$$

$$\rho_{u,k} \quad = \quad \chi_{u,k}^2 \quad \forall\ k = 1, \ldots, N \tag{7.65}$$

$$\rho_{s,k} \quad = \quad \chi_{s,k}^2 \quad \forall\ k = 1, \ldots, N \quad . \tag{7.66}$$

This allows Equations 7.61 and 7.62 to be written as

$$\|\bar{\dot{\sigma}}_k\|_2 + \chi_{u,k}\sqrt{2\log\frac{1}{\delta_{\dot{\sigma}}}} \quad \leq \quad \dot{\sigma}_{\max} \quad \forall\ k = 1, \ldots, N \tag{7.67}$$

$$\|\bar{s}_k\|_2 + \chi_{s,k}\sqrt{2\log\frac{1}{\delta_s}} \quad \leq \quad s_{\max} \quad \forall\ k = 1, \ldots, N \quad . \tag{7.68}$$

However, Equations 7.65 and 7.66 are now non-convex. As such, they are approximated using a first-order Taylor series expansion about a reference $\bar{\chi}_k$

$$\rho_{u,k} \quad = \quad \bar{\chi}_{u,k} + 2\bar{\chi}_{u,k}\left(\chi_{u,k} - \bar{\chi}_{u,k}\right) + \nu_{u,k} \quad \forall\ k = 1, \ldots, N \tag{7.69}$$

$$\rho_{s,k} \quad = \quad \bar{\chi}_{s,k} + 2\bar{\chi}_{s,k}\left(\chi_{s,k} - \bar{\chi}_{s,k}\right) + \nu_{s,k} \qquad \forall\, k = 1, \ldots, N \tag{7.70}$$

where $\nu_{u,k}$ and $\nu_{s,k}$ are slack variables introduced due to the linearization of Equations 7.65 and 7.66. The initial guesses for $\bar{\chi}_{u,k}$ and $\bar{\chi}_{s,k}$ are taken to be $0.3u_{\max}$ and $0.3s_{\max}$, respectively, as recommended by [225].

Next, the constraint on the integrated heat load at the final time must be addressed. Because the heat load is included in the state vector, it is assumed to be normally distributed. Therefore, the following relationship holds

$$\mathbb{P}\left[Q_f \leq Q_{\max}\right] = \text{cdf}\left(\frac{Q_{\max} - \mathbb{E}\left[Q_f\right]}{\sqrt{\text{Cov}(Q_f)}}\right) \tag{7.71}$$

which allows the constraint in Equation 7.22 to be written as

$$\mathbb{P}[Q_N \leq Q_{\max}] \geq 1 - \delta_Q \Leftrightarrow \mathbb{E}\left[Q_f\right] + \text{cdf}^{-1}(1-\delta_Q)\sqrt{\text{Cov}(Q_f)} \leq Q_{\max} \tag{7.72}$$

where $\text{cdf}^{-1}(\cdot)$ is the inverse function of the standard Gaussian cumulative distribution function. Substituting the decision variable P_k into Equation 7.72 yields

$$\boldsymbol{\xi}_Q^T \bar{\boldsymbol{x}}_N + \text{cdf}^{-1}(1-\delta_Q)\sqrt{\boldsymbol{\xi}_Q^T P_k \boldsymbol{\xi}_Q} \leq Q_{\max} \tag{7.73}$$

where $\boldsymbol{\xi}_Q \in \mathbb{R}^{n+2}$ is used to query the element of the state $\bar{\boldsymbol{x}}_k$ that corresponds to the heat load. Again, the square root in Equation 7.73 makes it non-convex. Instead, following [226], Equation 7.73 is written as a difference of convex (DC) functions by squaring both sides

$$\left(\text{cdf}^{-1}(1-\delta_Q)\right)^2 \boldsymbol{\xi}_Q^T P_k \boldsymbol{\xi}_Q - \left(Q_{\max} - \boldsymbol{\xi}_Q^T \bar{\boldsymbol{x}}_N\right)^2 \quad \leq \quad 0 \tag{7.74}$$

$$Q_{\max} - \boldsymbol{\xi}_Q^T \bar{\boldsymbol{x}}_N \quad \geq \quad 0 \tag{7.75}$$

where Equation 7.75 is needed to enforce equivalence with Equation 7.73. DC constraints are only convex if the second term is affine in the decision variables [227] which is not the case for Equation 7.74. Non-convex DC constraints may be handled using the convex-concave procedure (CCP) [228]. The CCP linearizes the second term around the reference trajectory $\hat{\boldsymbol{x}}$ such that

$$\left(\text{cdf}^{-1}(1-\delta_Q)\right)^2 \boldsymbol{\xi}_Q^T P_k \boldsymbol{\xi}_Q - \left[\left(Q_{\max} - \boldsymbol{\xi}_Q^T \hat{\boldsymbol{x}}_N\right)^2 - 2\boldsymbol{\xi}_Q^T\left(Q_{\max} - \boldsymbol{\xi}_Q^T \hat{\boldsymbol{x}}_N\right)\left(\bar{\boldsymbol{x}}_N - \hat{\boldsymbol{x}}_N\right)\right] \leq 0 \tag{7.76}$$

which is affine in P_N and $\bar{\boldsymbol{x}}_N$. Finally, to alleviate *artificial infeasibility* concerns, a slack variable ν_Q is introduced to relax Equation 7.76 to the form

$$\left(\text{cdf}^{-1}(1-\delta_Q)\right)^2 \boldsymbol{\xi}_Q^T P_k \boldsymbol{\xi}_Q - \left[\left(Q_{\max} - \boldsymbol{\xi}_Q^T \hat{\boldsymbol{x}}_N\right)^2 - 2\boldsymbol{\xi}_Q^T \left(Q_{\max} - \boldsymbol{\xi}_Q^T \hat{\boldsymbol{x}}_N\right)\left(\bar{\boldsymbol{x}}_N - \hat{\boldsymbol{x}}_N\right)\right] \leq \nu_Q \quad . \tag{7.77}$$

The slack variable ν_Q is then heavily penalized in the cost function to encourage constraint satisfaction at convergence.

To express the variance in the final apoapsis, the terminal apoapsis is first approximated around the reference as

$$r_a(\boldsymbol{x}_N) \approx r_a(\hat{\boldsymbol{x}}_N) + \left.\frac{\partial r_a}{\partial \boldsymbol{x}}\right|_{\text{ref}} (\boldsymbol{x}_N - \hat{\boldsymbol{x}}_N) \tag{7.78}$$

where the partial derivative is of the Keplerian apoapsis radius expression in Equation 7.17, not the energy-based expression in Equation 7.19, as Equation 7.19 does not explicitly solve for the apoapsis radius. Because the state is assumed to be normally distributed, it follows that the variance of the apoapsis radius is

$$\text{Var}(r_a(\boldsymbol{x}_f)) = S P_N S^T \tag{7.79}$$

where $S = \left.\dfrac{\partial r_a}{\partial \boldsymbol{x}}\right|_{\text{ref}}$. Using the cyclic property of the trace operator, a deterministic, convex form of the cost function in Equation 7.21 may be written as

$$J_{\text{cvx}} = S P_N S^T + \alpha_0 \sum_{k=1}^{N} \|\bar{\boldsymbol{u}}_k\|_2^2 + \sum_{k=1}^{N} \text{tr}(R_{u,k} Y_k) + \sum_{k=1}^{N-1} \text{tr}(Q_{x,k} P_k) + \alpha_1 \|\boldsymbol{\nu}\|_1 + \alpha_2 \nu_Q \tag{7.80}$$

where $\alpha_i \in \mathbb{R}_+$ are penalty weights, and $R_{u,k}$, $Q_{x,k}$ are weighting matrices for the control and state, respectively, and $\boldsymbol{\nu} = [\boldsymbol{\nu}_u, \boldsymbol{\nu}_s]$. As recommended by [74], the state covariance is penalized based on the local dynamic pressure

$$Q_{x,k} = \frac{1}{\hat{q}^2} \left(\frac{\partial q}{\partial \boldsymbol{x}}\right)^T \left(\frac{\partial q}{\partial \boldsymbol{x}}\right) \tag{7.81}$$

where $q = \frac{1}{2}\rho v^2$ is the dynamic pressure, and $\hat{q}$ is the dynamic pressure evaluated along the reference trajectory $\hat{\boldsymbol{x}}$.

7.5.5 Scaling of the Optimization Variables

In spacecraft trajectory optimization problems, the magnitude of the optimization variables can vary significantly. For example, in Eq. 5.1, the radius $r \sim \mathcal{O}(10^6)$, whereas, the bank angle rate $\dot{\sigma} \sim \mathcal{O}(10^0)$. This large discrepancy in variable magnitude can lead to significant buildup of numerical error from iteration to iteration of the optimization process. It has been shown that proper scaling of the optimization variables can increase both the accuracy and efficiency of the optimization process [202]. For this reason, the optimization variables for the mean state and control are scaled by introducing the following affine transformation [191]

$$\bar{\boldsymbol{x}}_k = S_x \tilde{\bar{\boldsymbol{x}}}_k + \boldsymbol{c}_x \tag{7.82a}$$

$$\bar{\boldsymbol{u}}_k = S_u \tilde{\bar{\boldsymbol{u}}}_k + \boldsymbol{c}_u \tag{7.82b}$$

where S_x, S_u, S_p are diagonal, square matrices, $\boldsymbol{c}_x$, $\boldsymbol{c}_u$, $\boldsymbol{c}_p$ are vectors that center the variables of commensurate size, and the tilde notation $\tilde{(\cdot)}$ denotes the scaled variable. The specific form of the i-th diagonal entry of S and ith entry of c for some quantity a is

$$S_{a,ii} = \frac{a_{i,\text{upper}} - a_{i,\text{lower}}}{2} \tag{7.83a}$$

$$c_{a,i} = \frac{a_{i,\text{upper}} + a_{i,\text{lower}}}{2} \quad . \tag{7.83b}$$

Here $a_{i,\text{upper}}$ and $a_{i,\text{upper}}$ are approximate upper and lower bounds on the i-th component of quantity a that will produce the scaled variables to lie in the range $\tilde{a}_i \in [-1, 1]$. A summary of the scaling values used in this work for the mean state and control is provided in Table 7.2. The optimization variables for the covariance dynamics are scaled via

$$P_k = D \tilde{P}_k D^T \tag{7.84}$$

$$U_k = \tilde{U}_k D \quad . \tag{7.85}$$

Here, the scaling matrix is taken to be $D = P_f^{1/2}$, where $P_f^{1/2}$ is a matrix that satisfies $P_f = P_f^{1/2} \left(P_f^{1/2}\right)^T$. Since P_f is not known ahead of time, it should be selected to be close to the desired closed-loop final covariance. Note that the matrix Y_k is not scaled, as recommended by [229].

Table 7.2: Values used to scale mean optimization variables in stochastic SCP algorithm.

Parameter	$a_{i,\text{upper}}$	$a_{i,\text{lower}}$
r (km)	1000	200
θ (deg)	0.2	-0.2
ϕ (deg)	0.2	- 0.2
v (km/s)	30	20
γ (deg)	10	-12
ψ (deg)	0.2	-0.2
σ (deg)	180	0.0
Q	100	0.0
t (s)	800	0.0
u (deg/s)	20	-20
s (s)	800	0.0

7.5.6 Complete Convex Subproblem

The first-order approximation of the dynamics is only valid in a neighborhood around the reference trajectory, so the solution must be kept close enough to the reference trajectory to avoid a phenomenon known as *artificial unboundedness*. This is accomplished by introducing the trust region constraints

$$\|\bar{\boldsymbol{x}}_k - \hat{\boldsymbol{x}}_k\|_2 + \|\bar{\boldsymbol{u}}_k - \hat{\boldsymbol{u}}_k\|_2 \leq \eta_k \quad \forall\ k = 1, \ldots, N \tag{7.86}$$

where the $\eta_k \in \mathbb{R}_+$ are the trust region radii. The convex cost function is augmented with the slack variables used to relax the control magnitude and heat-load constraint, and the trust region radii, such that the final cost function is written as

$$J_{\text{cvx,aug}} = J_{\text{cvx}} + \boldsymbol{w}_{tr}^T \boldsymbol{\eta} \tag{7.87}$$

$\boldsymbol{w}_{tr} \in \mathbb{R}^N_{++}$ is a vector of trust region weights. Finally, the solution variables are scaled by substituting $\bar{\boldsymbol{x}}_k \leftarrow S_x \tilde{\bar{\boldsymbol{x}}}_k + \boldsymbol{c}_x$, $\bar{\boldsymbol{u}}_k \leftarrow S_u \tilde{\bar{\boldsymbol{u}}}_k + \boldsymbol{c}_u$, $P_k \leftarrow D P_k D^T$, and $U_k \leftarrow \tilde{U}_k D$ wherever they appear, so that the decision variables become the scaled state, control, and covariance. The convex subproblem to be solved at each iteration in the SCP process is

Problem 4. *Find the values of* $\tilde{\bar{\boldsymbol{u}}}_k$, $\tilde{\bar{\boldsymbol{x}}}_k$, $\tilde{P}_k$, $\tilde{U}_k$, Y_K, χ_k, ν_k, ν_Q, *and* η_k *that minimize the cost Eq. 7.87 subject to equality constraints on the mean dynamics Eq. 7.44, covariance dynam-*

ics Eq. 7.51, initial condition Eq. 7.59, terminal apoapsis radius Eq. 7.58, and control magnitude Eqs. 7.69 and 7.70, and inequality constraints on the control magnitude Eqs. 7.63, 7.64, 7.67 and 7.68, terminal heat-load Eqs. 7.75 and 7.76, and trust regions Eq. 7.86.

7.5.7 Sequential Convex Programming Formulation

The SCP requires an initial reference trajectory about which the convex subproblem is formed. To improve the convergence behavior of the algorithm, the deterministic version of Problem 4 is first solved. This deterministic solution then provides a starting point for the stochastic optimization process. At each iteration, the nonlinear equations of motion are linearized around the current reference trajectory to formulate a convex optimization problem. Solving the convex subproblem yields an updated solution that is used to refine the feedforward control inputs. Subsequently, the nonlinear equations of motion are integrated numerically using the updated controls to produce a new reference trajectory. After the first iteration, the cross-covariance terms in Equation 7.56 are relaxed into the covariance dynamics constraint using the solution from the previous iteration. This iterative procedure continues until the change in the mean trajectory $\|\bar{\boldsymbol{x}}_k - \hat{\boldsymbol{x}}_k\|_2$ is below some threshold, and the slack variables $\boldsymbol{\nu}$ and ν_Q are sufficiently small. Algorithm 2 summarizes this process.

Algorithm 2: Sequential convex programming approach for solving Problem 2

Input: Initial mean state $\bar{\boldsymbol{x}}_0$ and forward control $\bar{\boldsymbol{u}}_k$, initial covariance P_0
for n = 1: max iterations **do**
  Integrate mean dynamics to get reference trajectory $\hat{\boldsymbol{x}}$ using $\boldsymbol{x}_0$ and $\bar{\boldsymbol{u}}$
  Compute linear system matrices A_k, B_k^+, B_k^-, $\boldsymbol{z}_k$, and $\boldsymbol{w}_k$ using Eqs. 7.41 and 7.42
  Compute GP covariance Σ_k using Eq. 7.43
  if $n \geq 2$ **then**
    Compute cross-covariance terms M_k using Eq. 7.56
  end
  Solve convex subproblem for $\tilde{\bar{\boldsymbol{x}}}_k$, $\tilde{\bar{\boldsymbol{u}}}_k$, $\tilde{P}_k$, $\tilde{U}_k$, and Y_k
  $\bar{\boldsymbol{u}}_k \mathrel{+}= S_u \tilde{\bar{\boldsymbol{u}}}_k + \boldsymbol{c}_u$, Update feedforward control profile
  $K_k = U_k P_k^{-1}$, Compute feedback gain matrices
end
Output: feedforward control, $\bar{\boldsymbol{u}}_k$ and feedback gain matrices, K_k

7.6 Results

This section applies the stochastic sequential convex programming algorithm described in Section 7.5 to a Neptune aerocapture scenario. Uncertainty in the vehicle's initial state, stagnation-point heat flux, aerodynamic coefficients, and the atmospheric density is considered. An optimal sequence of feedforward controls and feedback matrices are obtained in order to guide the entire initial distribution to a terminal distribution while satisfying probabilistic constraints on the control magnitude and terminal heat load. The values used to scale the cost function in Equation 7.87 are summarized in Table 7.3. An approximate final covariance of $P_f = \text{diag}[\sigma_r^2, \sigma_\theta^2, \sigma_\phi^2, \sigma_v^2, \sigma_\gamma^2, \sigma_\psi^2, \sigma_\sigma^2, \sigma_Q^2, \sigma_t^2]$, where $3\sigma_r = 10,000$ m, $3\sigma_\theta = 0.2$ deg, $3\sigma_\phi = 0.2$ deg, $3\sigma_v = 100$ m/s, $3\sigma_\gamma = 0.1$ deg, $3\sigma_\psi = 0.2$ deg, $3\sigma_\sigma = 180$ deg, $3\sigma_Q = 40$ kJ/cm^2, and $3\sigma_t = 50$ s, is used to scale the covariance dynamics in the subproblem.

Table 7.3: Values used to scale cost function in SCP algorithm.

Term	Value
Control, α_0	1×10^{-3}
Control Cov., $R_{u,k}$	$\text{diag}([1 \times 10^{-3},\ 1 \times 10^{-3}])$
State Cov., $Q_{x,k}$	From Eq. 7.81
Control slack, α_1	100
Heat load slack, α_2	1,000
Trust region, $\boldsymbol{w}_{tr}$	1×10^{-3}

7.6.1 Neptune Aerocapture Reference Mission

The entry state is derived from an existing Neptune aerocapture systems analysis [13], and the nominal initial conditions $\bar{\boldsymbol{x}}_0 = [r_0, \phi_0, \theta_0, v_0, \gamma_0, \psi_0, \sigma_0]$ are summarized in Table 7.4. The initial flight path angle is chosen to be in the middle of the nominal aerocapture corridor for a $100,000 \times 3,000$ km target orbit. The atmospheric density is modeled using the nominal profile from NeptuneGRAM [121]. To be consistent with the geometry studied in Ref. [45], the entry vehicle is modeled as an MSL-style, 70-degree sphere cone with a base diameter of 5 meters and

a nose radius of 1.25 m. The nominal lift-to-drag ratio and ballistic coefficient are $L/D = 0.3$ and $\beta = 110$ kg/m^2, respectively, as summarized in Table 7.5. To provide the initial guess for the stochastic sequential convex programming algorithm, the deterministic version of Problem 4 is first solved. A value of $N = 50$ is used, and Problem 4 is solved in Julia [205] using the `Convex.jl` [188] toolbox and `MOSEK` solver [230]. Figure 7.3 shows the nominal trajectory and corresponding heat-load profile. The resulting apoapsis altitude error is less than 100 m, and the final heat load is 84 kJ/cm^2.

Table 7.4: Nominal inertial initial conditions for Neptune mission.

h_0 (km)	ϕ_0 (deg.)	θ_0 (deg.)	v_0 (km/s)	γ_0 (deg.)	ψ_0 (deg.)	σ_0 (deg.)
1000	0	0	29.0	-11.64	0	10

Table 7.5: Nominal mass and aerodynamic properties for entry vehicle.

