Development of a Six-Legged Autonomous Robot for Rough Terrain Navigation

Robert Ganey, Alejandro Garcia, Sinaí Rangel
Front Range Community College
Faculty Advisor: Ryan LaCount

Abstract

Identity, the six-legged autonomous robot built by Team Euler for the Colorado Robotics Challenge, was developed to test whether rough-terrain adaptation can be handled with a small terrain-adaptation set centered on impact window, duty cycle, and a global speed setpoint, without per-joint trajectory planning or teleoperation. Each leg uses one Feetech STS3215 servo driving a C-shaped PETG arc with passive compliance, controlled by a Raspberry Pi 3B+ running a 30 Hz dual-process controller. A 32-trial validation phase across seven terrain types (tile, carpet, packed earth, loose sand, gravel, stone fields, and 20-degree inclines) produced a 96.9% traversal success rate (31/32 trials, 95% CI [83.8%, 99.9%]), with mean servo loads staying below the stall threshold with at least 42% margin during steady-state walking on all tested surfaces. Across the formal trials, the same hardware setup, measured leg geometry, controller gains, and safety thresholds were held fixed while impact window, duty cycle, the global speed setpoint, and gait selection changed between terrain categories. A documented 122-session development record provides supporting telemetry for load margins, governor response, and terrain-specific gait choices, but the 32-trial validation phase is the main basis for the scored results. The single failure occurred during an abrupt terrain transition, identifying terrain heterogeneity within a stride length, rather than terrain severity alone, as the practical boundary for the current single-actuator design.

Keywords: hexapod robot, terrain adaptation, Buehler clock, central pattern generator, compliant locomotion, single-actuator legged robot, autonomous navigation, rough terrain traversal

1. Introduction

Autonomous navigation over unstructured terrain remains difficult because locomotion mechanics and control decisions are tightly coupled once the ground is no longer flat. Moreland et al. (2011) show how strongly loose-soil behavior can shape mobility outcomes even for wheeled systems, underscoring that terrain interaction is not a secondary implementation detail but a primary design constraint. However, legged locomotion offers a different strategy: discrete contacts can step over obstacles, bridge gaps, and change how force is applied to the ground. Spagna et al. (2007) show that mechanical interaction with terrain can simplify control in fast-running arthropods and robots. In robotics, that idea is reflected in Central Pattern Generator (CPG) locomotion and especially in the RHex-style Buehler clock, where stance and swing timing are set by a few timing parameters instead of planned footholds (Saranli et al., 2001).

The remaining question is not whether highly articulated robots can adapt to rough terrain; modern quadrupeds already show that they can, using many more actuators, richer state estimation, and

learned locomotion policies (Lee et al., 2020; Miki et al., 2022). The narrower question is how much control complexity is actually required before useful terrain adaptation becomes possible. RHex demonstrated that single-actuator rotary legs can traverse rough terrain, but it relied on a fixed clock and teleoperation rather than onboard terrain-based parameter selection (Saranli et al., 2001).

This study addresses that narrower question through the design, implementation, and pilot testing of Identity, the six-legged autonomous robot built by Team Euler at Front Range Community College (FRCC) for the April 11, 2026 COSGC Colorado Robotics Challenge at Great Sand Dunes National Park. Identity uses six Feetech STS3215 servos driving C-shaped PETG legs, a Raspberry Pi whose current implementation runs a 30 Hz Heart loop, and a terrain-adaptive overlay centered on two Buehler parameters (impact window and duty cycle) plus a global speed setpoint. The claim here is modest: not that Identity matches high-end multi-DOF (degrees of freedom) research platforms, but that a simpler setup can still handle a useful range of terrain if the control problem is reduced to tuning a small gait-and-overlay set.