Mass (kg)	C_L	C_D	A_{ref} (m^2)
3563.7	0.495	1.65	19.6

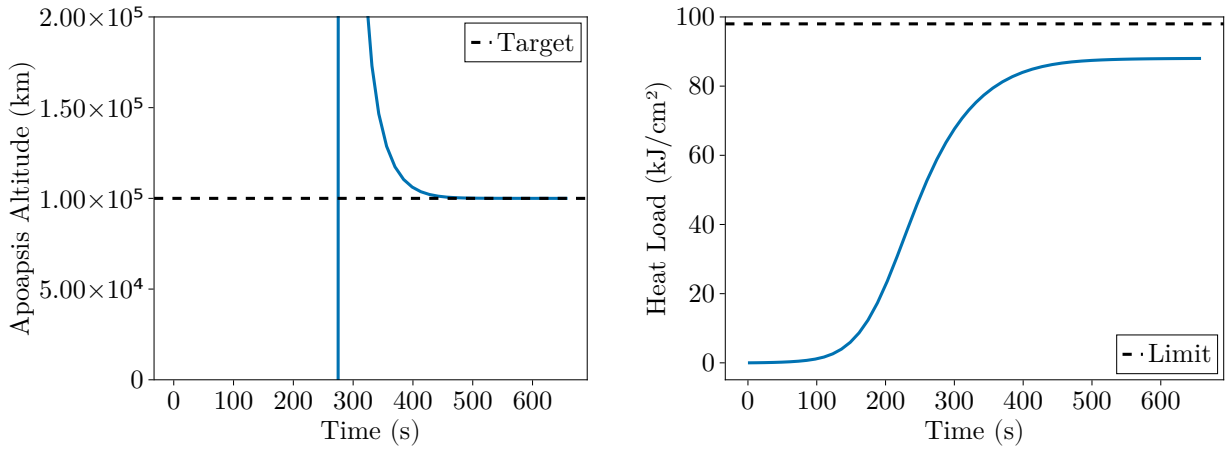

(a) Apoapsis altitude.

(b) Heat load.

Figure 7.3: Nominal apoapsis altitude (a) and heat-load profile (b) obtained from solution to deterministic problem.

7.6.2 Uncertainties

Four key uncertainties are considered: initial conditions, stagnation-point heat flux, aerodynamics, and the atmospheric density. Each is integrated into the algorithm to guide the optimization process toward solutions that are both feasible and robust under real-world variability.

Initial condition uncertainty is modeled as a multivariate Gaussian distribution, with the mean state taken as the values provided in Table 7.4, and the covariance is specified using the initial 3σ values in Table 7.6. Only the initial altitude, velocity, flight-path angle, and bank angle are dispersed, as these states define the majority of the aerothermodynamic environment. This probabilistic representation captures plausible deviations in the initial conditions, ensuring that the optimization accounts for potential variations in the trajectory's starting state due to delivery error. Uncertainty in the stagnation-point heat flux, aerodynamics coefficients, and atmospheric density is modeled using the GP models described in Section 7.3

The control magnitude and time-dilation factor magnitude are constrained such that the probability of exceeding the allowable limits is below a specified threshold, ensuring that the requested control authority is feasible, and the time dilation produces a physical time mapping. Similarly, the final heat load is constrained to ensure that aerothermal loads remain below the limitations of the TPS. These chance constraints enable the algorithm to generate solutions that balance performance and robustness, providing trajectories that are not only optimal under nominal conditions but also resilient to significant uncertainties. This ensures safety and reliability in both control actuation and thermal protection. The constraint boundary and confidence interval for each chance constraint are summarized in Table 7.7.

Table 7.6: Initial covariance used in optimization and Monte Carlo analysis.

	h (km)	v (km/s)	γ (deg.)	σ (deg.)	Q (kJ/cm^2)	t (s)
$3\sigma_0$	1.0	0.1	0.2	10	1.0	1.0

Table 7.7: Control and problem-specific constraints for Neptune aerocapture.

Parameter	Max. Value	Confidence Interval
Bank-angle rate, $\dot{\sigma}_{\max}$	20 deg/s	$\delta_{\dot{\sigma}} = 1 - 0.9973$
Time dilation factor, $s_{\max}$	800 s	$\delta_s = 1 - 0.9973$
Heat load, $Q_{\max}$	98 kJ/cm^2	$\delta_Q = 1 - 0.9973$

7.6.3 Monte Carlo Analysis

Algorithm 2 is run for a total of five iterations, after which the optimized feedback gains are computed. The termination is dictated by user feedback. To assess the performance of the optimized control policy, a 3,000 case Monte Carlo analysis is performed. For each of the cases, the initial conditions are sampled from the distribution in Table 7.6, the stagnation-point heat flux and aerodynamic coefficients are modeled using a single sample from the respective GP models, and the atmosphere is dispersed using NeptuneGRAM. The simulation for each case is terminated when the spacecraft returns to the original altitude or when the final time is reached, whichever occurs first. The dispersed profiles are integrated through the nonlinear equations of motion in two different scenarios. In the first scenario, only the optimized feedforward control is used. This is referred to as the "open-loop" profile in the coming results. The second scenario uses both the feedforward control and feedback gain matrices and is referred to as "closed-loop." Using the first-order hold parameterization on the control, both the feedforward and feedback controls are computed using linear interpolation at intermediate times $\tau \in [\tau_k, \tau_{k+1}]$. Figures 7.4–7.9 summarize the results. In Figures 7.4–7.8, the nominal profile is shown in black, 20 of the 3,000 dispersed profiles are shown in red, and the 99% confidence intervals are shown in the shaded region. Figure 7.9 shows all 3,000 data points.

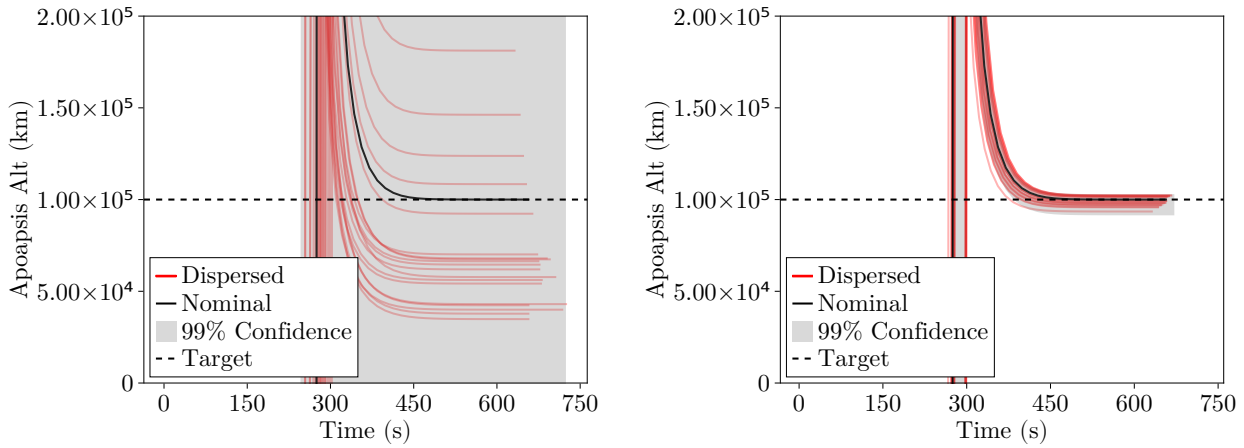

(a) Open loop.

(b) Closed loop.

Figure 7.4: Apoapsis-altitude profiles produced using open-loop (a) and closed-loop (b) control policies.

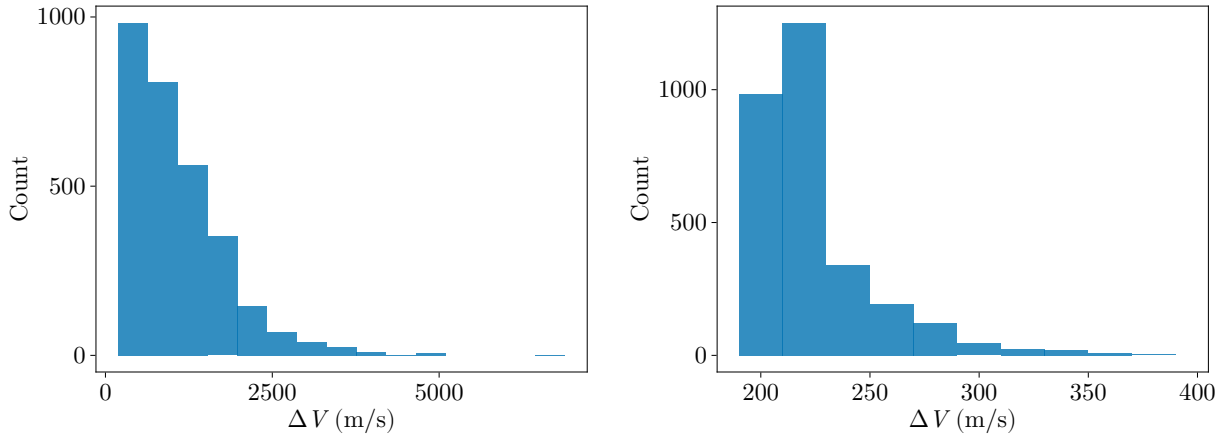

(a) Open loop.

(b) Closed loop.

Figure 7.5: Distribution of ΔV using open-loop (a) and closed-loop (b) control policies.

In the open-loop results, only the nominal case is able to reach the target apoapsis altitude, whereas the closed-loop control successfully guides the entire initial distribution to the desired target. From Figure 7.4, the average open-loop apoapsis altitude error is in excess of 100,000 km, while the average closed-loop apoapsis altitude error is less than 40 km. Additionally, the open-loop control results in 36 cases that remain hyperbolic (i.e. do not capture), while the closed-loop policy

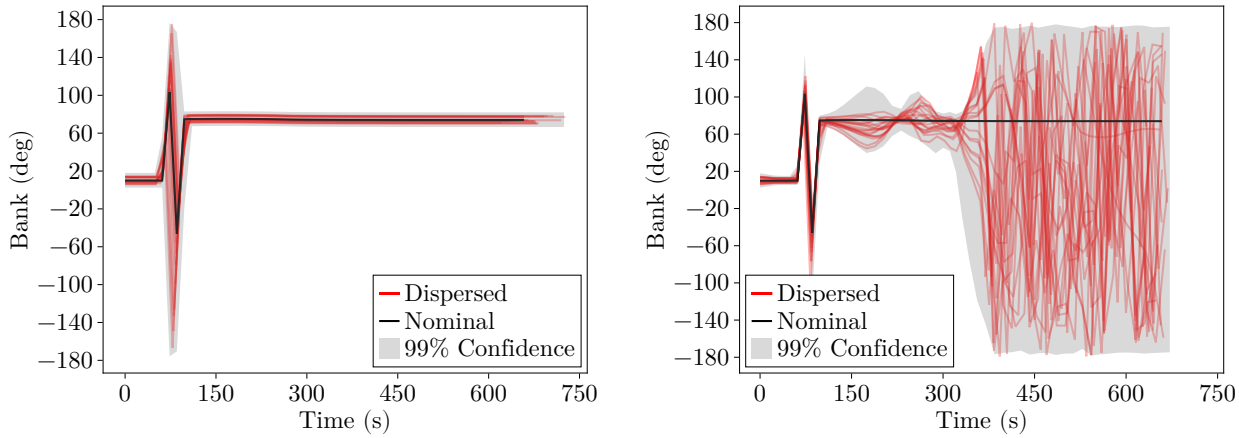

(a) Open loop.

(b) Closed loop.

Figure 7.6: Bank-angle profiles produced using open-loop (a) and closed-loop (b) control policies.

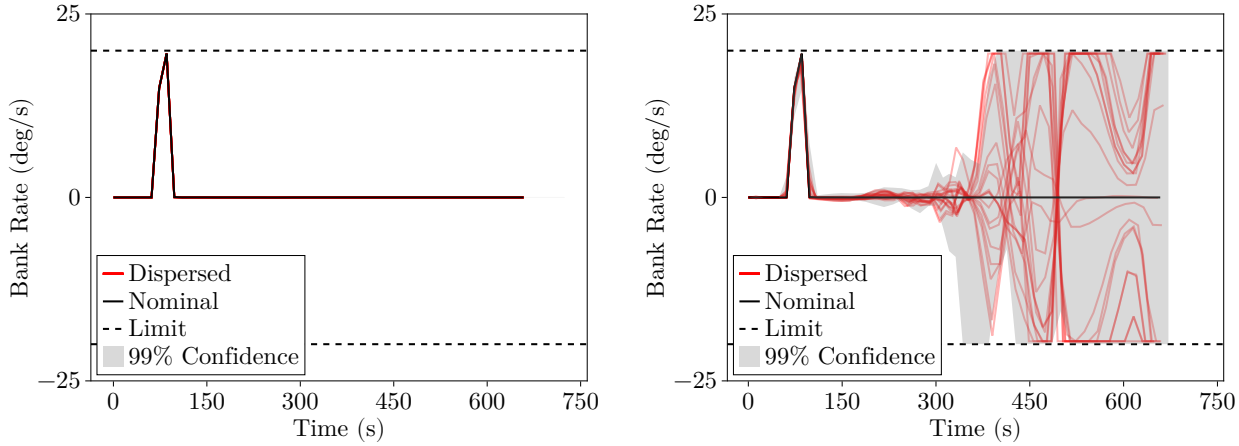

(a) Open loop.

(b) Closed loop.

Figure 7.7: Bank-angle rate profiles produced using open-loop (a) and closed-loop (b) control policies.

produces a 100% capture rate. The substantial increase in apoapsis accuracy is reflected in Figure 7.5, which compares the ΔV requirements for two control policies. For the closed- and open-loop control policies, the 99-th percentile ΔV requirements are 3502 m/s and 337 m/s, respectively, indicating that the closed-loop policy produces over an order of magnitude improvement.

The nominal bank-angle profile resembles a bang-bang form, with a sharp rise around $t = 85$

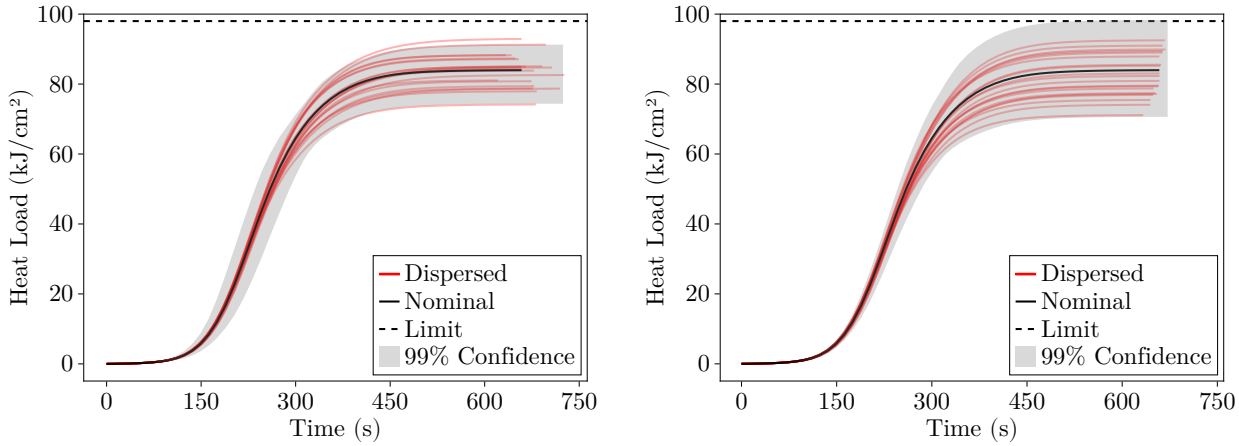

(a) Open loop. (b) Closed loop.

Figure 7.8: Heat-load profiles produced using open-loop (a) and closed-loop (b) control policies.

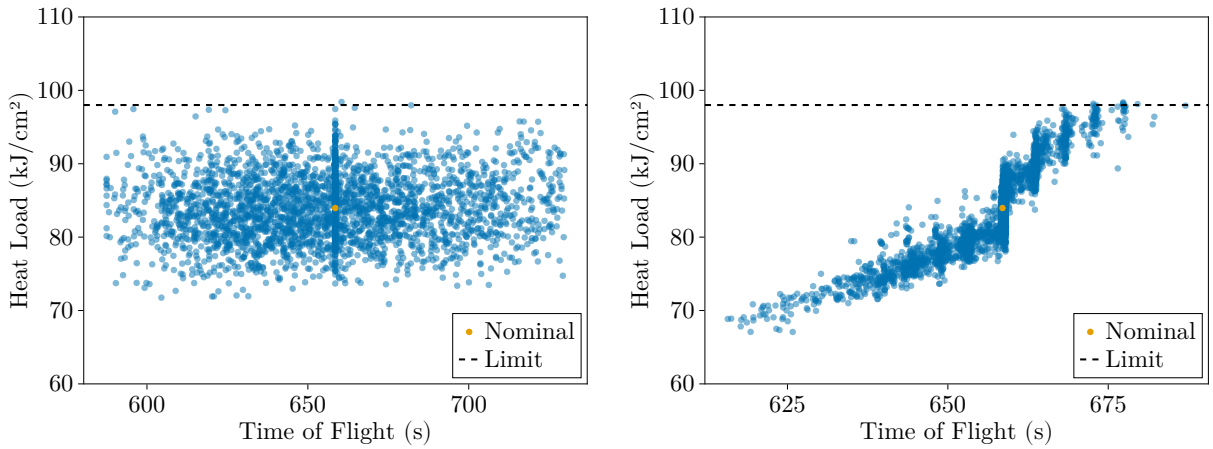

(a) Open loop. (b) Closed loop.

Figure 7.9: Distributions in final time and heat load using open-loop (a) and closed-loop (b) control policies.

seconds, after which it settles at a constant value of approximately 75 degrees. A bang-bang bank-angle profile is known to minimize the ΔV required to correct the in-plane orbital elements [55]. As such, the trends observed in Figure 7.6 demonstrate that choosing to target an apoapsis radius, and not directly minimize ΔV, is not resulting in significant loss of optimality. The effect of the feedback control on the bank angle is clearly observed in Figure 7.6. Because the open-loop

scenario only uses the nominal feedforward control, the resulting distribution of open-loop bank-angle profiles is relatively small. The variation in Figure 7.6a is solely a product of the uncertainty in the initial bank angle. In contrast, the closed-loop bank-angle profiles in Figure 7.6b exhibit greater variability. Feedback control is acting directly on the bank-angle rate, so this behavior is expected and demonstrates the ability of the feedback controller to respond to dispersions in the initial conditions and heat flux. Additionally, the first-order hold parameterization of the bank-angle rate is apparent, as all bank-angle profiles are smooth. Some jaggedness is observed in Figure 7.6b due to the wrapping of the bank angle to the range $\sigma \in [-180, 180]$ degrees. A zero-order hold on the bank-angle rate would also produce smooth bank-angle profiles (although slightly less smooth than a first-order hold), however, it would produce a discontinuous bank-angle rate (i.e. infinite bank acceleration) which cannot be achieved by a spacecraft's control systems.

From Figures 7.7, the bank-angle rate magnitude chance constraint is satisfied for the open-loop *and* closed-loop profiles. The nominal bank-angle rate peaks just below the 20 deg/s limit, preserving margin for dispersions. Because the open-loop results only use the nominal feedforward bank-angle rate and time-dilation factor, it is expected that there is no control magnitude constraint violation. Generally, the closed-loop bank-angle rate profiles remain close to the nominal feedforward profile. Small, relatively, constant deviations are observed which produce the smoothly varying bank-angle profiles in Figure 7.6b. The bank-angle rate profiles produced in all cases are well within the capabilities of existing entry systems [194].

Interestingly, the open-loop and closed-loop control policies produce a similar level of heat-load constraint satisfaction, as shown in Figures 7.8 and 7.9. Specifically, the open-loop policy results in seven of the 3,000 cases exceeding the heat-load limit of 98 kJ/cm^2, and the closed-loop policy produces six cases that exceed the limit. These correspond to 99.76% and 99.80% constraint satisfaction rates, respectively, which are both within the specified 3σ confidence interval. As illustrated in Figure 7.9, the closed-loop heat-load shows significant correlation with the final time of flight, as expected. On the other hand, Figure 7.9a shows that the open-loop policy produces a much broader and less correlated distribution of the time of flight and the heat load. This difference

arises because open-loop control produces greater variability in the altitudes and velocities traversed along the trajectory, which are stronger influences on the final heat load than variations in the final time. Finally, the control magnitude and time-dilation factor chance constraints are satisfied in 99.97%, and 100% of the 3,000 cases, respectively. Table 7.8 summarizes the constraint satisfaction rates for all chance constraints in this analysis.