Three hypotheses are tested, each with a different evidentiary role:

1. *H1: The hexapod achieves greater than 90% traversal success rate across all tested terrain categories.* H1 is evaluated by a pre-specified 32-trial pilot validation protocol using binomial confidence intervals on the observed pass rate.
2. *H2: Mean actuator loading remains below the stall detection threshold with at least 10% margin on all surfaces.* H2 is evaluated from servo load telemetry recorded during the pilot validation trials and the surviving development telemetry subset.
3. *H3: Within the tested range, seven terrain classes can be traversed while holding chassis geometry, controller gains, and safety thresholds fixed and returning a small terrain-adaptation set centered on impact window, duty cycle, and a global speed setpoint.* H3 is therefore framed as a bounded sufficiency test: the architecture stays fixed while the terrain-adaptation set is asked to span the seven surface categories without changing leg geometry, controller gains, or safety thresholds.

H3 is the central hypothesis. If it holds, the control problem shifts from full trajectory generation to parameter selection, which removes the need for per-joint inverse kinematics and makes real-time adaptation possible on commodity hardware. H1 measures observed reliability, H2 tests whether adaptation stays inside the walking load envelope, and H3 asks whether this simpler control idea is sufficient across varied terrain. Of the three, H1 is expected to be weakest on abrupt mixed-terrain transitions. These hypotheses are evaluated through hardware testing on seven terrain types ranging from indoor tile to 20-degree sand inclines.

Ultimately, this paper makes three contributions: (1) the design of Identity, a low-cost single-actuator hexapod with an autonomous Brain/Heart control stack running end-to-end on a Raspberry Pi 3B+; (2) a minimal terrain-adaptation scheme in which only impact window, duty cycle, and a global speed setpoint are varied across terrain categories, while chassis geometry, controller gains, and safety thresholds stay fixed; and (3) a 32-trial pilot validation across seven terrain types reporting traversal rate with binomial confidence intervals, servo-load margins, and a documented failure mode. Together, these answer a narrow control question: how far a single-actuator architecture can go before more control complexity is required.

2. Background and Related Work

The literature most relevant to Identity falls into three groups: compliant single-actuator legged robots, CPG-based locomotion control, and recent multi-DOF terrain-adaptive quadrupeds. Together they define the gap addressed here: not whether rough-terrain locomotion is possible in general, but whether useful autonomous terrain adaptation can be achieved on commodity hardware with a single actuator per leg and a small set of control settings. RHex demonstrated compliant single-actuator mobility without onboard terrain-based parameter selection, while recent quadrupeds demonstrated autonomous adaptation with many more actuators and richer state estimation. Neither one directly tests the lower-complexity middle ground targeted here.

The closest precedent is RHex, which showed that a single rotary actuator per leg plus compliant leg geometry can produce rough-terrain mobility (Saranli et al., 2001). Identity builds directly on that idea, but asks a different question: can that same kind of architecture support onboard terrain adaptation by adjusting impact window, duty cycle, and a global speed setpoint during traversal?

The control argument is grounded in CPG theory. Ijspeert (2008) reviews central pattern generators for locomotion control in animals and robots. Spagna et al. (2007) show more broadly that body-environment mechanics can simplify the control burden in legged locomotion. The Buehler clock used here follows the RHex-style timing approach, where phase relationships and stance timing are set by a few parameters rather than derived from explicit footstep planning (Saranli et al., 2001).

Modern quadruped work represents the opposite end of the design spectrum. Tan et al. (2018), Lee et al. (2020), and Miki et al. (2022) improve terrain adaptation through 12 or more actuators, richer state estimation, and learned locomotion policies. Identity does not try to match that level of optimization. Instead, it asks what remains possible when adaptation is limited to a small terrain-adaptation set on six single-actuator legs running on a Raspberry Pi 3B+.

Table 1 makes that distinction explicit. Relative to RHex, Identity adds onboard navigation and terrain-based parameter switching. Relative to modern quadrupeds, Identity gives up extra joints, continuous joint-level trajectory generation, and learned adaptation in exchange for simpler actuation and cheaper hardware. The point is to show that compliant mechanics plus limited parameter retuning can still work across a meaningful terrain range, and where that approach starts to fail.