Table 7.8: Chance constraint satisfaction rates using closed-loop control.

Constraint	Max. Value	Specified Satisfaction Rate (%)	Actual Satisfaction Rate (%)
Heat load	98 kJ/cm^2	99.73	99.80
Bank-angle rate	20 deg/s	99.73	99.97
Time-dilation	800 s	99.73	100.0

The developed algorithm demonstrates sufficient robustness to the uncertainties under investigation, effectively managing variability in initial conditions, stagnation-point heat flux, aerodynamics, and atmospheric density, while satisfying chance constraints on critical parameters. By explicitly coupling the aerocapture trajectory to the aerothermodynamic environment, the algorithm integrates the vehicle's dynamics and aerothermal loads in a more cohesive and realistic manner. This explicit coupling enables the stochastic SCP framework to address challenges in a highly coupled system, ensuring reliable trajectory design and thermal protection even under significant uncertainty. The direct consideration of uncertainty differs from existing closed-loop controllers, such as the fully numerical predictor-corrector guidance (FNPAG) [55] algorithm, which are unable to account for aerothermodynamic constraints or aerothermodynamic uncertainty.

Table 7.9 provides a final comparison of the statistics for all key variables derived from the stochastic algorithm and FNPAG (as discussed in Chapter 6). It's worth noting that the heat-load statistics for FNPAG do not incorporate uncertainties in the aerothermodynamic environment; including these would increase FNPAG's total heat-load. Even without accounting for heat-flux uncertainty, FNPAG exceeds the heat-load constraint in approximately 50% of cases, whereas the stochastic algorithm meets the constraint with 99.8% confidence. Overall, the closed-loop con-

trol policy and FNPAG yield comparable ΔV requirements and apoapsis accuracy, with FNPAG performing slightly better. Since the closed-loop controller explicitly addresses aerothermodynamic uncertainty, relaxing the heat-load constraint and related uncertainties could enhance the stochastic algorithm's performance in terms of ΔV and apoapsis targeting.

Table 7.9: Comparison of Neptune aerocapture statistics for the stochastic algorithm and FNPAG.

Closed Loop (Capture Rate = 100.0%)	Mean	Std. Dev.	1st Perc.	50th Perc.	99th Perc.
In-Plane ΔV (m/s)	224.1	29.59	192.3	217.1	337.6
Apoapsis Error (%)	1.690	1.398	3.232×10^{-2}	1.427	6.892
Heat Load (kJ/cm^2)	83.94	6.637	70.63	83.27	96.91
Open Loop (Capture Rate = 98.8%)	**Mean**	**Std. Dev.**	**1st Perc.**	**50th Perc.**	**99th Perc.**
In-Plane ΔV (m/s)	1071	723.6	204.7	909.0	3502
Apoapsis Error (%)	140.7	1400	7.570×10^{-1}	36.94	1628
Heat Load (kJ/cm^2)	83.97	4.536	74.38	83.80	94.59
FNPAG (Capture Rate = 99.8%)	**Mean**	**Std. Dev.**	**1st Perc.**	**50th Perc.**	**99th Perc.**
In-Plane ΔV (m/s)	219.6	216.8	187.9	202.6	306.7
Apoapsis Error (%)	1.543	9.503	1.612×10^{-2}	7.443×10^{-1}	6.293
Heat Load (kJ/cm^2)	98.03	2.236	93.94	98.10	102.2

7.7 Summary

This chapter developed a robust and efficient stochastic aerocapture guidance algorithm designed to respond to uncertainties in delivery state, aerothermodynamics, and atmospheric density. The algorithm leveraged Gaussian-process (GP) surrogates to model uncertainty in the entry vehicle's stagnation-point convective heating, aerodynamic coefficients, and atmospheric properties, which provided a computationally efficient and accurate representation of the uncertainties as a function of their relative input spaces. The stochastic optimal control problem was solved using a stochastic sequential convex programming (SSCP) framework, in which both feedforward control inputs and feedback gain matrices were optimized. Chance constraints were enforced to ensure probabilistic satisfaction of critical constraints—such as control authority and terminal heat load—while guiding the spacecraft to its target orbit. A distinguishing feature of this approach was the incor-

poration of variable time of flight within the feedback control design, enabling dynamic adaptation to uncertainties and enhancing the framework's flexibility.

By directly coupling aerocapture dynamics with aerothermal uncertainty, this method established an efficient stochastic framework that captured the complex interplay between trajectory and aerothermodynamic environment. A 3,000-case Monte Carlo analysis was used to assess the algorithm's robustness for a Neptune aerocapture scenario. Open-loop results exhibited inconsistent orbit targeting, with significant deviations from the desired apoapsis altitude and frequent violations of heat-load constraints. In contrast, the closed-loop control strategy successfully mitigated all uncertainty sources, achieved precise orbit insertion, and maintained constraint compliance within specified tolerances. Furthermore, this approach required substantially less ΔV for final orbit adjustments compared to the closed-loop policy, with over an order of magnitude improvement observed when using the closed-loop policy.

Chapter 8

Conclusions

"What we know is a drop, what we don't know is an ocean."
— *Isaac Newton*

This chapter provides an overview of the research conducted in this dissertation. It begins by briefly outlining the methodologies and key results from each preceding chapter. Then, it highlights the distinct contributions this work offers to the field. Finally, it explores the broader trajectory of this research and offers specific suggestions for future investigations.

8.1 Dissertation Summary

Chapter 2 described the physical concepts and numerical methodologies required to model the thermochemical nonequilibrium phenomena that occur behind the strong shock waves produced along an aerocapture trajectory. It began by exploring the physical processes in a hypersonic environment, highlighting their importance for aerocapture vehicles, and clarified the difference between thermochemical equilibrium and nonequilibrium. It then described two numerical tools: LeMANS, a CFD code used to predict aerodynamic and convective-heating environments, and NEQAIR, which modeled the radiative heating. Lastly, it outlined the mathematical foundations of two surrogate-modeling methods: polynomial chaos expansion (PCE) and Gaussian process regression (GPR). It presented a global sensitivity analysis (GSA) approach based on PCE and explained the core application of GPR in probabilistic surrogate modeling, emphasizing their importance for aerothermodynamics uncertainty quantification (UQ).

In Chapter 3, an aerothermodynamic UQ/GSA study was presented for an entry capsule performing aerocapture at Neptune. Convective heating showed low uncertainty (4.10–11.7%), primarily influenced by hydrogen electron-impact-ionization (EII) rate coefficient and collision integral variability. Radiative heating, significant at velocities above 25 km/s, exhibited high uncertainty (>160%), driven mainly by H_2 dissociation rate coefficients, with an increased shoulder sensitivity to hydrogen EII. Variability in the freestream CH_4 concentration, which is currently speculated to strongly affect the radiative heating, produced little impact on the radiative-heating uncertainties. Including $H_2{}^+$ in the model notably reduced radiative heating by up to 40%. The findings of Chapter 3 suggested a need for conservative thermal protection system (TPS) design due to radiative uncertainties and a call for further research into hydrogen-related reactions.

A Gaussian Process Regression (GPR) framework was presented in Chapter 4 for creating efficient surrogate models capable of approximating the high-fidelity aerothermodynamic uncertainty predicted in Chapter 3 along an entire aerocapture trajectory. These models accurately predicted key quantities, and their associated uncertainties, including stagnation-point heat flux, aerodynamic coefficients, and surface heat-flux distributions. A multi-output GPR approach was utilized to identify underlying data correlations, and a multi-fidelity approach enabled the efficient blending of low- and high-fidelity data, balancing cost and accuracy. The surrogates were then integrated into a trajectory simulation framework, showcasing their ability to quantify hysteresis effects from aerothermodynamic uncertainty. To ensure physical consistency, GPR kernel functions were modified to enforce invariances, such as symmetry in aerodynamic coefficients. These modifications greatly improved the predictive accuracy, especially when extrapolating beyond the training data.

Chapter 5 laid out the fundamental concepts of aerocapture trajectories and the goals of aerocapture guidance. It started by detailing the equations of motion for simulating a spacecraft's trajectory, accounting for gravitational and aerodynamic forces, and introduced the "entry corridor" concept, which bounds the feasible trajectory space and helps inform the selection of atmospheric-entry conditions. It then explained the fundamental goals of guidance, including a description of the fully numerical predictor-corrector aerocapture guidance (FNPAG) algorithm. The chapter

concluded by presenting the fundamental aspects of convex optimization and explained how sequential convex programming (SCP) is used for trajectory optimization, including several introductory tutorials for using SCP to solve both deterministic and stochastic trajectory-optimization problems.

The first aerocapture guidance algorithm developed in this dissertation was presented in Chapter 6. It introduced the convex predictor-corrector aerocapture guidance (CPAG) algorithm, a predictor-corrector aerocapture guidance algorithm built on the principles of SCP. By solving a convex optimization problem in the correction phase, CPAG was able to reason about more complex control profiles, removing the need for a separate lateral logic algorithm to manage out-of-plane error. It also seamlessly integrated system-level constraints like aerodynamic loading and heat load into the correction phase -— something that existing NPC algorithms like FNPAG cannot do. Numerical tests showed CPAG effectively guided vehicles with low to mid lift-to-drag ratios to their target orbits under aerodynamic, delivery, and atmospheric uncertainties, with Monte Carlo analyses comparing it to FNPAG for Neptune-, Mars-, and Earth-aerocapture scenarios. CPAG slightly underperformed FNPAG for in-plane ΔV requirements at Neptune and Mars but outperformed it at Earth, while showcasing a substantial improvement in inclination targeting and constraint satisfaction rate relative to FNPAG. The convex framework managed uncertainties effectively and made constraint integration straightforward. Finally, an energy-based method for targeting an apoapsis radius that accounted for non-spherical gravity perturbations was presented. It was shown that, for many conditions, standard Keplerian expressions are inaccurate, and using the energy-based expressed improved ΔV prediction accuracy by up to 25%.

An alternative, stochastic optimization approach, was taken for the aerocapture guidance problem in Chapter 7. A chance-constrained covariance steering problem was defined, subject to spatially-varying uncertainty in the aerothermodynamics, aerodynamic, and atmospheric density. The GPR framework developed in Chapter 4 was leveraged to model the uncertainty in the aerodynamics and aerothermodynamics, and a Gaussian-process based approximation of the atmospheric density was used to model the uncertainty in the planetary atmosphere. A state-feedback control law was optimized to simultaneously compute both feedforward and feedback control terms to

guide an uncertain atmospheric entry distribution to the target orbit, while probabilistically satisfying aerothermodynamic constraints to a prescribed tolerance. In doing so, this chapter explicitly coupled trajectory dynamics with aerothermal uncertainty in a robust and efficient manner.

8.2 Contributions

Within the context of studying the interplay between hypersonic aerothermodynamics and trajectory/guidance design for aerocapture, the major contributions of this research may be summarized as follows:

- **Aerothermodynamic uncertainty quantification and sensitivity analysis for Ice Giant aerocapture**: A modern and comprehensive assessment of the aerothermodynamic uncertainty in Ice Giant aerocapture trajectories was performed, supporting the current interest in Ice Giant exploration.
- **Development of a reduced-order modeling framework using Gaussian-process regression for aerothermodynamic uncertainty propagation**: A probabilistic approach was employed to develop surrogate models that reliably captured aerothermodynamic uncertainty, with physical constraints incorporated to enhance model accuracy and decrease the required training data volume.
- **Predictor-corrector aerocapture guidance algorithm using convex programming**: Principles of SCP were used to design an NPC guidance algorithm that easily incorporated system-level constraints and eliminated the need for discrete lateral logic. Its capabilities were compared to the current state of the art, highlighting its advantages and disadvantages.
- **Robust apoapsis targeting for NPC guidance**: An energy-based expression for targeting an apoapsis radius in aerocapture guidance algorithms was formulated. Because energy is well-defined for any orbit, it substantially improves the convergence properties of NPC guidance algorithms, especially for highly elliptical target orbits. It is compatible with any NPC guidance, not just CPAG.

- **Stochastic aerocapture guidance**: A stochastic optimization problem that considers aerodynamic, atmospheric, and aerothermodynamic uncertainty was used to solve for a feedback control policy that could be implemented onboard a spacecraft. Explicit treatment of the aerothermodynamic uncertainty within the guidance law provides a more comprehensive approach to the design and analysis of aerocapture guidance and aerothermodynamics.

- **Development of a generalized trajectory-optimization program**: A general purpose trajectory simulation and optimization program for broader use in the NGPDL hypersonics research group was designed. A particle swarm optimization algorithm, combined with a modular optimization-problem definition format, allows for rapid exploration of the hypersonic trajectory space [231, 232].

- **Tutorials on convex optimization for trajectory optimization**: Solving a non-convex optimal control problem using SCP is challenging, both numerically and in implementation. In an attempt to help others avoid many of the headaches associated with using SCP, this dissertation produced several introductory tutorials for solving both deterministic and stochastic trajectory-optimization problems. They have been made open source and are available on the author's GitHub.

8.3 Recommendations for Future Work

The research carried out in this dissertation addressed the interplay of hypersonic aerothermodynamics and trajectories at a fundamental level, and there are still many new interesting directions to explore. Some of these are listed below.

Aerothermodynamic Uncertainty Quantification for Ice Giant Aerocapture

- Inclusion of turbulence: The UG/GSA study performed in Chapter 3 assumed laminar flow. For MSL-scale entry capsules, the flow will likely be turbulent throughout much of the heat pulse, so it will be necessary to quantify the level of heating and uncertainty augmentation

introduced by turbulent physics. West and Johnston [42] provide a starting point for key parameters in turbulent flow, but an adaptation to $H_2/He/CH_4$ mixtures would be necessary.

- Coupled CFD-Material-Response-Radiation UQ: As noted in Chapter 3, the radiative heating in the earlier portions of the Ice Giant aerocapture is a similar order of magnitude as the convective heating. As a result, the exclusion of the radiative source term in LeMANS is likely incorrect. The presence of the additional source term in the energy balance will produce a net cooling of the flow, which could have substantial effects on both the convective- and radiative-heating environments. Furthermore, because the UQ study in Chapter 3 is along a trajectory, it would be more representative to include the transient material response in the analysis. Carter and Boyd [47] and Erb et al. [117] provide an excellent review of carbon-based reactions and their uncertainties for ablative materials used in an entry-capsule TPS. The work of Erb et al. [117] would be more representative of the work done in this dissertation, as they investigated a Jovian entry mission. This approach has further implications beyond the uncertainty in ablation reactions, as an ablating TPS serves to cool the surface of the vehicle. Generally, a cooler surface temperature produces higher convective heat fluxes, which would affect the feasible trajectory space.

- Investigation of higher-fidelity physics: For the high velocities in Ice Giant aerocapture, it may be necessary to consider more advanced physics within the computational models. The first area for investigation would be the effects of ionization potential lowering, as a substantial sensitivity to the hydrogen ionization rate coefficients was observed. As the electron number density increases, the ionization potential for H decreases, as originally bound orbits past a certain principal quantum number can no longer exist (i.e. the electron becomes unbound). This reduces not only the number of electronic levels treated for evaluating the thermodynamic properties of H, but also the heat of formation for H^+ [233]. Erb et al. [117] found the effects of ionization potential lowering to be substantial for Jovian entry, however, the conclusions of such an extreme environment cannot necessarily be generalized

to all H_2/He atmospheres. The next key area to investigate would be the usage of the full Stefan-Maxwell multi-component diffusion model. While there is a large difference in the molecular weights between hydrogen/helium and carbon-based species, an effective binary diffusion coefficient was used in Chapter 3, as the amount of carbon present in the boundary layer is small. This assumption would break down if ablation products are modeled, because a larger concentration of heavier hydrocarbons would be injected into the boundary layer.

- Rarefied conditions: The goal of this dissertation was to take steps toward a more coupled treatment of the aerothermodynamics and trajectory of a hypersonic vehicle, yet only select locations along a trajectory were analyzed in detail. To further characteristic the coupling between the trajectory/guidance design and aerothermodynamic uncertainty, it would be necessary to extend the UQ/GSA analysis to higher-altitude, rarefied regimes. In Martian atmospheres, the motion of the sonic line around the shoulder of an MSL-style capsule can produce aerodynamic instabilities [234, 235], so it would be interesting to quantify the extent of this effect, if anything, for Ice Giant atmospheres. This would require the use of a direct simulation Monte Carlo (DSMC) solver, which is not a trivial undertaking.

Usage of Gaussian-Processes for Aerothermodynamics

- Bayesian inference of aerothermal quantities of interest Bayes' theorem is stated as

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \tag{8.1}$$

where the $P(A|B)$ is the probability of observing event "A" given that event "B" has occurred, $P(B|A)$ is the likelihood of observing event "B" (i.e. the conditional, joint probability), $P(A)$ is the prior probability density of event A, and $P(B)$ is the prior probability density of event B. A benefit of GPR is that the surrogate predicts a conditional distribution, meaning $P(B|A)$ is obtained for free. In the context of aerothermodynamic modeling, the event "B" could be the heat flux, and event "A" represents the CFD inputs and GP hyperparameters. Chapter 3 identified several parameters to which the convective and radiative heating were

sensitive, and the GPR framework developed in Chapter 4 could be used to infer these parameters using surface measurements of the heating environment. This approach could be extended to other QoIs, such as trim conditions [137] and freestream compositions, provided there is a strong enough correlation.

- Explicit constraints on GPR posterior distribution: Chapter 4 presented a technique for encoding aerodynamic symmetry into the GP kernel, which drastically improved the three-dimensional aerodynamic surrogate model's out-of-distribution predictive capabilities. Other constraints of interest would be inequality constraints on the GP and its derivatives, specifically boundedness and monotonicity. Agrell [236] presents a method for enforcing these types of constraints in a GPR framework, although it requires knowledge of the kernel's derivatives which makes it challenging to apply to the multi-output and multi-fidelity models developed in Sections 4.4.1 and 4.4.3, respectively.

Aerocapture Guidance

- Direct-force control: The guidance algorithms developed in this dissertation used bank-angle modulation to control the spacecraft. As highlighted in Chapter 6, using the bank angle to steer inherently couples the in- and out-of-plane motion, making the balancing of apoapsis and inclination accuracy difficult. An alternative to bank-angle modulation is direct-force control (DFC) [169, 209, 237]. DFC independently controls the angle of attack and sideslip angle, which decouples the in- and out-of-plane guidance channels. A DFC version of CPAG would be particularly interesting, as the convex-optimization foundation would allow for complete angle-of-attack and side-slip angle profiles to be optimized simultaneously.

- More detailed aerothermal constraints: In Chapters 6 and 7, the aerothermodynamic constraints were imposed on the integrated heat load. The heat load is correlated to the thickness of the TPS, as a thicker TPS is required to prevent the bond-line temperature limit from being exceeded in harsher aerothermal environments. Accounting for a bond-line temperature constraint in a trajectory optimization problem, while more realistic, introduces

additional computational complexity because it requires computing the transient material response of the heat shield along the entire trajectory. However, the GPR framework in Chapter 4 could be used to rapidly propagate the surface-heating environment along the trajectory, allowing a computationally efficient, one-dimensional heat-conduction code to be included within the optimization framework.

- Stochastic density model within NPC guidance: The two guidance algorithms developed in this dissertation solved either a deterministic or a stochastic optimal control problem. Each approach has their benefits, so it would be interesting to pursue a hybrid deterministic-stochastic guidance approach. One route for this could be including the stochastic density model used in Chapter 7 within the prediction phase of an NPC guidance like CPAG or FNPAG. Typically, NPC guidance algorithms use a nominal density profile in the prediction phase that is scaled up or down using sensed acceleration data. An alternative to this would be to use an unscented transform (UT) to propagate multiple trajectories using several realizations of the stochastic density model [238]. Instead of predicting a single, deterministic final state, the UT would enable the NPC to construct a terminal state distribution, potentially improving its robustness to density variability. Because each sample used to construct the UT is independent, this is a process that could be parallelized.

- Out-of-plane targets in stochastic guidance: The only target for the guidance in Chapter 7 was an apoapsis radius, and it would be interesting to extend the stochastic algorithm to additionally consider an inclination target.