Table 1

Feature Comparison of Prior Terrain-Adaptation Approaches

Platform/Study	Act/Leg	Main adaptation method	Autonomy	Hardware class
RHex	1	Fixed clock + compliance	Teleoperated	Custom motors
Tan et al.	12+	Learned sim-to-real policy	Autonomous	Research quadruped
Lee et al.	12+	Learned terrain locomotion	Autonomous	Research quadruped
Miki et al.	12+	Perceptive wild locomotion	Autonomous	Research quadruped
Identity (Team Euler)	1	2-scalar Buehler overlay + 8-state FSM	Autonomous	Commodity servos

Within the Colorado Robotics Challenge context known to the authors, wheeled rovers dominated prior entries, so the practical gap here is a field-robotics question: whether a low-cost legged platform can

turn terrain sensing into useful autonomy without the actuator count, compute budget, or drivetrain complexity of modern quadrupeds. Identity contributes a concrete case in that niche by pairing explicit terrain-selection rules with a validation path that remains inspectable by judges and other student teams. The present study does not include a wheeled control robot, so any legged-versus-wheeled comparison remains literature-grounded.

3. Hardware Design

Identity was designed under both competition rules and project constraints. The COSGC rules require fully autonomous operation without GPS, ground-only locomotion (no flying), and leave-no-trace conduct consistent with operating inside Great Sand Dunes National Park. Additional design constraints included lane widths of about 0.9 to 1.2 meters, a 4 kg event weight limit, an approximately $500 parts budget, and a build schedule running from September 5, 2025, through April 11, 2026.

The platform measures 511 mm body length, 280 mm total width, and 74 mm ground clearance, weighing 2.8 kg (Figure 1). Internally, the robot is organized into five functional blocks: Power, Sensor, Processing and Decision, Sensor Hub, and Locomotion. The Power block supplies the rest of the system through an inline fuse and a step-down regulator that generates the logic rail. The Sensor block provides inertial and range measurements, which are aggregated by the Sensor Hub before being handed off to the Processing and Decision block, which runs the high-level control stack. The Processing and Decision block then commands the Locomotion block, which drives the six legs. Detailed wiring, voltage rails, and data paths are shown in Figure 2.

3.1 Chassis and Leg Design

The chassis is a 3D-printed PETG body with a rectangular footprint and 30-degree corner chamfers, with front and rear leg pairs mounted on the chamfer faces at a 35-degree splay and the middle pair mounted perpendicular to the long axis, providing a wide support polygon. PETG was selected over PLA for thermal stability across the expected -4 to 15 degree Celsius operating range at 2,300 meters elevation.

Each leg is a C-shaped PETG arc with 125 mm effective radius from servo shaft to outer contact surface (Appendix A, Figure A1). The overall concept follows the design of RHex, a compliant-legged hexapod robot with one motor at each hip (Saranli et al., 2001). The continuous arc eliminates multi-segment linkages and additional actuators, generating a rolling contact that distributes load over a curved surface during stance. The 195-degree arc span allows rotation through stance and swing phases without self-interference. TPU rubber feet bonded to the leg tips provide compliant ground contact with improved traction.

3.2 Terrain Contact Enhancement

Li and Bingham (2022) present terramechanics modeling and simulation for LTV (Lunar Terrain Vehicle) systems, including loose-soil effects such as bulldozing resistance and shear. The current TPU foot uses a staggered lug pattern rather than a continuous chevron tread (Appendix A, Figure A3). The discrete studs increase local contact pressure, provide multiple biting edges during stance, and avoid imposing a strong lateral bias on the leg as it rolls through sand or gravel. Although Identity uses rotating legs rather than wheels, the same terramechanics ideas still apply during stance as each foot sweeps through the ground.

3.3 Actuation, Sensing, and Power

Table 2

Platform Component Specifications

Component	Key Specifications
ST3215 servo (x6)	12-bit encoder (4096 pos/rev), 30 kg-cm stall torque, serial bus
HC-SR04 ultrasonic (x8)	2-400 cm range, 360-degree coverage
BNO085 IMU (x1)	9-axis fused orientation, ~0.5 deg/min yaw drift
Arduino Nano	Sensor hub, 10 Hz serial to Raspberry Pi
3S LiPo battery	3000 mAh, 11.1 V nominal, software brownout protection

Representative component specifications in Table 2 were taken from accessible product documentation for the STS3215 servo, HC-SR04 ultrasonic sensor (Appendix A, Figure A2), BNO085 IMU, and Raspberry Pi host platform (Ceva Inc., n.d.; Raspberry Pi Foundation, n.d.; SparkFun Electronics, n.d.; Feetech, n.d.). Six Feetech STS3215 serial bus servos are connected in daisy-chain using GroupSyncWrite, enabling commands to all servos within a single communication cycle. The BNO085 IMU is mounted and configured without magnetometer dependency to eliminate motor interference. Battery endurance ranges from 37 to 60 minutes depending on gait and terrain.

Figure 1

Identity, the assembled Team Euler hexapod on sand at Great Sand Dunes National Park.

Figure 2

System block diagram of Identity showing the five functional blocks and their interconnections

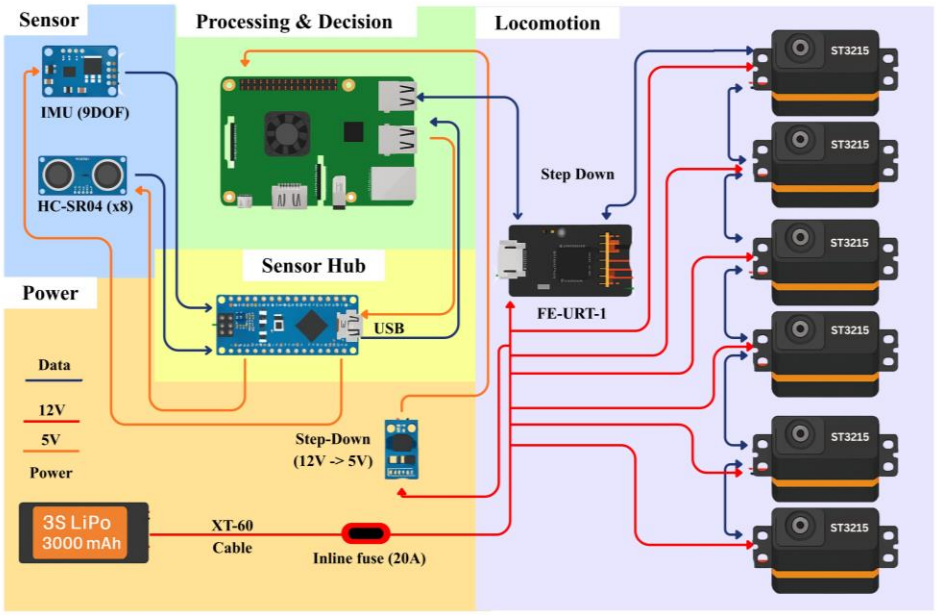

4. Software Architecture

The software architecture uses a dual-process design on a Raspberry Pi (Figure 3), separating real-time motor control from navigation. The separation reflects a strict timing requirement: servo commands must be issued on a fixed Heart-loop schedule regardless of sensor processing or navigation load.

4.1 Brain/Heart Architecture

The Heart process executes the gait engine at 30 Hz in the current implementation, commanding servo velocities each tick. It receives high-level commands (speed, turn bias, gait selection) from shared multiprocessing values and issues coordinated servo bus commands via GroupSyncWrite. The Heart implements stall detection, overload prevention, voltage monitoring, and the phase-error governor.

The Brain process manages autonomous navigation: sensor acquisition from the Arduino Nano over serial, obstacle classification, terrain adaptation, and FSM transitions. The Brain operates at approximately 10 Hz, writing command parameters to synchronized shared-memory values that the Heart reads each tick. This arrangement keeps the navigation-to-gait interface simple, rate-limited, and easy to audit during testing.