- Auto-Tuned Weights: Both algorithms developed in this dissertation required the cost-function and constraint penalty weights to be painstakingly tuned by hand. Recent work from Mceowen et al. [239] has developed techniques for auto-tuning these weights, which decreases the amount of effort required to tune the weights for a given problem. This could potentially improve the robustness and generalization of both CPAG and the stochastic aerocapture guidance algorithm.

Bibliography

[1] Kevin Bowcutt. A Perspective on the Future of Aerospace Vehicle Design. In 12th AIAA International Space Planes and Hypersonic Systems and Technologies, Norfolk, Virginia, December 2003. American Institute of Aeronautics and Astronautics.

[2] Roberto Armellin and Michele Lavagna. Multidisciplinary optimization of aerocapture maneuvers. Journal of Artificial Evolution and Applications, 2008, December 2008.

[3] Howard S. London. Change of Satellite Orbit Plane by Aerodynamic Maneuvering. Journal of the Aerospace Sciences, 29(3):323–332, 1962.

[4] E. M. Repic, M. G. Boobar, and F. G. Chapel. Aerobraking as a potential planetary capture mode. Journal of Spacecraft and Rockets, 5(8):921–926, 1968.

[5] Gerald D. Walberg. A survey of aeroassisted orbit transfer. Journal of Spacecraft and Rockets, 22(1):3–18, 1985.

[6] G. D. Walberg, P. M. Siemers, III, R. L. Calloway, and J. J. Jones. The Aeroassist Flight Experiment. Brighton International Astronautical Federation Congress, October 1987.

[7] Jim Jones. The rationale for an aeroassist flight experiment. In 22nd Thermophysics Conference, Honolulu,HI,U.S.A., June 1987. American Institute of Aeronautics and Astronautics.

[8] Richard Powell. Numerical roll reversal predictor corrector aerocapture and precision landing guidance algorithms for the Mars Surveyor Program 2001 missions. In 23rd Atmospheric Flight Mechanics Conference, Boston,MA,U.S.A., August 1998. American Institute of Aeronautics and Astronautics.

[9] Christian Cazaux, Firouz Naderi, Charles Whetsel, Dave Beaty, Bob Gershman, Richard Kornfeld, Bob Mitcheltree, and Bob Sackheim. The NASA/CNES Mars sample return— a status report. Acta Astronautica, 54(8):601–617, April 2004.

[10] Henry S. Wright, David Y. Oh, Carlos H. Westhelle, Jody L. Fisher, R. Eric Dyke, Karl T. Edquist, James L. Brown, Hilary L. Justh, and Michelle M. Munk. Mars Aerocapture Systems Study. Technical Report L-19290, November 2006. NTRS Author Affiliations: NASA Langley Research Center, Jet Propulsion Lab., California Inst. of Tech., NASA Johnson Space Center, National Inst. of Aerospace, Swales Aerospace, NASA Ames Research Center, NASA Marshall Space Flight Center NTRS Document ID: 20060056070 NTRS Research Center: Langley Research Center (LaRC).

[11] Mary Kae Lockwood, Brett R. Starr, John W. Paulson, Dean A. Kontinos, Y. K. Chen, Bernard Laub, Joseph Olejniczak, Michael J. Wright, Naruhisa Takashima, and Carl G. Justus. Systems Analysis for a Venus Aerocapture Mission. Technical Report NASA/TM-2006-214291, March 2006. NTRS Author Affiliations: NASA Langley Research Center, Science Applications International Corp., NASA Ames Research Center, Analytical Mechanics Associates, Inc., Computer Sciences Corp. NTRS Document ID: 20060010899 NTRS Research Center: Langley Research Center (LaRC).

[12] Mary K. Lockwood, Eric M. Queen, David W. Way, Richard W. Powell, Karl Edquist, Brett W. Starr, Brian R. Hollis, E. Vincent Zoby, Glenn A. Hrinda, and Robert W. Bailey. Aerocapture Systems Analysis for a Titan Mission. Technical Report L-19194, February 2006. NTRS Author Affiliations: NASA Langley Research Center, Jet Propulsion Lab., California Inst. of Tech. NTRS Document ID: 20060007561 NTRS Research Center: Langley Research Center (LaRC).

[13] Mary Kae Lockwood. Neptune Aerocapture Systems Analysis. NTRS Author Affiliations: NASA Langley Research Center NTRS Meeting Information: AIAA Atmospheric Flight Mechanics Conference and Exhibit; 2004-08-16 to 2004-08-19; undefined NTRS Report/Patent Number: AIAA Paper 2004-4951 NTRS Document ID: 20040111217 NTRS Research Center: Langley Research Center (LaRC), January 2004.

[14] National Research Council. Vision and Voyages for Planetary Science in the Decade 2013-2022. The National Academies Press, Washington, DC, 2011.

[15] National Academies of Sciences, Engineering, and Medicine. Origins, Worlds, and Life: A Decadal Strategy for Planetary Science and Astrobiology 2023-2032. The National Academies Press, Washington, DC, 2023.

[16] A Simon, F Nimmo, and R Anderson. Uranus orbiter and probe: Journey to an ice giants system. Technical report, National Aeronautrics and Space Administration, 2021.

[17] Thomas R. Spilker, Mark Adler, Nitin Arora, Patricia M. Beauchamp, James A. Cutts, Michelle M. Munk, Richard W. Powell, Robert D. Braun, and Paul F. Wercinski. Qualitative Assessment of Aerocapture and Applications to Future Missions. Journal of Spacecraft and Rockets, 56(2):536–545, 2019.

[18] Soumyo Dutta, Gonçalo Afonso, Samuel W. Albert, Hisham K. Ali, Gary A. Allen, Antonella I. Alunni, James O. Arnold, Alexander Austin, Gilles Bailet, Shyam Bhaskaran, Alan M. Cassell, George T. Chen, Ian J. Cohen, James A. Cutts, Rohan G. Deshmukh, Robert A. Dillman, Guillermo Dominguez Calabuig, Sarah N. D'Souza, Donald T. Ellerby, Giusy Falcone, Alberto Fedele, Jay Feldman, Roberto Gardi, Athul P. Girija, Tiago Hormigo, Jeffrey P. Hill, Shayna Hume, Christopher Jelloian, Vandana Jha, Breanna J. Johnson, Craig A. Kluever, Jean-Pierre Lebreton, Marcus A. Lobbia, Ping Lu, Ye Lu, Rafael A. Lugo, Daniel A. Matz, Robert W. Moses, Michelle M. Munk, Adam P. Nelessen, Miguel Perez-Ayucar, Richard W. Powell, Zachary R. Putnam, Jeremy R. Rea, Sachin Alexander Reddy, Thomas Reimer, Sarag J. Saikia, Isil Sakraker Özmen, Kunio Sayanagi, Stephan Schuster, Jennifer Scully, Ronald R. Sostaric, Christophe Sotin, David A. Spencer, Benjamin M. Tackett, Nikolas Trawny, Ethiraj Venkatapathy, Paul F. Wercinski, and Cindy L. Young. Aerocapture as an Enhancing Option for Ice Giants Missions. Technical Report 20205002647, NASA, July 2020.

[19] Sarag J. Saikia, James Millane, Ye Lu, Alec Mudek, Archit Arora, Paul Witsberge, Kyle Hughes, James M. Longuski, Thomas Spilker, Anastassios Petropoulos, Nitin Arora, James Cutts, John Elliott, Jon Sims, and Kim Reh. Aerocapture Assessment for NASA Ice Giants Pre-Decadal Survey Mission Study. Journal of Spacecraft and Rockets, 58(2):505–515, March 2021.

[20] Athul Pradeepkumar Girija, Sarag J. Saikia, James M. Longuski, Ye Lu, and James A. Cutts. Quantitative Assessment of Aerocapture and Applications to Future Solar System Exploration. Journal of Spacecraft and Rockets, 59(4):1074–1095, 2022.

[21] Athul Pradeepkumar Girija. A Flagship-class Uranus Orbiter and Probe mission concept using aerocapture. Acta Astronautica, 202:104–118, January 2023.

[22] Soumyo Dutta, Eli Shellabarger, James B. Scoggins, Andrew Gomez-Delrio, Rafael Lugo, Rohan Deshmukh, Benjamin Tackett, Joseph Williams, Breanna Johnson, Daniel Matz, Joshua Geiser, Jonathan Morgan, Ricardo Restrepo, and Declan Mages. Uranus Flagship-class Orbiter and Probe Using Aerocapture. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[23] Soumyo Dutta, Rohan Deshmukh, James B. Scoggins, Eli Shellabarger, Andrew J. Gomez-Delrio, Rafael Lugo, Breanna J. Johnson, Daniel Matz, Sergio A. Sandoval, Joshua Geiser, Jonathan Morgan, Ricardo L. Restrepo, Declan Mages, Pardha Chadalavada, Joseph Williams, and Justin Garland. Aerocapture Solutions for Uranus Flagship-Class Orbiter and Probe. In AIAA SCITECH 2025 Forum. American Institute of Aeronautics and Astronautics, 2025.

[24] Chul Park. Nonequilibrium hypersonic aerothermodynamics. John Wiley and Sons, 1990.

[25] Michael J Wright. Aerothermal Modeling for Entry and Aerocapture. 2007.

[26] Lewis P. Leibowitz. Measurements of the structure of an ionizing shock wave in a hydrogen-helium mixture. The Physics of Fluids, 16(1):59–68, January 1973.

[27] Floyd R. Livingston and P. T. Y. Poon. Relaxation Distance and Equilibrium Electron Density Measurements in Hydrogen-Helium Plasmas. AIAA Journal, 14(9):1335–1337, 1976.

[28] D.W. Bogdanoff and C. Park. Radiative interaction between driver and driven gases in an arc-driven shock tube. Shock Waves, 12(3):205–214, November 2002.

[29] Chul Park. Nonequilibrium Ionization and Radiation in Hydrogen-Helium Mixtures. Journal of Thermophysics and Heat Transfer, 26(2):231–243, April 2012.

[30] Lewis P. Leibowitz and Ta-Jin Kuo. Ionizational Nonequilibrium Heating During Outer Planetary Entries. AIAA Journal, 14(9):1324–1329, September 1976.

[31] Chul Park. Effect of Lymann Radiation on Nonequilibrium Ionization of Atomic Hydrogen. In 37th AIAA Thermophysics Conference, Portland, Oregon, June 2004. American Institute of Aeronautics and Astronautics.

[32] Chul Park. Nonequilibrium Chemistry and Radiation for Neptune Entry. Journal of Spacecraft and Rockets, 48(6):897–903, November 2011.

[33] Chul Park. Viscous Shock-Layer Calculation of Stagnation-Region Heating Environment in Neptune Aerocapture. Journal of Spacecraft and Rockets, 51(2):635–640, March 2014.

[34] Brian Hollis, Michael Wright, Joseph Olejniczak, Naruhisa Takashima, Kenneth Sutton, and Dinesh Prabhu. Preliminary Convective-Radiative Heating Environments for a Neptune Aerocapture Mission. In AIAA Atmospheric Flight Mechanics Conference and Exhibit, Guidance, Navigation, and Control and Co-located Conferences. American Institute of Aeronautics and Astronautics, August 2004.

[35] João Coelho and Mário Lino da Silva. Aerothermodynamic analysis of Neptune ballistic entry and aerocapture flows. Advances in Space Research, 71(8):3408–3432, April 2023.

[36] Joseph Steer, Peter L Collen, Alex B Glenn, Christopher Hambidge, Luke J Doherty, Matthew McGilvray, Stefan Loehle, and Louis Walpot. Shock radiation tests for ice giant entry probes including ch4 in the t6 free-piston driven wind tunnel. In AIAA SCITECH 2023 Forum, page 1729, 2023.

[37] Brett A. Cruden and Augustin C. Tibère-Inglesse. Impact of Trace CH_4 on Shock Layer Radiation in Outer Planet Entry. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[38] Kaelan Hansson, Alex T. Carroll, Iain D. Boyd, and Savio J. Poovathingal. Analysis of chemical kinetic parameters for hydrogen atmospheres. AIAA SciTech Forum, January 2021.

[39] Alex T Carroll, Guillaume Blanquart, Aaron M Brandis, and Brett A Cruden. Kinetic and transport modeling for entry flows in hydrogen-helium atmospheres. In AIAA AVIATION 2023 Forum, page 3728, 2023.

[40] Alex T. Carroll, Guillaume Blanquart, Aaron M. Brandis, and Brett A. Cruden. State-Specific Kinetic Modeling for Predictions of Radiative Heating in H2/He Entry Flows. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[41] Thomas K West IV, Christopher O Johnston, and Serhat Hosder. Uncertainty and sensitivity analysis of afterbody radiative heating predictions for earth entry. Journal of Thermophysics and Heat Transfer, 31(2):294–306, 2017.

[42] Thomas K. West and Christopher O. Johnston. Efficient Parametric Uncertainty Analysis of an Earth Entry Vehicle Concept using Least Angle Regression. In AIAA SCITECH 2023 Forum, National Harbor, MD & Online, January 2023. American Institute of Aeronautics and Astronautics.

[43] Aaron Brandis, Justin Haskins, Tom West, and Monica Hughes. Entry systems modeling: Progress in aerothermal predictions for flight and new research directions. In 21st International Planetary Probe Workshop (IPPW), 2024.

[44] Timothy T. Aiken and Iain D. Boyd. Analysis of Critical Rate Processes for Ionization in Shock-Heated Air. In AIAA AVIATION 2023 Forum. American Institute of Aeronautics and Astronautics, 2023.

[45] Jens A. Rataczak, Iain D. Boyd, and Jay W. McMahon. Parametric Sensitivity Analysis and Uncertainty Quantification of Ice Giant Aerocapture Aerothermodynamics. In AIAA AVIATION FORUM AND ASCEND, 2024.

[46] Timothy T. Aiken and Iain. D. Boyd. Sensitivity Analysis of Ionization in Two-Temperature Models of Hypersonic Air Flows. Journal of Thermophysics and Heat Transfer, 0(0):1–13, 2024.

[47] Nicholas Carter and Iain D. Boyd. Sensitivity Analyses of Thermochemical Kinetics for Radiating Ablative Nonequilibrium Hypersonic Flows. In AIAA AVIATION FORUM AND ASCEND 2024. American Institute of Aeronautics and Astronautics, 2024.

[48] Theodore Ro and Eric Queen. Study of Martian aerocapture terminal point guidance. In 23rd Atmospheric Flight Mechanics Conference. American Institute of Aeronautics and Astronautics, 1998.

[49] Rafael A. Lugo, Richard Powell, and Alicia M. Dwyer-Cianciolo. Overview of a Generalized Numerical Predictor-Corrector Targeting Guidance with Application to Human-Scale Mars Entry, Descent, and Landing. In AIAA Scitech 2020 Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2020.

[50] Nguyen X Vinh, Adolf Busemann, and Robert D Culp. Hypersonic and planetary entry flight mechanics. NASA Sti/Recon Technical Report A, 81:16245, 1980.

[51] H. Julian Allen and A. J. Eggers. A study of the motion and aerodynamic heating of ballistic missiles entering the earth's atmosphere at high supersonic speeds, January 1958.

[52] Ping Lu. Entry Guidance: A Unified Method. Journal of Guidance, Control, and Dynamics, 37(3):713–728, 2014.

[53] Jeremy Rea, Luke McNamara, and Mark Kane. ORION ARTEMIS I ENTRY PERFORMANCE. In 46th Annual American Astronautical Society Guidance, Navigation and Control Conference, 2024.

[54] Jarret M Lafleur. The Conditional Equivalence of ΔV Minimization and Apoapsis Targeting in Numerical Predictor- Corrector Aerocapture Guidance. NASA TM-2011-216156, Johnson Space Center, Houston, TX, August 2011.

[55] Ping Lu, Christopher J. Cerimele, Michael A. Tigges, and Daniel A. Matz. Optimal Aerocapture Guidance. Journal of Guidance, Control, and Dynamics, 38(4):553–565, 2015.

[56] Kyle D Webb, Ping Lu, and Alicia M Dwyer Cianciolo. Aerocapture guidance for a human mars mission. In AIAA Guidance, Navigation, and Control Conference, page 1900, 2017.

[57] Daniel A Matz, Ping Lu, Gavin F Mendeck, and Ronald R Sostaric. Application of a fully numerical guidance to mars aerocapture. In AIAA Guidance, Navigation, and Control Conference, page 1901, 2017.

[58] C R Heidrich, S Dutta, and R D Braun. Modern aerocapture guidance to enable reduced-lift vehicles at neptune. In AAS/AIAA Space Flight Mechanics Conference, number AAS 19-221, 2019.

[59] Daniel Louis Engel, Soumyo Dutta, and Zachary Putnam. Numerical predictor-corrector based guidance scheme for aero-gravity assist at titan for enceladus missions. In AAS Space Flight Mechanics Meeting, 2021.

[60] Behcet Acikmese, M Aung, Jordi Casoliva, Swati Mohan, Andrew Johnson, Daniel Scharf, David Masten, Joel Scotkin, Aron Wolf, and Martin W Regehr. Flight testing of trajectories computed by g-fold: Fuel optimal large divert guidance algorithm for planetary landing. In AAS/AIAA spaceflight mechanics meeting, page 386, 2013.

[61] Behcet Acikmese and Scott R. Ploen. Convex Programming Approach to Powered Descent Guidance for Mars Landing. Journal of Guidance, Control, and Dynamics, 30(5):1353–1366, 2007.

[62] Lars Blackmore, Behçet Açikmeşe, and Daniel P. Scharf. Minimum-Landing-Error Powered-Descent Guidance for Mars Landing Using Convex Optimization. Journal of Guidance, Control, and Dynamics, 33(4):1161–1171, 2010.

[63] Marco Sagliano. Pseudospectral Convex Optimization for Powered Descent and Landing. Journal of Guidance, Control, and Dynamics, 41(2):320–334, February 2018.

[64] Xinfu Liu, Zuojun Shen, and Ping Lu. Entry Trajectory Optimization by Second-Order Cone Programming. Journal of Guidance, Control, and Dynamics, 39(2):227–241, 2016.

[65] Zhenbo Wang and Michael J. Grant. Hypersonic Trajectory Optimization by Sequential Semidefinite Programming. In AIAA Atmospheric Flight Mechanics Conference, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2017.

[66] Zhenbo Wang and Ye Lu. Improved Sequential Convex Programming Algorithms for Entry Trajectory Optimization. Journal of Spacecraft and Rockets, 57(6):1373–1386, November 2020.

[67] Christian Hofmann and Francesco Topputo. Rapid Low-Thrust Trajectory Optimization in Deep Space Based on Convex Programming. Journal of Guidance, Control, and Dynamics, 44(7):1379–1388, July 2021.

[68] Serena Jing Zhang, Behçet Açikmeşe, Sean Shan-Min Swei, and Dinesh Prabhu. Convex programming approach to real-time trajectory optimization for mars aerocapture. In 2015 IEEE Aerospace Conference, pages 1–7. IEEE, 2015.

[69] Hongwei Han, Dong Qiao, Hongbo Chen, and Xiangyu Li. Rapid planning for aerocapture trajectory via convex optimization. Aerospace Science and Technology, 84:763–775, January 2019.

[70] Jens A. Rataczak, Iain D. Boyd, and Jay W. McMahon. Stochastic Aerocapture Trajectory Optimization Under Aerothermodynamic Uncertainty. In AIAA SCITECH 2025 Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2025.

[71] Jens A. Rataczak, Jay W. McMahon, and Iain D. Boyd. Predictor-Corrector Aerocapture Guidance Using Convex Programming. In AIAA SCITECH 2025 Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2025.

[72] Kenshiro Oguri and Jay W. McMahon. Robust Spacecraft Guidance Around Small Bodies Under Uncertainty: Stochastic Optimal Control Approach. Journal of Guidance, Control, and Dynamics, 44(7):1295–1313, 2021.

[73] Divija Aleti, Kenshiro Oguri, and Naoya Kumagai. Chance-Constrained Output-Feedback Control without History Feedback: Application to NRHO Stationkeeping. In AAS/AIAA Astrodynamics Specialist Conference, Big Sky, MT. August 2023.

[74] Jack Ridderhof and Panagiotis Tsiotras. Chance-Constrained Covariance Steering in a Gaussian Random Field via Successive Convex Programming. Journal of Guidance, Control, and Dynamics, 45(4):599–610, April 2022.

[75] Iain D. Boyd and Thomas E. Schwartzentruber. Nonequilibrium Gas Dynamics and Molecular Simulation. Cambridge Aerospace Series. Cambridge University Press, 2017.

[76] Christopher O. Johnston, Brian R. Hollis, and Kenneth Sutton. Non-Boltzmann Modeling for Air Shock-Layer Radiation at Lunar-Return Conditions. Journal of Spacecraft and Rockets, 45(5):879–890, September 2008.

[77] Chul Park. Assessment of two-temperature kinetic model for ionizing air. Journal of Thermophysics and Heat Transfer, 3(3):233–244, 1989.

[78] Chul Park. The Limits of Two-Temperature Kinetic Model in Air. In 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, Aerospace Sciences Meetings. American Institute of Aeronautics and Astronautics, January 2010.

[79] Leonardo C. Scalabrin. Numerical simulation of weakly ionized hypersonic flow over re-entry capsules. Ph.D. Thesis, University of Michigan-Ann Arbor, 2007.

[80] Alexandre Martin, Leonardo C. Scalabrin, and Iain D. Boyd. High performance modeling of atmospheric re-entry vehicles. Journal of Physics: Conference Series, 341(1):012002, February 2012.

[81] W. G. Vincenti and C. H. Kruger. Introduction to Physical Gas Dynamics. Krieger Publishing Company, 1982.

[82] Roger C. Millikan and Donald R. White. Systematics of Vibrational Relaxation. The Journal of Chemical Physics, 39(12):3209–3213, June 1963.

[83] Hicham Alkandry, Iain D. Boyd, and Alexandre Martin. Comparison of Transport Properties Models for Flowfield Simulations of Ablative Heat Shields. Journal of Thermophysics and Heat Transfer, 28(4):569–582, 2014.

[84] C. R. Wilke. A Viscosity Equation for Gas Mixtures. The Journal of Chemical Physics, 18(4):517–519, April 1950.

[85] Roop N. Gupta, Jerrold M. Yos, Richard A. Thompson, and Kam-Pui Lee. A review of reaction rates and thermodynamic and transport properties for an 11-species air model for chemical and thermal nonequilibrium calculations to 30000 K, August 1990.

[86] Grant Palmer, Dinesh Prabhu, and Brett A. Cruden. Aeroheating Uncertainties in Uranus and Saturn Entries by the Monte Carlo Method. Journal of Spacecraft and Rockets, 51(3):801–814, May 2014.

[87] Chul Park, Richard L. Jaffe, and Harry Partridge. Chemical-Kinetic Parameters of Hyperbolic Earth Entry. Journal of Thermophysics and Heat Transfer, 15(1):76–90, 2001.

[88] Michael J. Wright, Deepak Bose, Grant E. Palmer, and Eugene Levin. Recommended Collision Integrals for Transport Property Computations Part 1: Air Species. AIAA Journal, 43(12):2558–2564, December 2005.

[89] A. M. Brandis, D. A. Saunders, C. O. Johnston, B. A. Cruden, and T. R. White. Radiative Heating on the After-Body of Martian Entry Vehicles. Journal of Thermophysics and Heat Transfer, 34(1):66–77, 2020.

[90] Kaelan B. Hansson. Electromagnetic energy transport through complex media. Ph.D. Thesis, University of Michigan-Ann Arbor, 2021.

[91] William Oberkampf, Jon Helton, and Kari Sentz. Mathematical representation of uncertainty. In 19th AIAA Applied Aerodynamics Conference, Fluid Dynamics and Co-located Conferences. American Institute of Aeronautics and Astronautics, June 2001.

[92] Thomas K. West, Serhat Hosder, and Christopher O. Johnston. Multistep Uncertainty Quantification Approach Applied to Hypersonic Reentry Flows. Journal of Spacecraft and Rockets, 51(1):296–310, 2014.

[93] Thomas K. West and Serhat Hosder. Uncertainty Quantification of Hypersonic Reentry Flows with Sparse Sampling and Stochastic Expansions. Journal of Spacecraft and Rockets, 52(1):120–133, 2015.

[94] Thomas K. West, Andrew J. Brune, Serhat Hosder, and Christopher O. Johnston. Uncertainty Analysis of Radiative Heating Predictions for Titan Entry. Journal of Thermophysics and Heat Transfer, 30(2):438–451, April 2016.

[95] Thomas K. West, Christopher O. Johnston, and Serhat Hosder. Uncertainty and Sensitivity Analysis of Afterbody Radiative Heating Predictions for Earth Entry. Journal of Thermophysics and Heat Transfer, 31(2):294–306, 2017.

[96] Timothy T. Aiken and Iain D. Boyd. Collisional-radiative modeling of shock-heated nitrogen mixtures. Journal of Applied Physics, 135(9):093301, March 2024.

[97] Dongbin Xiu and George Em Karniadakis. The Wiener–Askey Polynomial Chaos for Stochastic Differential Equations. SIAM Journal on Scientific Computing, 24(2):619–644, January 2002.

[98] Ji Peng, Jerrad Hampton, and Alireza Doostan. On polynomial chaos expansion via gradient-enhanced l1-minimization. Journal of Computational Physics, 310:440–458, April 2016.

[99] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. The Annals of Statistics, 32(2):407–499, April 2004.

[100] Stefano Marelli and Bruno Sudret. UQLab: A Framework for Uncertainty Quantification in Matlab. pages 2554–2563, July 2014. Publisher: American Society of Civil Engineers.

[101] Carl Edward Rasmussen and Christopher KI Williams. Gaussian processes for machine learning, volume 1. Springer, 2006.

[102] David Duvenaud. Automatic model construction with Gaussian processes. PhD thesis, University of Cambridge, 2014.

[103] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.

[104] Alexander G. de G. Matthews, Mark van der Wilk, Tom Nickson, Keisuke. Fujii, Alexis Boukouvalas, Pablo León-Villagrá, Zoubin Ghahramani, and James Hensman. GPflow: A Gaussian process library using TensorFlow. Journal of Machine Learning Research, 18(40):1–6, apr 2017.

[105] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: a system for large-scale machine learning. In Osdi, volume 16, pages 265–283. Savannah, GA, USA, 2016.

[106] Iain D. Boyd. Monte Carlo simulation of nonequilibrium flow in a low-power hydrogen arcjet. Physics of Fluids, 9(10):3086–3095, October 1997.

[107] Kaelan Hansson, Alex T. Carroll, Savio J. Poovathingal, and Iain D. Boyd. Analysis of Chemical Kinetic Parameters for Hydrogen Atmospheres. In AIAA Scitech 2021 Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2021.

[108] G. Colonna, G. D'Ammando, L. D. Pietanza, and M. Capitelli. Excited-state kinetics and radiation transport in low-temperature plasmas. Plasma Physics and Controlled Fusion, 57(1):014009, November 2014.

[109] Gianpiero Colonna, Lucia D. Pietanza, Giuliano D'Ammando, Roberto Celiberto, Mario Capitelli, and Annarita Laricchiuta. Vibrational kinetics of electronically excited states in H2 discharges. The European Physical Journal D, 71(11):279, November 2017.

[110] Gianpiero Colonna, Annarita Laricchiuta, and Lucia Daniela Pietanza. Modeling plasma heating by ns laser pulse. Spectrochimica Acta Part B: Atomic Spectroscopy, 141:85–93, March 2018.

[111] Gianpiero Colonna, Lucia Daniela Pietanza, and Annarita Laricchiuta. Ionization kinetic model for Hydrogen-Helium atmospheres in hypersonic shock tubes. International Journal of Heat and Mass Transfer, 156:119916, August 2020.

[112] Alex T. Carroll, Guillaume Blanquart, Aaron M. Brandis, and Brett A. Cruden. Kinetic and Transport Modeling for Entry Flows in Hydrogen-Helium Atmospheres. In AIAA AVIATION 2023 Forum. American Institute of Aeronautics and Astronautics, 2023.

[113] Brett A. Cruden and David W. Bogdanoff. Shock Radiation Tests for Saturn and Uranus Entry Probes. Journal of Spacecraft and Rockets, 54(6):1246–1257, 2017.

[114] Joseph Steer, Peter L. Collen, Alex B. Glenn, Christopher Hambidge, Luke J. Doherty, Matthew McGilvray, Stefan Loehle, and Louis Walpot. Shock Radiation Tests for Ice Giant Entry Probes Including CH4 in the T6 Free-Piston Driven Wind Tunnel. In AIAA SCITECH 2023 Forum. American Institute of Aeronautics and Astronautics, 2023.

[115] J. I. Moses, T. Cavalié, L. N. Fletcher, and M. T. Roman. Atmospheric chemistry on Uranus and Neptune. Philosophical transactions. Series A, Mathematical, physical, and engineering sciences, 378(2187):20190477, December 2020.

[116] Ricardo Restrepo, Declan Mages, Matthew Smith, Rohan Deshmukh, Soumyo Dutta, and Lylia Benhacine. Mission Design and Navigation Solutions for Uranus Aerocapture. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[117] Aaron J. Erb, Thomas K. West, and Christopher O. Johnston. Investigation of Galileo Probe Entry Heating with Coupled Radiation and Ablation. Journal of Spacecraft and Rockets, May 2020.

[118] Mário Lino da Silva and João Vargas. A Physically-Consistent Chemical Dataset for the Simulation of N_2-CH_4 Shocked Flows Up to T=100,000K, December 2022. arXiv:2212.09911 [physics].

[119] Aaron M Brandis and Brett A Cruden. Neqair v15. 0 release notes: Nonequilibrium and equilibrium radiative transport and spectra program. Technical report, 2019.

[120] Carl D. Scott, Samir Farhat, Alix Gicquel, Khaled Hassouni, and Michel Lefebvre. Determining electron temperature and density in a hydrogen microwave plasma. Journal of Thermophysics and Heat Transfer, 10(3):426–435, July 1996.

[121] H L Justh and J Hoffman. Neptune Global Reference Atmospheric Model (Neptune-GRAM): User Guide, 2020.

[122] James B. Scoggins, Andrew D. Hinkle, and Eli Shellabarger. Aeroheating Environment of Aerocapture Systems for Uranus Orbiters. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[123] Jonathan Morgan, Joseph Williams, Ethiraj Venkataphy, Matthew Gasch, Rohan Deshmukh, Eli Shellabarger, James B. Scoggins, Andrew J. Gomez-Delrio, Benjamin Tackett, and Soumyo Dutta. Thermal Protection System Design of Aerocapture Systems for Uranus Orbiters. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[124] Eli Shellabarger, James B. Scoggins, Andrew D. Hinkle, Soumyo Dutta, Rohan Deshmukh, Mihirkumar Patel, and Sarah Agam. Aerodynamic Implications of Aerocapture Systems for Uranus Orbiters. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[125] Iain D. Boyd and Eswar Josyula. Analysis of Associative Ionization Rates for Hypersonic Flows. Journal of Thermophysics and Heat Transfer, 35(3):484–493, 2021.

[126] V A Srećković, M S Dimitrijević, and Lj M Ignjatović. Atom–Rydberg atom chemi-ionization/recombination processes in the hydrogen clouds in broad-line region of AGNs. Monthly Notices of the Royal Astronomical Society, 480(4):5078–5083, November 2018.

[127] D. Auerbach, R. Cacak, R. Caudano, T. D. Gaily, C. J. Keyser, J. W. McGowan, J. B. A. Mitchell, and S. F. J. Wilk. Merged electron-ion beam experiments. I. Method and measurements of (e-H2+) and (e-H3+) dissociative-recombination cross sections. Journal of Physics B: Atomic and Molecular Physics, 10(18):3797, December 1977.

[128] Rohan Deshmukh, Pardha Chadalavada, Soumyo Dutta, and Rafael Lugo. 6-DoF Uranus Aerocapture Trajectory Analysis. In AIAA SCITECH 2025 Forum. American Institute of Aeronautics and Astronautics, 2025.

[129] Karen Bibb, Eric Walker, Greg Brauckmann, and Philip Robinson. Development of the orion crew module static aerodynamic database, part ii: Subsonic/supersonic. 29th AIAA Applied Aerodynamics Conference, 2011.

[130] Karen Bibb, Greg Brauckmann, Eric Walker, and Philip Robinson. Development of the orion crew module static aerodynamic database, part i: Hypersonic. 29th AIAA Applied Aerodynamics Conference, 2011.

[131] Julien Laurenceau and P Sagaut. Building efficient response surfaces of aerodynamic functions with kriging and cokriging. AIAA Journal, 46(2):498–507, 2008.

[132] Yuichi Kuya, Kenji Takeda, Xin Zhang, and Alexander IJ Forrester. Multifidelity surrogate modeling of experimental and computational aerodynamic data sets. AIAA journal, 49(2):289–298, 2011.

[133] David JJ Toal and Andy J Keane. Efficient multipoint aerodynamic design optimization via cokriging. Journal of Aircraft, 48(5):1685–1695, 2011.

[134] Meysam Mohammadi-Amin, Moein M. Entezari, and Alireza Alikhani. An efficient surrogate-based framework for aerodynamic database development of manned reentry vehicles. Advances in Space Research, 62(5):997–1014, September 2018.

[135] Thomas J Wignall, Tenavi Nakamura-Zimmerer, James B Scoggins, Steven Snyder, Satvik Kumar, and Brendon K Colbert. Comparisons of performance metrics and machine learning methods on an entry descent and landing database. In AIAA SCITECH 2023 Forum, page 1183, 2023.

[136] Steven Snyder, Thomas J Wignall, Justin S Green, Satvik Kumar, Michael W Lee, Tenavi Nakamura-Zimmerer, James B Scoggins, and Robert A Williams. Aerofusion: Data fusion and uncertainty quantification for entry vehicles. In AIAA SCITECH 2023 Forum, page 1182, 2023.

[137] James B. Scoggins, Thomas J. Wignall, Tenavi Nakamura-Zimmerer, and Karen L. Bibb. Multihierarchy Gaussian Process Models for Probabilistic Aerodynamic Databases using Uncertain Nominal and Off-Nominal Configuration Data. In AIAA SCITECH 2023 Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2023.

[138] Tenavi Nakamura-Zimmerer, Mary T Stringer, Brendon K Colbert, and James B Scoggins. Structured covariance gaussian networks for orion crew module aerodynamic uncertainty quantification. In AIAA SCITECH 2023 Forum, page 1184, 2023.

[139] James B. Scoggins, Andrew Hinkle, Sara Agam, and Mihirkumar Patel. Forebody Aeroheating Environments for a Uranus Aerocapture System Using a Mars-Heritage Entry Vehicle. In AIAA SCITECH 2025 Forum. American Institute of Aeronautics and Astronautics, 2025.

[140] Jens A. Rataczak, Iain D. Boyd, and Jay W. McMahon. Surrogate Models for Hypersonic Aerothermodynamics and Aerodynamics using Gaussian Process Regression. In AIAA SCITECH 2024 Forum, AIAA SciTech Forum, January 2024.

[141] Mauricio A. Alvarez, Lorenzo Rosasco, and Neil D. Lawrence. Kernels for vector-valued functions: a review, 2012.

[142] Loïc Brevault, Mathieu Balesdent, and Ali Hebbal. Overview of gaussian process based multi-fidelity techniques with variable relationship between fidelities, application to aerospace systems. Aerospace Science and Technology, 107:106339, 2020.

[143] Emily Paxton, Arturo Villasenor, Timothy Gallagher, and Peter Chun. Multi fidelity aerodynamic and aerothermal boundary layer and near body characterization of hypersonic configurations. In AIAA AVIATION 2020 FORUM, page 2721, 2020.

[144] Mario J Santos, Serhat Hosder, and Thomas K West. Multifidelity turbulent heating prediction of hypersonic inflatable aerodynamic decelerators with surface scalloping. Journal of Spacecraft and Rockets, 58(5):1325–1338, 2021.

[145] Kevin R Quinlan, Jagadeesh Movva, Elizabeth V Stein, and Ana Kupresanin. Leveraging multi-fidelity aerodynamic databasing to efficiently represent a hypersonic design space. ASCEND 2021, 2021.

[146] Marc C Kennedy and Anthony O'Hagan. Predicting the output from a complex computer code when fast approximations are available. Biometrika, 87(1):1–13, 2000.

[147] Loic Le Gratiet and Josselin Garnier. Recursive co-kriging model for design of computer experiments with multiple levels of fidelity. International Journal for Uncertainty Quantification, 4(5), 2014.

[148] Edwin V Bonilla, Kian Chai, and Christopher Williams. Multi-task gaussian process prediction. Advances in neural information processing systems, 20, 2007.

[149] Phillip Boyle and Marcus Frean. Dependent gaussian processes. Advances in neural information processing systems, 17, 2004.

[150] Andy J Keane. Cokriging for robust design optimization. AIAA journal, 50(11):2351–2364, 2012.

[151] M Giselle Fernández-Godino, Chanyoung Park, Nam-Ho Kim, and Raphael T Haftka. Review of multi-fidelity models. arXiv preprint arXiv:1609.07196, 2016.

[152] KR Quinlan, J Movva, EV Stein, and A Kupresanin. Multi-fidelity aerodynamic databases for efficient representation of hypersonic design spaces. Technical report, Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2021.

[153] Paris Perdikaris, Maziar Raissi, Andreas Damianou, Neil D Lawrence, and George Em Karniadakis. Nonlinear information fusion algorithms for data-efficient multi-fidelity modelling. Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 473(2198):20160751, 2017.

[154] Andrew J. Lofthouse, Iain D. Boyd, and Michael J. Wright. Effects of continuum breakdown on hypersonic aerothermodynamics. Physics of Fluids, 19(2):027105, February 2007.

[155] Mark E Johnson, Leslie M Moore, and Donald Ylvisaker. Minimax and maximin distance designs. Journal of statistical planning and inference, 26(2):131–148, 1990.

[156] Grgoire Montavon, Genevive Orr, and Klaus-Robert Mller. Neural Networks: Tricks of the Trade. Springer Publishing Company, Incorporated, 2nd edition, 2012.

[157] J. A. Fay and F. R. Riddell. Theory of stagnation point heat transfer in dissociated air. Journal of the Aerospace Sciences, 25(2):73–85, 1958.

[158] Kenneth Sutton and Randolph A Graves Jr. A general stagnation-point convective heating equation for arbitrary gas mixtures. Technical report, 1971.

[159] John D. Anderson. Hypersonic and High-Temperature Gas Dynamics, Third Edition. AIAA Education Series, 2019.

[160] David Ginsbourger, Olivier Roustant, and Nicolas Durrande. Invariances of random fields paths, with applications in Gaussian Process Regression, August 2013. arXiv:1308.1359 [math, stat].

[161] David Ginsbourger, Xavier Bay, Olivier Roustant, and Laurent Carraro. Argumentwise invariant kernels for the approximation of invariant functions. Annales de la Faculté des sciences de Toulouse : Mathématiques, 21(3):501–527, October 2012.

[162] A. Miele, Z. G. Zhao, and W. Y. Lee. Optimal trajectories for the Aeroassisted Flight Experiment. Part 1: Equations of motion in an Earth-fixed system. Technical Report NAS 1.26:186134, January 1989.

[163] Jon C. Harpold and Donald E. Gavert. Space Shuttle entry guidance performance results. Journal of Guidance, Control, and Dynamics, 6(6):442–447, 1983.

[164] Gavin Mendeck and Lynn Craig. Entry Guidance for the 2011 Mars Science Laboratory Mission. In AIAA Atmospheric Flight Mechanics Conference, Guidance, Navigation, and Control and Co-located Conferences. American Institute of Aeronautics and Astronautics, August 2011.

[165] Lynn C. McGrew, Brian J. Killeen, and Gavin F. Mendeck. Entry Guidance Design and Post-Flight Performance of the Mars 2020 Mission. Kaua'i, HI, 2025.

[166] Ian M Meginnis, Zachary R Putnam, Ian G Clark, Robert D Braun, and Gregg H Barton. Guided entry performance of low ballistic coefficient vehicles at mars. Journal of Spacecraft and Rockets, 50(5):1047–1059, 2013.

[167] Soumyo Dutta, Robert D Braun, and Christopher D Karlgaard. Uncertainty quantification for mars entry, descent, and landing reconstruction using adaptive filtering. Journal of Spacecraft and Rockets, 51(3):967–977, 2014.

[168] Geethu Lisba Jacob, Geethu Neeler, and RV Ramanan. Mars entry mission bank profile optimization. Journal of Guidance, Control, and Dynamics, 37(4):1305–1316, 2014.

[169] Rohan G. Deshmukh, David A. Spencer, and Soumyo Dutta. Investigation of direct force control for aerocapture at neptune. Acta Astronautica, 175:375–386, October 2020.

[170] Casey R Heidrich, Marcus J Holzinger, and Robert D Braun. Optimal information filtering for robust aerocapture trajectory generation and guidance. Journal of Spacecraft and Rockets, 59(2):524–537, 2022.

[171] Samuel W. Albert, Alireza Doostan, and Hanspeter Schaub. Onboard density modeling for planetary entry via karhunen-loève expansion. In 2023 IEEE Aerospace Conference, pages 1–12, 2023.

[172] Ping Lu, Sergio Sandoval, and Christopher Davami. Fast and robust optimization of full trajectory from entry through powered descent. Journal of Guidance, Control, and Dynamics, 47(2):203–216, 2024.

[173] Ping Lu. What Is Guidance? Journal of Guidance, Control, and Dynamics, 44(7):1237–1238, 2021. Publisher: American Institute of Aeronautics and Astronautics _eprint: https://doi.org/10.2514/1.G006191.

[174] Daniel Matz, Breanna J. Johnson, Joshua Geiser, Sergio Sandoval, Rohan Deshmukh, Rafael Lugo, Soumyo Dutta, and Pardhasai Chadalavada. Analysis of a Bank Control Guidance for Aerocapture at Uranus. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[175] L. S. Pontryagin. Mathematical Theory of Optimal Processes. Routledge, London, May 1987.

[176] Kelly M. Smith. Predictive Lateral Logic for Numerical Entry Guidance Algorithms. Napa, CA, February 2016. NTRS Author Affiliations: NASA Johnson Space Center NTRS Report/Patent Number: JSC-CN-35110-1 NTRS Document ID: 20160001182 NTRS Research Center: Johnson Space Center (JSC).

[177] Stephen P Boyd and Lieven Vandenberghe. Convex optimization. Cambridge university press, 2004.

[178] Jorge Nocedal and Stephen J Wright. Numerical optimization. Springer, 1999.

[179] Xinfu Liu and Ping Lu. Solving Nonconvex Optimal Control Problems by Convex Optimization. Journal of Guidance, Control, and Dynamics, 37(3):750–765, 2014.

[180] Zhenbo Wang. A survey on convex optimization for guidance and control of vehicular systems. Annual Reviews in Control, 57:100957, January 2024.

[181] Behçet Açıkmeşe and Lars Blackmore. Lossless convexification of a class of optimal control problems with non-convex control constraints. Automatica, 47(2):341–347, February 2011.

[182] Danylo Malyuta, Taylor P. Reynolds, Michael Szmuk, Thomas Lew, Riccardo Bonalli, Marco Pavone, and Behçet Açıkmeşe. Convex Optimization for Trajectory Generation: A Tutorial on Generating Dynamically Feasible Trajectories Reliably and Efficiently. IEEE Control Systems Magazine, 42(5):40–113, October 2022.

[183] Michael Szmuk, Taylor P. Reynolds, and Behçet Açıkmeşe. Successive Convexification for Real-Time Six-Degree-of-Freedom Powered Descent Guidance with State-Triggered Constraints. Journal of Guidance, Control, and Dynamics, 43(8):1399–1413, 2020.

[184] Marco Sagliano, David Seelbinder, Stephan Theil, and Ping Lu. Six-Degree-of-Freedom Rocket Landing Optimization via Augmented Convex–Concave Decomposition. Journal of Guidance, Control, and Dynamics, pages 1–16, September 2023.

[185] Marco Sagliano, Ping Lu, Breanna Johnson, David Seelbinder, and Stephan Theil. Six-Degrees-of-Freedom Aero-Propulsive Entry Trajectory Optimization. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[186] Inc. CVX Research. CVX: Matlab software for disciplined convex programming, version 2.0. https://cvxr.com/cvx, August 2012.

[187] Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. Journal of Machine Learning Research, 2016. To appear.

[188] Madeleine Udell, Karanveer Mohan, David Zeng, Jenny Hong, Steven Diamond, and Stephen Boyd. Convex optimization in Julia. SC14 Workshop on High Performance Technical Computing in Dynamic Languages, 2014.

[189] Z. R. Putnam, S. H. Bairstow, R. D. Braun, and G. H. Barton. Improving Lunar Return Entry Range Capability Using Enhanced Skip Trajectory Guidance. Journal of Spacecraft and Rockets, 45(2):309–315, March 2008.

[190] Athul Pradeepkumar Girija, Sarag J. Saikia, James M. Longuski, Shyam Bhaskaran, Matthew S. Smith, and James A. Cutts. Feasibility and performance analysis of neptune aerocapture using heritage blunt-body aeroshells. Journal of Spacecraft and Rockets, 57(6):1186–1203, November–December 2020.

[191] Taylor Reynolds, Danylo Malyuta, Mehran Mesbahi, Behcet Acikmese, and John M. Carson. A Real-Time Algorithm for Non-Convex Powered Descent Guidance. In AIAA Scitech 2020 Forum. American Institute of Aeronautics and Astronautics, 2020.

[192] Kevin Tracy and Zachary Manchester. CPEG: A Convex Predictor-corrector Entry Guidance Algorithm. In 2022 IEEE Aerospace Conference (AERO), pages 1–10, Big Sky, MT, USA, March 2022. IEEE.

[193] Kevin S. Tracy, Giusy Falcone, and Zachary Manchester. Robust Entry Guidance with Atmospheric Adaptation. In AIAA SCITECH 2023 Forum. American Institute of Aeronautics and Astronautics, 2023.

[194] Rohan Deshmukh, Soumyo Dutta, Rafael Lugo, Ricardo Restrepo, Declan Mages, Breanna Johnson, Daniel Matz, Joshua Geiser, James B. Scoggins, Eli Shellabarger, Andrew Gomez-Delrio, and Joseph Williams. Performance Analysis of Aerocapture Systems for Uranus Orbiter. In AIAA SCITECH 2024 Forum. American Institute of Aeronautics and Astronautics, 2024.

[195] Purnanand Elango, Dayou Luo, Abhinav G. Kamath, Samet Uzun, Taewan Kim, and Behçet Açıkmeşe. Successive Convexification for Trajectory Optimization with Continuous-Time Constraint Satisfaction, April 2024. arXiv:2404.16826 [math].

[196] Daniel P Scharf, Behçet Açıkmeşe, Daniel Dueri, Joel Benito, and Jordi Casoliva. Implementation and experimental demonstration of onboard powered-descent guidance. Journal of Guidance, Control, and Dynamics, 40(2):213–229, 2017.

[197] Danylo Malyuta, Taylor Reynolds, Michael Szmuk, Mehran Mesbahi, Behcet Acikmese, and John M. Carson. Discretization Performance and Accuracy Analysis for the Rocket Powered Descent Guidance Problem. In AIAA Scitech 2019 Forum, San Diego, California, January 2019.

[198] Yuanqi Mao, Michael Szmuk, and Behçet Açıkmeşe. Successive convexification of non-convex optimal control problems and its convergence properties. In 2016 IEEE 55th Conference on Decision and Control (CDC), pages 3636–3641, December 2016.

[199] Christodoulos A Floudas. Nonlinear and Mixed-Integer Optimization: Fundamentals and Applications. Oxford University Press, November 1995.

[200] Danylo Malyuta, Taylor Reynolds, Michael Szmuk, Behcet Acikmese, and Mehran Mesbahi. Fast Trajectory Optimization via Successive Convexification for Spacecraft Rendezvous with Integer Constraints. In AIAA Scitech 2020 Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2020.

[201] Ping Lu. Convex–Concave Decomposition of Nonlinear Equality Constraints in Optimal Control. Journal of Guidance, Control, and Dynamics, 44(1):4–14, January 2021. Publisher: American Institute of Aeronautics and Astronautics.

[202] I. M. Ross, Q. Gong, M. Karpenko, and R. J. Proulx. Scaling and Balancing for High-Performance Computation of Optimal Controls. Journal of Guidance, Control, and Dynamics, 41(10):2086–2097, 2018.

[203] Ping Lu and Ryan Callan. Propellant-Optimal Powered Descent Guidance Revisited. Journal of Guidance, Control, and Dynamics, 46(2):215–230, February 2023.

[204] Marco Sagliano. Performance analysis of linear and nonlinear techniques for automatic scaling of discretized control problems. Operations Research Letters, 42(3):213–216, May 2014.

[205] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. SIAM Review, 59(1):65–98, 2017.

[206] Paul J. Goulart and Yuwen Chen. Clarabel: An interior-point solver for conic programs with quadratic objectives, 2024.

[207] Tyler D. Stoffel, Christopher D. Karlgaard, Todd R. White, and Thomas K. West. Fusion of In-Flight Aerothermodynamic Heating Sensor Measurements Using Kalman Filtering. Journal of Spacecraft and Rockets, 61(2):599–610, 2024.

[208] Daniel A. Matz, Christopher Cerimele, and Ronald R. Sostaric. Comparison of Aerocapture Performance Using Bank Control and Direct Force Control with Two Human-Scale Vehicles at Mars. In AIAA Scitech 2020 Forum, AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, January 2020.

[209] Sergio A. Sandoval and Daniel Matz. A Comparison of Bank Control and Direct Force Control for Aerocapture at Uranus. In AIAA SCITECH 2025 Forum. American Institute of Aeronautics and Astronautics, 2025.

[210] Alicia M. Dwyer-Cianciolo and Tara Polsgrove. Human Mars Entry, Descent, and Landing Architecture Study: Phase 2 Summary. In 2018 AIAA SPACE and Astronautics Forum and Exposition, AIAA SPACE Forum. American Institute of Aeronautics and Astronautics, September 2018.

[211] H. L. Justh, A. M. Dwyer Cianciolo, and J. Hoffman. Mars global reference atmospheric model (mars-gram): User guide. Technical report, National Aeronautics and Space Administration, NASA/TM-20210023957, November 2021.

[212] PW White and J Hoffman. Earth global reference atmospheric model (earth-gram): User guide. Technical report, National Aeronautics and Space Administration, NASA/TM-20210022157, November 2021.

[213] John T. Betts. Survey of numerical methods for trajectory optimization. Journal of Guidance, Control, and Dynamics, 21(2):193–207, 1998.

[214] Jeng-Shing Chern and Nguyen Xuan Vinh. Optimum reentry trajectories of a lifting vehicle. Technical Report 19790013968, University of Michigan Ann Arbor, NASA-CR-158486, January 1978.

[215] N.X. Vinh. Optimal trajectories in atmospheric flight. In L.G. Napolitano, editor, Space Mankind's Fourth Environment, pages 449–468. Pergamon, 1982.

[216] Abebe Geletu, Michael Klöppel, Hui Zhang, and Pu Li. Advances and applications of chance-constrained approaches to systems optimisation under uncertainty. International Journal of Systems Science, 44(7):1209–1232, 2013.

[217] Kenshiro Oguri and Jay W McMahon. Robust spacecraft guidance around small bodies under uncertainty: Stochastic optimal control approach. Journal of Guidance, Control, and Dynamics, 44(7):1295–1313, 2021.

[218] Kenshiro Oguri and Gregory Lantoine. Stochastic sequential convex programming for robust low-thrust trajectory design under uncertainty. In AAS/AIAA Astrodynamics Conference, pages 1–20, 2022.

[219] Jack Ridderhof, Joshua Pilipovsky, and Panagiotis Tsiotras. Chance-constrained covariance control for low-thrust minimum-fuel trajectory optimization. In 2021 AIAA/AAS Astrodynamics Specialists Conference. January 2021.

[220] Spencer Boone and Jay McMahon. Orbital guidance using higher-order state transition tensors. Journal of Guidance, Control, and Dynamics, 44(3):493–504, 2021.

[221] Kazuhide Okamoto, Maxim Goldshtein, and Panagiotis Tsiotras. Optimal Covariance Control for Stochastic Systems Under Chance Constraints. IEEE Control Systems Letters, 2(2):266–271, April 2018.

[222] Kenshiro Oguri, Masahiro Ono, and Jay W. McMahon. Convex Optimization over Sequential Linear Feedback Policies with Continuous-time Chance Constraints. In 2019 IEEE 58th Conference on Decision and Control (CDC), pages 6325–6331, December 2019. ISSN: 2576-2370.

[223] Michael Szmuk, Taylor P. Reynolds, and Behçet Açıkmeşe. Successive Convexification for Real-Time Six-Degree-of-Freedom Powered Descent Guidance with State-Triggered Constraints. Journal of Guidance, Control, and Dynamics, 43(8):1399–1413, 2020.

[224] Fengjiao Liu, George Rapakoulias, and Panagiotis Tsiotras. Optimal Covariance Steering for Discrete-Time Linear Stochastic Systems, October 2024. arXiv:2211.00618.

[225] Boris Benedikter, Alessandro Zavoli, Zhenbo Wang, Simone Pizzurro, and Enrico Cavallini. Convex Approach to Covariance Control with Application to Stochastic Low-Thrust Trajectory Optimization. Journal of Guidance, Control, and Dynamics, 45(11):2061–2075, 2022.

[226] Joshua Pilipovsky and Panagiotis Tsiotras. Computationally efficient chance constrained covariance control with output feedback, 2024. arXiv.2310.02485.

[227] Thomas Lipp and Stephen Boyd. Variations and extension of the convex–concave procedure. Optimization and Engineering, 17(2):263–287, June 2016.

[228] A. L. Yuille and Anand Rangarajan. The concave-convex procedure. Neural Computation, 15(4):915–936, 04 2003.

[229] Naoya Kumagai and Kenshiro Oguri. SEQUENTIAL CHANCE-CONSTRAINED COVARIANCE STEERING FOR ROBUST CISLUNAR TRAJECTORY DESIGN UNDER UNCERTAINTIES. In AAS/AIAA Astrodynamics Specialist Conference, August 2024.

[230] MOSEK ApS. MOSEK Optimizer API for Julia, 2024.

[231] Jens A. Rataczak, Jay W. McMahon, and Iain D. Boyd. Reachability Analysis of a Hypersonic Glide Vehicle using Particle Swarm Optimization. In AIAA SCITECH 2023 Forum. American Institute of Aeronautics and Astronautics.

[232] Jens A. Rataczak, Ross S. Chaudhry, Jay W. McMahon, and Iain D. Boyd. Investigation of Surface-Catalycity Effects on Hypersonic Glide Vehicle Trajectory Optimization. Journal of Spacecraft and Rockets, 0(0):1–13, 2024.

[233] H. F. Nelson. Effects of reduction of the ionization potential on solutions of the Rankine-Hugoniot equations for Jovian entry. Journal of Quantitative Spectroscopy and Radiative Transfer, 16(12):1061–1064, December 1976.

[234] Peter A. Gnoffo, K. James Weilmuenster, Robert D. Braun, and Christopher I. Cruz. Influence of sonic-line location on Mars Pathfinder Probe aerothermodynamics. Journal of Spacecraft and Rockets, 33(2):169–177, 1996. Publisher: American Institute of Aeronautics and Astronautics _eprint: https://doi.org/10.2514/3.26737.

[235] Karl T. Edquist, Prasun N. Desai, and Mark Schoenenberger. Aerodynamics for Mars Phoenix Entry Capsule. Journal of Spacecraft and Rockets, 48(5):713–726, 2011.

[236] Christian Agrell. Gaussian processes with linear operator inequality constraints, September 2019. arXiv:1901.03134 [stat].

[237] Daniel L. Engel and Zachary R. Putnam. Optimal Range Capabilities for Low-Lift-to-Drag Ratio Mars Entry Vehicles. Journal of Spacecraft and Rockets, 62(1):223–235, 2025.

[238] Jay W. McMahon, Davide Amato, Donald Kuettel, and Melis J. Grace. Stochastic Predictor-Corrector Guidance. In AIAA SCITECH 2022 Forum. American Institute of Aeronautics and Astronautics, 2022.

[239] Skye Mceowen, Daniel J. Calderone, Aman Tiwary, Jason S. Zhou, Taewan Kim, Purnanand Elango, and Behcet Acikmese. Auto-Tuned Primal-Dual Successive Convexification for Hypersonic Reentry Guidance. In AIAA SCITECH 2025 Forum. American Institute of Aeronautics and Astronautics, 2025.

[240] National Institute of Standards and Technology. Atomic spectra database, 2021.

Appendix A

Hydrogen-Helium-Methane Chemical Kinetic Models

A.1 Chemical Rate Model

Table A.1: Flow field chemical reactions for Ice Giant aerocapture aerothermodynamic analysis.

Number	Reaction	Third Body	A_f (cm^3/mol/s)	η_f	D_f (K)	U_i	Ref.
Dissociation							
1	$H_2 + M \leftrightarrow H + H + M$	Molecules	1.04e19	-1.00	51,958	[0.25, 3.0]	[26]
	$H_2 + M \leftrightarrow H + H + M$	Atoms except He	8.35e19	-1.00	51,958	[0.25, 3.0]	[26]
	$H_2 + M \leftrightarrow H + H + M$	He	4.17e18	-1.00	51,958	[0.25, 3.0]	[26]
2	$CH_4 + M \leftrightarrow CH_3 + H + M$	All	1.22e28	-2.91	54,000	[0.1, 10]	[118]
3	$CH_3 + M \leftrightarrow CH_2 + H + M$	All	2.82e14	0.00	42,460	[0.1, 10]	[118]
4	$CH_3 + M \leftrightarrow CH + H_2 + M$	All	5.00e15	0.00	42,800	[0.1, 10]	[118]
5	$CH_2 + M \leftrightarrow CH + H + M$	All	4.00e15	0.00	41,800	[0.1, 10]	[118]
6	$CH_2 + M \leftrightarrow C + H_2 + M$	All	1.30e14	0.00	29,700	[0.1, 10]	[118]
7	$CH + M \leftrightarrow C + H + M$	All	1.90e14	0.00	33,700	[0.5, 2]	[118]
8	$C_2 + M \leftrightarrow C+C+M$	All	3.72e14	0.00	69,800	[0.1, 10]	[118]
Ionization							
9	$H + H \leftrightarrow H^+ + e^- + H$		3.09e11	0.50	116,100	[0.1, 10]	[30, 112]
10	$H + He \leftrightarrow H^+ + e^- + He$		2.44e11	0.50	116,100	[0.1, 10]	[30, 112]
11	$H + e^- \leftrightarrow H^+ + e^- + e^-$		1.14e14	0.50	157,800	[0.1, 10]	[30, 112]
12	$H + e^- \leftrightarrow H^+ + e^- + e^-$		2.10e14	0.50	116,100	[0.1, 10]	[30, 112]
13	$C + e^- \leftrightarrow C^+ + e^- + e^-$		1.24e15	0.28	142,700	[0.5, 10]	[118]
Exchange							
14	$H_2 + C \leftrightarrow CH + H$		4.00e14	0.00	11,700	[0.5, 2]	[118]
15	$CH_4 + H \leftrightarrow CH_3 + H_2$		1.54e14	0.00	6,874	[0.1, 10]	[118]
16	$CH_3 + H \leftrightarrow CH_2 + H_2$		6.03e13	0.00	76,000	[0.5, 2]	[118]
17	$CH_2 + H \leftrightarrow CH + H_2$		4.21e08	-0.09	-1,560	[0.1, 10]	[118]
18	$CH + C \leftrightarrow C_2 + H$		2.00e14	0.00	0.0000	[0.1, 10]	[118]

Table A.2: Additional chemical reactions included when considering H_2^+.

Number	Reaction	A_f (cm^3/mol/s)	η_f	D_f (K)	U_i	Ref.
Dissociation						
1	$H_2^+ + e^- \leftrightarrow H^+ + H + e^-$	1.46e17	0.0	37,460	[0.1, 10]	[120]
Ionization						
2	$H + H \leftrightarrow H_2^+ + e^-$	1.13e15	-0.06	129,060	[0.5, 100]	[118, 126]
3	$H_2 + e^- \leftrightarrow H_2^+ + e^- + e^-$	4.05e13	0.52	180,767	[0.5, 2]	[118]
Exchange						
4	$H_2 + H^+ \leftrightarrow H_2^+ + H$	1.90e14	0.00	21,902	[0.1, 10]	[120]

A.2 Boltzmann-Weighted Rate Coefficient

The associative ionization (AI) reaction studied in this research is $H(n = 1) + H(n = i > 1) \leftrightarrow H_2^+ + e^-$, where n is the principal quantum number. In [40], the first 18 electronic states of hydrogen are explicitly modeled, and the state-specific AI rate data of Srećković et al. [126] is used. However, because LeMANS uses Park-style two-temperature kinetics, the individual electronic states of hydrogen are not tracked. Instead, a "Boltzmann-weighted" rate coefficient may be defined as [125]

$$k_{BW}(T_{tr}, T_v) = \sum_{n=2}^{18} F_1(T_v) F_n(T_v) k_n(T_{tr}) \tag{A.1}$$

where $k_n(T_{tr})$ is the AI rate coefficient of the channel corresponding to the n-th excited state of hydrogen, and the F_n is the Boltzmann fraction of the n-th electronic state. The Boltzmann fraction is defined as

$$F_n(T_v) = \frac{g_n \exp\left(-\theta_n / T_v\right)}{\sum_{i=1}^{18} g_i \exp\left(-\theta_i / T_v\right)} \tag{A.2}$$

where θ_n and g_n are the characteristic electronic temperature and degeneracy of the n-th electronic state, respectively. Values of θ_n and g_n are taken from the NIST database [240] and are summarized in Appendix A.3. Because ionization occurs preferentially from electronically excited configurations, whose populations are often below that of a Boltzmann distribution at $T_v = T_e$ in a net-ionizing environment, the Boltzmann-weighted rate can serve as an upper bound in an ionizing flow.

To validate the implementation of the Boltzmann-weighted rate, its prediction of the rate of associative ionization is compared to the net state-specific rate of associative ionization. Number

densities of the electronic states of atomic hydrogen are computed using NEQAIR along two lines of sight. The first line of sight corresponds to the stagnation line, and the second corresponds to the shoulder. The net state-specific rate is give by

$$r_{\text{net}}(T) = \sum_{i=2}^{18} n_{\text{H}(n=1)} n_{\text{H}(n=i)} k_i(T) \tag{A.3}$$

where $n_{\text{H}(n=i)}$ is the density of the i-th excited state of hydrogen predicted by NEQAIR and k_i is the state-specific rate corresponding to the i-th channel. Figure A.1 displays the AI rates predicted by the Boltzmann-weighted rate and the net state-to-state rate. In these results, NEQAIR assumes all state populations are Boltzmann, so the agreement between the two rates is expected. Figure A.2 relaxes the Boltzmann population assumption in NEQAIR and instead uses the QSS approximation. The gray lines in each figure denote the AI rate for each electronic state of atomic hydrogen, and the blue line denotes the AI rate from Lino da Silva [118].

Generally, the Boltzmann-weighted rate overpredicts the rate of AI relative to the net state-to-state rate when the gas is net ionizing and underpredicts the rate of AI when the flow is net recombining. As shown in Figure A.2, the excited states of atomic hydrogen are depleted relative to the corresponding Boltzmann value in Figure A.1. In a net ionization flow, there are fewer ions than at equilibrium, so the Saha-Boltzmann state population will be below the Boltzmann population. Therefore, the period of net ionization should see an overprediction from the Boltzmann-weighted rate, as the excited states have not had time to reach equilibrium. Along the stagnation line, the net state-to-state rate and Boltzmann-weighted rate converge near $x = -0.1$ m, indicating that the state populations have reached equilibrium. As the gas recombines in the boundary layer, the opposite phenomenon occurs, and the excited states are overpopulated relative to Boltzmann (i.e. Saha-Boltzmann above Boltzmann). For all cases, the rate from Lino da Silva is lower than the Boltzmann-weighted rate.

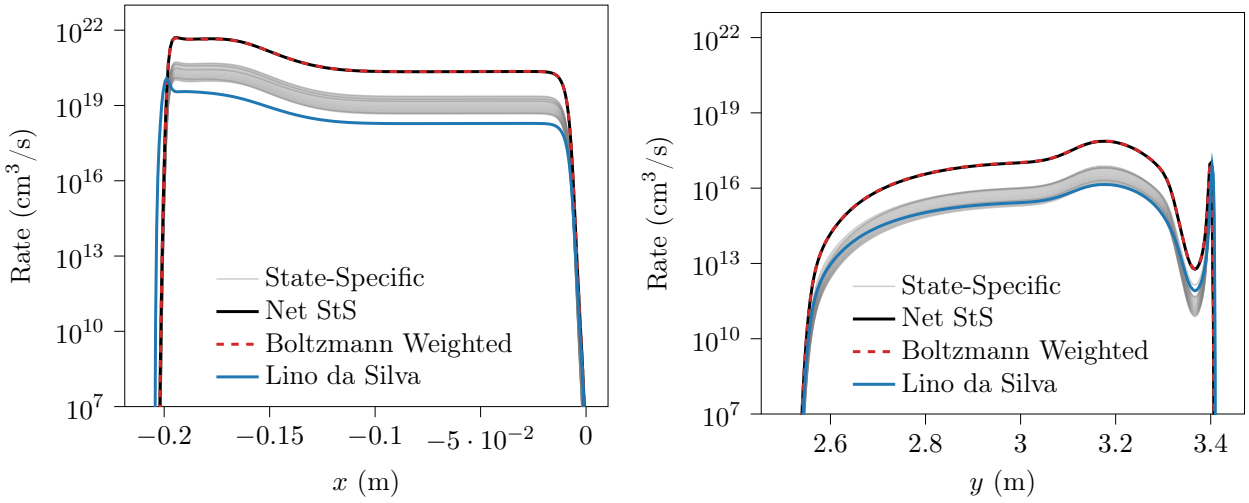

(a) Stagnation line.

(b) Shoulder.

Figure A.1: Comparison of AI rate predicted by Boltzmann-weighted and net state-to-state rate at the stagnation line (a) and shoulder (b) using a Boltzmann assumption in NEQAIR.

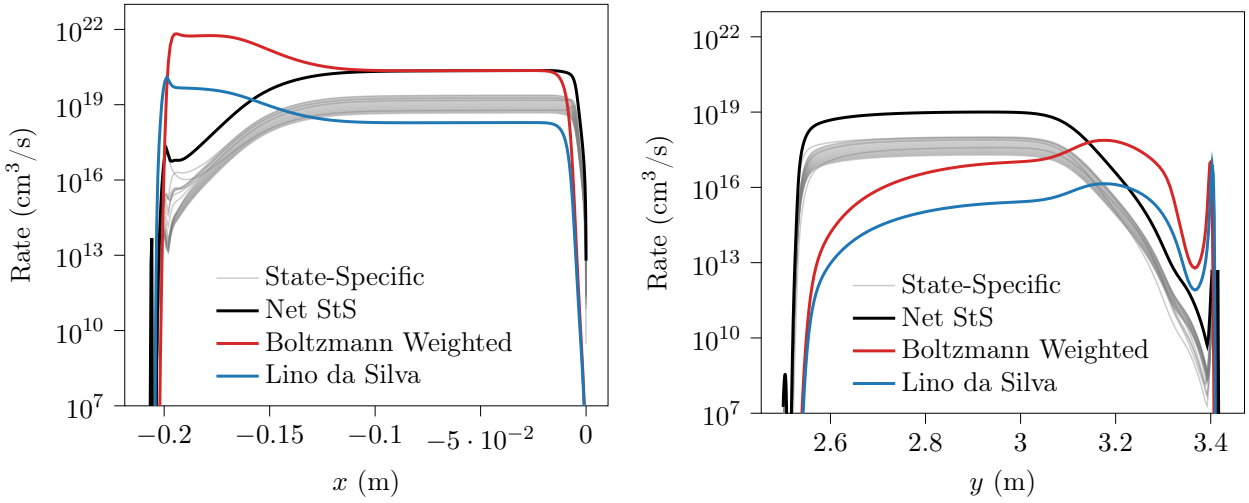

(a) Stagnation line.

(b) Shoulder.

Figure A.2: Comparison of AI rate predicted by Boltzmann-weighted and net state-to-state rate at the stagnation line (a) and shoulder (b) using a QSS model in NEQAIR.

A.3 Atomic Hydrogen Electronic Energy Data

Electronic-energy level data for neutral atomic hydrogen, taken from the NIST [240], is summarized in Table A.3. The degeneracy g of the hydrogen atom may be expressed generally as $g_n = 2n^2$, where n is the principal quantum number. Characteristic electronic temperatures θ_n are computed from the corresponding energy via $\theta_n = 1.986 \times 10^{-23} E_n/k_B$, where the $k_B = 1.3806 \times 10^{-23}$ J/K is the Boltzmann constant. The factor of 1.986×10^{-23} comes from the assumption that E_n is expressed with units of cm^{-1}.

Table A.3: Atomic hydrogen electronic energy levels, taken from [240].

Energy Level n	Energy (cm^{-1})	Characteristic Electronic Temp θ_n (K)	Degeneracy g_n
1	0	0	2
2	82259.1	118231.7981	8
3	97492.3	140126.6235	18
4	102824	147789.9273	32
5	105292	151337.2076	50
6	106632	153263.2025	72
7	107440	154424.5487	98
8	107965	155179.1363	128
9	108325	155696.5677	162
10	108582	156065.9563	200
11	108772	156339.0451	242
12	108917	156547.455	288
13	109030	156709.871	338
14	109119	156837.7916	392
15	109191	156941.2779	450
16	109250	157026.0792	512
17	109299	157096.5073	578
18	109340	157155.437	648

Appendix B

Useful Mathematical Results

B.1 Analytical Jacobians

This section provides expressions for the Jacobians of various functions used throughout this dissertation. These include the equations of motion and terminal constraints used in the development of the aerocapture trajectory optimization problems.

B.1.1 Equations of Motion

The equations of motion which describe the motion of a spacecraft in spherical coordinates about an oblate, rotating planet are

$$\dot{r} = v \sin\gamma \tag{B.1a}$$

$$\dot{\theta} = \frac{v \cos\gamma \cos\psi}{r \cos\phi} \tag{B.1b}$$

$$\dot{\phi} = \frac{v \cos\gamma \sin\psi}{r} \tag{B.1c}$$

$$\dot{v} = -D - g_r \sin\gamma - g_\phi \cos\gamma \sin\psi + \Omega^2 r \cos\phi \left(\sin\gamma \cos\phi - \cos\gamma \sin\phi \sin\psi\right) \tag{B.1d}$$

$$v\dot{\gamma} = L \cos\sigma - g_r \cos\gamma + \frac{v^2}{r} \cos\gamma + g_\phi \sin\gamma \sin\psi + 2\Omega v \cos\phi \cos\psi + \tag{B.1e}$$
$$\Omega^2 r \cos\phi \left(\cos\gamma \cos\phi + \sin\gamma \sin\phi \sin\psi\right)$$

$$v\dot{\psi} = \frac{L \sin\sigma}{\cos\gamma} - \frac{v^2}{r} \cos\gamma \cos\psi \tan\phi - g_\phi \frac{\cos\psi}{\cos\gamma} + 2\Omega v \left(\tan\gamma \cos\phi \sin\psi - \sin\phi\right) \tag{B.1f}$$
$$- \Omega^2 r \frac{\sin\phi \cos\phi \cos\psi}{\cos\gamma}$$

$$\dot{\sigma} = u \quad . \tag{B.1g}$$

These may be compactly represented as $\boldsymbol{f}(\boldsymbol{x}(t), u(t))$, where $\boldsymbol{x} = [r, \theta, \phi, v, \gamma, \psi, \sigma]$ is the state, and the bank-angle rate $\dot{\sigma}$ is taken to be the control u. A first-order Taylor series expansion about a reference trajectory $\bar{\boldsymbol{x}}(t)$ and control $\bar{u}(t)$ yields

$$\boldsymbol{f}(\boldsymbol{x}(t), u(t)) \approx \boldsymbol{f}(\bar{\boldsymbol{x}}(t), \bar{u}(t)) + \left.\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}\right|_{\bar{\boldsymbol{x}}, u} (\boldsymbol{x}(t) - \bar{\boldsymbol{x}}(t)) + \left.\frac{\partial \boldsymbol{f}}{\partial u}\right|_{\bar{\boldsymbol{x}}, u} (u(t) - \bar{u}(t)) \quad . \tag{B.2}$$

The partial derivatives $\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}$ and $\frac{\partial \boldsymbol{f}}{\partial u}$ are the Jacobians of the equations of motion with respect to the state and control, respectively. They are expressed as

$$\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} = A = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} & a_{17} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} & a_{27} \\ a_{31} & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} & a_{37} \\ a_{41} & a_{42} & a_{43} & a_{44} & a_{45} & a_{46} & a_{47} \\ a_{51} & a_{52} & a_{53} & a_{54} & a_{55} & a_{56} & a_{57} \\ a_{61} & a_{62} & a_{63} & a_{64} & a_{65} & a_{66} & a_{67} \\ a_{71} & a_{72} & a_{73} & a_{74} & a_{75} & a_{76} & a_{77} \end{bmatrix} \tag{B.3}$$

$$\frac{\partial \boldsymbol{f}}{\partial u} = B = \begin{bmatrix} 0_{6\times 1} \\ 1 \end{bmatrix} \tag{B.4}$$

where

$$a_{11} = a_{12} = a_{13} = a_{16} = a_{17} = 0 \tag{B.5}$$

$$a_{14} = \sin\gamma \tag{B.6}$$

$$a_{15} = v\cos\gamma \tag{B.7}$$

$$a_{21} = -\frac{v\cos\gamma\cos\psi\sec\phi}{r^2} \tag{B.8}$$

$$a_{22} = a_{27} = 0 \tag{B.9}$$

$$a_{23} = \frac{v\cos\gamma\cos\psi\sec\phi\tan\phi}{r} \tag{B.10}$$

$$a_{24} = \frac{\cos\gamma\cos\psi\sec\phi}{r} \tag{B.11}$$

$$a_{25} = -\frac{v\cos\psi\sec\phi\sin\gamma}{r} \tag{B.12}$$

$$a_{26} = -\frac{v\cos\gamma\sec\phi\sin\psi}{r} \tag{B.13}$$

$$a_{31} = -\frac{v\cos\gamma\sin\psi}{r^2} \tag{B.14}$$

$$a_{32} = a_{33} = a_{37} = 0 \tag{B.15}$$

$$a_{34} = \frac{\cos\gamma\sin\psi}{r} \tag{B.16}$$

$$a_{35} = -\frac{v\sin\gamma\sin\psi}{r} \tag{B.17}$$

$$a_{36} = \frac{v\cos\gamma\cos\psi}{r} \tag{B.18}$$

$$a_{41} = -\frac{1}{2}\frac{\partial\rho}{\partial r}\frac{v^2}{\beta}\rho_D + \sin\gamma\left(\Omega^2\cos^2\phi + \frac{\mu}{r^5}\left(2r^2 + 6J_2R_e^2 - 18J_2R_e^2\sin^2\phi\right)\right) + \frac{1}{r^5}\left(12J_2R_e^2\mu - r^5\Omega^2\right)\cos\gamma\cos\phi\sin\phi\sin\psi \tag{B.19}$$

$$a_{42} = a_{47} = 0 \tag{B.20}$$

$$a_{43} = \frac{1}{r^4}\left[\left(9J_2R_e^2\mu - 2r^5\Omega^2\right)\cos\phi\sin\gamma\sin\phi + \left(-3J_2R_e^2\mu - r^5\Omega^2\right)\cos\gamma\cos^2\phi\sin\psi + \left(3J_2Re^2\mu + r^5\Omega^2\right)\cos\gamma\sin^2\phi\sin\psi\right] \tag{B.21}$$

$$a_{44} = -\frac{v\rho}{\beta}\rho_D \tag{B.22}$$

$$a_{45} = \frac{1}{r^4}\left[-\mu\cos\gamma\left(r^2 + J_2R_e^2\left(1.5 - 4.5\sin^2\phi\right)\right) + 3J_2R_e^2\mu\cos\phi\sin\gamma\sin\phi\sin\psi + r^5\Omega^2\cos\phi\left(\cos\gamma\cos\phi + \sin\gamma\sin\phi\sin\psi\right)\right] \tag{B.23}$$

$$a_{46} = -\frac{1}{r^4}\left(3J_2R_e^2\mu + r^5\Omega^2\right)\cos\gamma\cos\phi\cos\psi\sin\phi \tag{B.24}$$

$$a_{51} = \frac{1}{2}\frac{\frac{\partial\rho}{\partial r}L/Dv^2\cos\sigma}{\beta v}\rho_L + \frac{1}{r^5v}\left[\cos\gamma\left(-r^3v^2 + 2r^2\mu + 6J_2R_e^2\mu + r^5\Omega^2\cos^2\phi - 18J_2R_e^2\mu\sin^2\phi\right) + \left(-12J_2R_e^2\mu + r^5\Omega^2\right)\cos\phi\sin\gamma\sin\phi\sin\psi\right] \tag{B.25}$$

$$a_{52} = 0 \tag{B.26}$$

$$a_{53} = \frac{1}{r^4v}\left[\left(9J_2R_e^2\mu - 2r^5\Omega^2\right)\cos\gamma\cos\phi\sin\phi - 2r^4v\Omega\cos\psi\sin\phi + \left(3J_2R_e^2\mu + r^5\Omega^2\right)\cos 2\phi\sin\gamma\sin\psi\right] \tag{B.27}$$

$$a_{54} = \frac{1}{2}\frac{\rho L/D\cos\sigma}{\beta}\rho_L + \frac{1}{r^4v^2}\left[\cos\gamma\left(r^3v^2 + r^2\mu + 1.5J_2R_e^2\mu - r^5\Omega^2\cos^2\phi - 4.5J_2R_e^2\mu\sin^2\phi\right) +\right.$$

$$\left(-3J_2R_e^2\mu - r^5\Omega^2\right)\cos\phi\sin\gamma\sin\phi\sin\psi\Big] \tag{B.28}$$

$$a_{56} = \cos\phi\left[\frac{1}{r^4v}\left(3J_2R_e^2\mu + r^5\Omega^2\right)\cos\psi\sin\gamma\sin\phi - 2\Omega\sin\psi\right] \tag{B.29}$$

$$a_{57} = -\frac{1}{2}\frac{L/Dv\rho\sin\sigma}{\beta}\rho_L \tag{B.30}$$

$$a_{61} = \frac{1}{2}\frac{\frac{\partial\rho}{\partial r}L/Dv\sec\gamma\sin\sigma}{\beta}\rho_L + \frac{\cos\psi}{r^5v}\left[\left(12J_2R_e^2\mu - r^5\Omega^2\right)\cos\phi\sec\gamma\sin\phi + r^3v^2\cos\gamma\tan\phi\right] \tag{B.31}$$

$$a_{62} = 0 \tag{B.32}$$

$$a_{63} = -\frac{1}{r^4v}\left(3J_2R_e^2\mu + r^5\Omega^2\right)\cos 2\phi\cos\psi\sec\gamma - \frac{1}{r}v\cos\gamma\cos\psi\sec^2\phi - 2\Omega\left(\cos\phi + \sin\phi\sin\psi\tan\gamma\right) \tag{B.33}$$

$$a_{64} = \sec\gamma\left[\frac{1}{2}\frac{L/D\rho\sin\sigma}{\beta}\rho_L + \frac{1}{r^4v^2}\left(3J_2R_e^2\mu + r^5\Omega^2\right)\cos\phi\cos\psi\sin\phi\right] - \frac{1}{r}\cos\gamma\cos\psi\tan\phi \tag{B.34}$$

$$a_{65} = \frac{1}{2}\frac{L/Dv\rho\sec\gamma\sin\sigma\tan\gamma}{\beta}\rho_L + \cos\phi\sec\gamma\left[2\Omega\sec\gamma\sin\phi - \frac{1}{r^4v}\left(3J_2R_e^2\mu + r^5\Omega^2\right)\cos\psi\sin\phi\tan\gamma\right] + \frac{1}{r}v\cos\gamma\cos\psi\tan\gamma \tag{B.35}$$

$$a_{66} = \cos\phi\left[\frac{1}{r^4v}\left(3J_2R_e^2\mu + r^5\Omega^2\right)\sec\gamma\sin\phi\sin\psi + 2\Omega\cos\psi\tan\gamma\right] + \frac{1}{r}v\cos\gamma\sin\psi\tan\phi \tag{B.36}$$

$$a_{67} = \frac{1}{2}\frac{L/Dv\rho\cos\sigma\sec\gamma}{\beta}\rho_L \tag{B.37}$$

$$a_{71} = a_{72} = a_{73} = a_{74} = a_{75} = a_{76} = a_{77} = 0 \quad . \tag{B.38}$$

Here, $L/D = \frac{C_L}{C_D}$ is the lift-to-drag ratio, $\beta = \frac{m}{C_D A_{\text{ref}}}$ is the ballistic coefficient, J_2 is the second spherical harmonic for the gravitational model, R_e is the equatorial radius of the planet, μ is the gravitational constant of the planet, and ρ_L and ρ_D are the first-order fading memory filter multipliers. The partial derivative of the density with respect to the radius depends on the density model. If the density is assumed to vary exponentially, such that

$$\rho(h) = \rho_0\exp\left(-h/H\right) \tag{B.39}$$

where H is the scale height, the Jacobian is

$$\frac{\partial \rho}{\partial r} = -\frac{\rho_0}{H} \exp\left(-h/H\right) \quad . \tag{B.40}$$

Otherwise, if the density is implemented as a lookup table, the Jacobian $\frac{\partial \rho}{\partial r}$ may be computed using either a finite-difference approximation or automatic differentiation. This work uses the `ForwardDiff.jl` toolbox in `Julia` for all automatic differentiation.

B.1.1.1 Consideration of Augmented State

When the state is augmented with the heat load and square of the load factor, as in Chapters 6 and 7, additional partial derivatives must be computed. The square of the load factor and the heat load are computed as

$$\tilde{a} = \frac{1}{g_0^2}\left(L^2 + D^2\right) \tag{B.41}$$

$$Q = \int_0^t q(t)\mathrm{d}t \quad . \tag{B.42}$$

Substitution for the lift and drag accelerations yields

$$\tilde{a} = \frac{1}{g_0^2}\left[\left(\frac{\rho v^2 L/D}{2\beta}\right)^2 + \left(\frac{\rho v^2}{2\beta}\right)^2\right] \tag{B.43}$$

which is simplified to

$$\tilde{a} = \left(\frac{L/D^2 + 1}{4 g_0^2 \beta^2}\right)\rho^2 v^4 \quad . \tag{B.44}$$

Differentiating with respect to time yields

$$\dot{\tilde{a}} = \left(\frac{L/D^2 + 1}{4 g_0^2 \beta^2}\right)\left(2\rho v^4 \frac{\mathrm{d}\rho}{\mathrm{d}t} + 4\rho^2 v^3 \frac{\mathrm{d}v}{\mathrm{d}t}\right) \quad . \tag{B.45}$$

It is assumed that the density only varies with altitude (i.e. $\rho = \rho(r)$) such that the dynamics of the density are written as

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = \frac{\partial \rho}{\partial r}\frac{\partial r}{\partial t} = \frac{\partial \rho}{\partial r} v \sin\gamma \tag{B.46}$$

where Eq. B.1a has been used, and the partial derivative $\frac{\partial \rho}{\partial r}$ is computed as previously described. Using Eq. B.1d, the dynamics of the squared load factor are then

$$\dot{\tilde{a}} = \left(\frac{L/D^2 + 1}{4 g_0^2 \beta^2}\right)\left[2\rho v^5 \sin\gamma \frac{\partial \rho}{\partial r} +\right.$$

$$4\rho^2 v^3 \left(-D - g_r \sin\gamma - g_\phi \cos\gamma \sin\psi + \Omega^2 r \cos\phi \left(\sin\gamma \cos\phi - \cos\gamma \sin\phi \sin\psi\right)\right)\Big] \quad . \tag{B.47}$$

Finally, differentiating Eq. B.42 with respect to time yields

$$\frac{\partial Q}{\partial t} = \dot{Q} = q \tag{B.48}$$

where q is the heat flux. Equations B.47 and B.48 are then differentiated with respect to the augmented state and included in matrix A. Because the density, aerodynamic, and heat-flux models used in this dissertation are not always analytic (i.e. exponential, Sutton-Graves, etc.), the partial derivatives of the squared load factor and heat load are computed using automatic differentiation.

B.1.1.2 Generalized Time-Dilation Formulation

A free time of flight problem may be cast as a fixed time of flight problem using a transformation called *generalized time dilation* [195]. In this approach, a strictly increasing, continuously differentiable mapping $t : [0, 1] \mapsto \mathbb{R}_+$, with boundary conditions $t(0) = t_0$ and $t(1) = t_f$, is defined. The derivative of this mapping is defined as

$$s(\tau) = \frac{\mathrm{d}t(\tau)}{\mathrm{d}\tau} = \overset{\circ}{t}(\tau) \tag{B.49}$$

for $\tau \in [0, 1]$, where $\overset{\circ}{(\cdot)}$ denotes the derivative with respect to the non-dimensional time. The term $s(\tau)$ is referred to as the *time dilation factor* and is treated as an additional control input. Next, the state $\boldsymbol{x} = [r, \theta, \phi, v, \gamma, \psi, \sigma]$ is augmented with the physical time t such that

$$\tilde{\boldsymbol{x}}(\tau) = [\boldsymbol{x}(t(\tau)),\ t(\tau)] \tag{B.50}$$

where $\tilde{\boldsymbol{x}}$ is the "augmented" state. The augmented state dynamics may then be written as

$$\overset{\circ}{\tilde{\boldsymbol{x}}}(\tau) = \begin{bmatrix} \dot{\boldsymbol{x}}(t) \\ 1 \end{bmatrix} \frac{\mathrm{d}t(\tau)}{\mathrm{d}\tau} = \begin{bmatrix} \boldsymbol{f}(\boldsymbol{x}(t), u(t)) \\ 1 \end{bmatrix} \frac{\mathrm{d}t(\tau)}{\mathrm{d}\tau} = \boldsymbol{F}(\tilde{\boldsymbol{x}}(\tau), \tilde{\boldsymbol{u}}(\tau)) \tag{B.51}$$

where $\tilde{\boldsymbol{u}} = [u(t(\tau)), s(\tau)]$ is the original control augmented with the time dilation factor. The Jacobian of the augmented state dynamics with respect to the augmented state and augmented

control are expressed as

$$\frac{\partial \boldsymbol{F}}{\partial \tilde{\boldsymbol{x}}} = A = \begin{bmatrix} s(\tau)\dfrac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} & s(\tau)\dfrac{\partial \boldsymbol{f}}{\partial t} \\ \dfrac{\partial s}{\partial \boldsymbol{x}} & \dfrac{\partial s}{\partial t} \end{bmatrix} = \begin{bmatrix} s(\tau)\dfrac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} & 0_{n\times 1} \\ 0_{1\times n} & 0 \end{bmatrix} \tag{B.52}$$

$$\frac{\partial \boldsymbol{F}}{\partial \tilde{\boldsymbol{u}}} = B = \begin{bmatrix} s(\tau)\dfrac{\partial \boldsymbol{f}}{\partial u} & \dfrac{\partial}{\partial s}\left(s(\tau)\boldsymbol{f}\right) \\ \dfrac{\partial s}{\partial u} & \dfrac{\partial s}{\partial s} \end{bmatrix} = \begin{bmatrix} s(\tau)\dfrac{\partial \boldsymbol{f}}{\partial u} & \boldsymbol{f} \\ 0 & 1 \end{bmatrix} \tag{B.53}$$

where n is the dimension of the non-augmented state, and the partial derivatives $\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}$ and $\frac{\partial \boldsymbol{f}}{\partial u}$ may be computed as in Section B.1.1. This formulation also works when the state is augmented with additional variables such as the heat load, and the Jacobian $\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}$ is recomputed accordingly.

B.1.2 Terminal Constraints

The terminal constraints considered in this dissertation include those on the final apoapsis radius and inclination. The expression used to enforce the apoapsis radius target is

$$\begin{aligned} \varepsilon(\boldsymbol{x}_f) = &\frac{\mu}{r_f}\left[1 - \frac{J_2}{2}\left(\frac{R_e}{r_f}\right)^2\left(3\sin^2\phi_f - 1\right)\right] - \frac{v_f^2}{2} \\ &- \frac{\mu}{r_{a,\text{target}}}\left[1 - \frac{J_2}{2}\left(\frac{R_e}{r_{a,\text{target}}}\right)^2\left(3\sin^2\phi_{\text{target}} - 1\right)\right] + \frac{r_f^2 v_f^2 \cos^2\gamma_f}{2r_{a,\text{target}}^2} = 0 \quad . \end{aligned} \tag{B.54}$$

The Jacobian of Eq. B.54 is written as

$$\frac{\partial \varepsilon}{\partial \boldsymbol{x}} = \begin{bmatrix} e_1 & e_2 & e_3 & e_4 & e_5 & e_6 & e_7 \end{bmatrix}^T \tag{B.55}$$

where

$$e_1 = \frac{r v^2 \cos^2\gamma}{r_{a,\text{target}}^2} - \frac{\mu}{4r^4}\left(4r^2 - 3J_2R_e^2 + 9J_2R_e^2\cos 2\phi\right) \tag{B.56}$$

$$e_2 = e_6 = e_7 = 0 \tag{B.57}$$

$$e_3 = -\frac{3J_2R_e^2\mu\cos\phi\sin\phi}{r^3} \tag{B.58}$$

$$e_4 = -v + \frac{r^2 v\cos^2\gamma}{r_{a,\text{target}}^2} \tag{B.59}$$

$$e_5 - \frac{r^2 v^2\cos\gamma\sin\gamma}{r_{a,\text{target}}^2} \tag{B.60}$$

The expression used to enforce the inclination constraint is

$$\cos\phi\cos\psi - \cos i_{\text{target}} = 0 \tag{B.61}$$

and its Jacobian is

$$\frac{\partial \cos i}{\partial \boldsymbol{x}} = \begin{bmatrix} i_1 & i_2 & i_3 & i_4 & i_5 & i_6 & i_7 \end{bmatrix}^T \tag{B.62}$$

where

$$i_1 = i_2 = i_4 = i_5 = i_7 = 0 \tag{B.63}$$

$$i_3 = -\sin\phi\cos\psi \tag{B.64}$$

$$i_6 = -\cos\phi\sin\psi \quad . \tag{B.65}$$

Finally, in cases where the periapsis maximization term is included in the cost function, its Jacobian must be computed. The inverse of the periapsis radius is expressed as

$$\frac{1}{r_p(\boldsymbol{x}_f)} = p(\boldsymbol{x}_f) = \left[a\left(1 - \sqrt{1 - \frac{v_f^2 r_f^2 \cos^2\gamma_f}{\mu a}}\right)\right]^{-1} \quad . \tag{B.66}$$

The Jacobian of Eq. B.66 is expressed as

$$\frac{\partial p}{\partial \boldsymbol{x}} = \begin{bmatrix} p_1 & p_2 & p_3 & p_4 & p_5 & p_6 & p_7 \end{bmatrix}^T \tag{B.67}$$

where

$$p_1 = \frac{1}{r^2\mu^3 P_2}\left[-rv^2\left(rv^2 - 3\mu\right)\left(rv^2 - 2\mu\right)\cos^2\gamma - 2\mu^3(P_1 - 1)\right] \tag{B.68}$$

$$p_2 = p_3 = p_6 = p_7 = 0 \tag{B.69}$$

$$p_4 = \frac{1}{\mu^2 P_2}\left[2rv^3\cos^2\gamma - 2v\mu\left(P_1 + \cos 2\gamma\right)\right] \tag{B.70}$$

$$p_5 = \frac{1}{\mu^3 P_2}\left[\left(rv^3 - 2v\mu\right)^2\cos\gamma\sin\gamma\right] \tag{B.71}$$

and

$$P_1 = \sqrt{1 + \frac{1}{\mu^2} rv^2\left(rv^2 - 2\mu\right)\cos^2\gamma} \tag{B.72}$$

$$P_2 = P_1\left(P_1 - 1\right)^2 \quad . \tag{B.73}$$

B.1.2.1 Consideration of the Augmented State

When the state is augmented with the heat load and square of the load factor, no additional partial derivatives are required for the terminal constraints. This is because they do not depend explicitly on either the squared load factor or the heat load. Additional zeros are included in the Jacobians to account for the increased state dimension.

B.2 Discretization of Continuous-Time Linear Systems

Consider the following equations of motion for a continuous-time, linear, time-varying system

$$\dot{\boldsymbol{x}}(t) = A(t)\boldsymbol{x}(t) + B(t)\boldsymbol{u}(t) + \boldsymbol{z}(t) \tag{B.74}$$

where $\boldsymbol{x} \in \mathbb{R}^n$, is the state, $\boldsymbol{u} \in \mathbb{R}^m$, is the control, $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$ are known system matrices, and $\boldsymbol{z}(t) \in \mathbb{R}^n$ is a known drift term. To convert Equation B.74 into an equivalent discrete-time formulation, the following result from linear system theory is utilized. Let $\Phi(t, t_0) : \mathbb{R} \mapsto \mathbb{R}^{n\times n}$ denote the state transition matrix (STM) for the system in Equation B.74 over the time interval $[t_0, t]$. From linear systems theory, the STM evolves according to

$$\dot{\Phi}(t, t_0) = A(t)\Phi(t, t_0), \qquad \Phi(t_0, t_0) = I_{n\times n} \qquad . \tag{B.75}$$

Using the STM, Equation B.74 is integrated from $t = t_k$ to $t = t_{k+1}$ which yields

$$\boldsymbol{x}_{k+1} = \Phi(t_{k+1}, t_k)\boldsymbol{x}_k + \int_{t_k}^{t_{k+1}} \Phi(t_{k+1}, \tau)B(\tau)\boldsymbol{u}(\tau)\mathrm{d}\tau + \int_{t_k}^{t_{k+1}} \Phi(t_{k+1}, \tau)\boldsymbol{z}(\tau)\mathrm{d}\tau \tag{B.76}$$

where $\boldsymbol{x}_k$ denotes the state at $t = t_k$. Note that integrating the right-hand side of Equation B.76 requires computing the STM $\Phi(t_{k+1}, \tau)$ for each intermediate time $\tau \in [t_k, t_{k+1}]$. This introduces substantially more computational complexity than desired, so an alternative formulation of Equation B.76 is defined. Recall that the following properties hold for the STM

$$\Phi(t_2, t_0) = \Phi(t_2, t_1)\Phi(t_1, t_0) \tag{B.77}$$

$$\Phi(t_1, t_0) = \Phi(t_0, t_1)^{-1} \quad . \tag{B.78}$$

This allows us to rewrite Equation B.76 as

$$\boldsymbol{x}_{k+1} = \Phi(t_{k+1}, t_k)\boldsymbol{x}_k + \int_{t_k}^{t_{k+1}} \Phi(t_{k+1}, t_k)\Phi(t_k, \tau)B(\tau)\boldsymbol{u}(\tau)\mathrm{d}\tau + \int_{t_k}^{t_k} \Phi(t_{k+1}, t_k)\Phi(t_k, \tau)\boldsymbol{z}(\tau)\mathrm{d}\tau \tag{B.79}$$

which is simplified to

$$\boldsymbol{x}_{k+1} = \Phi(t_{k+1}, t_k)\boldsymbol{x}_k + \Phi(t_{k+1}, t_k)\int_{t_k}^{t_{k+1}} \Phi(\tau, t_k)^{-1}B(\tau)\boldsymbol{u}(\tau)\mathrm{d}\tau + \Phi(t_{k+1}, t_k)\int_{t_k}^{t_k} \Phi(\tau, t_k)^{-1}\boldsymbol{z}(\tau)\mathrm{d}\tau \quad . \tag{B.80}$$

The form of Equation B.80 only requires the integration of the STM from t_k to t_{k+1} which is more computationally efficient than using Equation B.76.

Next, an assumption on the control $\boldsymbol{u}(\tau)$ is required. This dissertation uses a first-order hold, which assumes a linear variation between temporal nodes k and $k+1$

$$u(\tau) = \sigma_k^-(\tau)u(\tau_k) + \sigma_k^+(\tau)u(\tau_{k+1}) \qquad \forall\ \tau \in [\tau_k, \tau_{k+1}) \tag{B.81}$$

where $\sigma_k^-(\tau)$ and $\sigma_k^+(\tau)$ are defined as

$$\sigma_k^-(\tau) = \frac{\tau_{k+1} - \tau}{\tau_{k+1} - \tau_k} \qquad \sigma_k^+(\tau) = \frac{\tau - \tau_k}{\tau_{k+1} - \tau_k} \quad . \tag{B.82}$$

This allows Equation B.80 to be simplified to

$$\boldsymbol{x}_{k+1} = A_k\boldsymbol{x}_k + B_k^-\boldsymbol{u}_k + B_k^+\boldsymbol{u}_{k+1} + \boldsymbol{z}_k \tag{B.83}$$

where

$$A_k = \Phi(\tau_{k+1}, \tau_k) \tag{B.84}$$

$$B_k^- = A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k)\sigma_k^-(\xi)B(\xi)\mathrm{d}\xi \tag{B.85}$$

$$B_k^+ = A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k)\sigma_k^+(\xi)B(\xi)\mathrm{d}\xi \tag{B.86}$$

$$\boldsymbol{z}_k = A_k \int_{\tau_k}^{\tau_{k+1}} \Phi^{-1}(\xi, \tau_k)\boldsymbol{z}(\xi)\mathrm{d}\xi \quad . \tag{B.87}$$

B.2.1 Stochastic Systems

Next, consider the following stochastic differential equation (SDE)

$$\mathrm{d}\boldsymbol{x}(t) = (A(t)\boldsymbol{x}(t) + B(t)\boldsymbol{u}(t) + \boldsymbol{z}(t))\mathrm{d}t + G(t)\mathrm{d}\boldsymbol{w}(t) \tag{B.88}$$

where $G(t) \in \mathbb{R}^{n\times n}$ is a known system matrix, and $\boldsymbol{w}(t)$ is an n-dimensional Brownian motion vector. Following the same procedure as above allows us to write

$$\boldsymbol{x}_{k+1} = A_k\boldsymbol{x}_k + B_k\boldsymbol{u}_k + \boldsymbol{z}_k + G_k\boldsymbol{w}_k \tag{B.89}$$

where G_k is any matrix such that

$$G_k\boldsymbol{w}_k \overset{d}{=} A_k \int_{t_k}^{t_{k+1}} \Phi(\tau, t_k)^{-1} G(\tau)\mathrm{d}\boldsymbol{w}(\tau) \tag{B.90}$$

and $\overset{d}{=}$ denotes equality in distribution. It follows that G_k satisfies

$$G_k G_k^T = A_k \left(\int_{t_k}^{t_{k+1}} \Phi(\tau, t_k)^{-1} G(\tau) G(\tau)^T \Phi(\tau, t_k)^{-T} \mathrm{d}\tau \right) A_k^T \quad . \tag{B.91}$$

As such, G_k can be obtained by taking any matrix square root, such as the Cholesky factorization, of the integral on the right-hand size of Equation B.91.