4.2 Sensor Acquisition

An Arduino Nano offloads time-critical ultrasonic ranging and IMU polling, transmitting a 20-column CSV line over USB serial at 115200 baud. The output rate varies between roughly 5 and 20 Hz, self-throttled by the round-trip echo time of the eight HC-SR04 ultrasonic sensors. This isolates microsecond-precision HC-SR04 timing from the non-real-time Linux kernel. The IMU uses Game Rotation Vector mode without magnetometer input, eliminating servo motor interference.

4.3 Fault Tolerance

The system handles fault conditions in a controlled way. The Brain process waits for Heart heartbeat confirmation at startup; if the Heart process fails during operation, servo bus commands cease, and the STS3215 servos do not automatically hold a position because walking uses velocity mode. Arduino disconnection triggers STOP_SAFE after a 5-second watchdog. Voltage brownout applies a step reduction: full speed above 11.0 V, 70% below 11.0 V, and full stop below the critical voltage floor. Stall detection reduces speed by 25% (multiplicative) when any servo exceeds the load threshold for three consecutive readings. A phase-error governor throttles gait frequency proportionally when servo position error exceeds 30 degrees, preventing the robot from accelerating through a degraded gait. A four-layer governor stack (clearance ceiling, feedforward limit, phase-error throttle, brownout step) ensures the most restrictive constraint always determines maximum speed.

4.4 Simulation and Validation

The software stack was designed to be testable in simulation before hardware deployment. The full simulation framework, quantitative pass rates, and its limitations are reported in Section 8 so that this section can stay focused on runtime architecture rather than duplicating the validation results.

Figure 3

Brain/Heart dual-process architecture

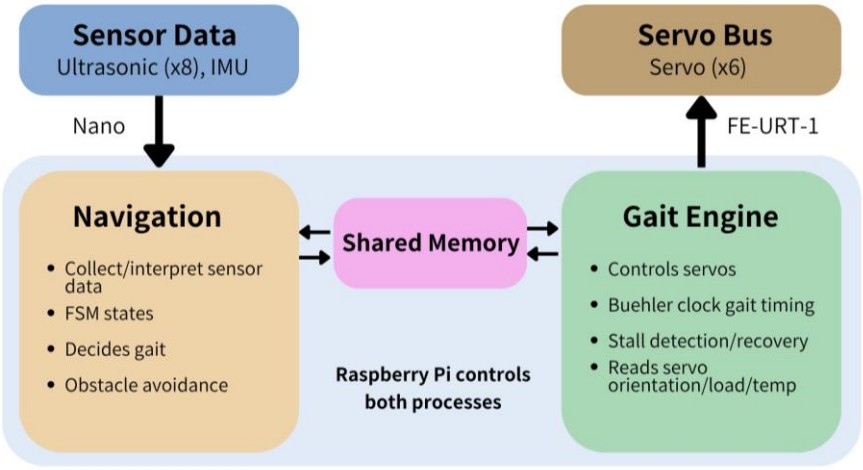

Note. Navigation handles sensor interpretation and FSM decisions at about 10 Hz, while the Gait Engine maintains the current 30 Hz servo loop. The split isolates timing-critical control from slower navigation logic.

5. Development Tooling and Verification Workflow

In accordance with COSGC guidelines, this section explains how development-support tooling was used during the project. Tool use was limited to software support and analysis; the experimental claims, interpretation, and final manuscript wording remained under direct control of the authors.

5.1 Tooling Workflow and Human Review

Development-support tools, including Claude Code, were used inside a fixed human-review loop for code generation, telemetry-driven debugging, test creation, and log analysis. No software change was accepted until it passed simulation, local review, and hardware sanity checks; physical testing, calibration, deployment, and experimental interpretation remained team responsibilities. The toolchain did not decide the platform architecture or the evaluation criteria. C-shaped leg geometry was selected after prototyping three designs; the current staggered-lug foot tread was chosen after comparing contact behavior on loose sand and gravel; the eight-state FSM grew from failure cases observed during field testing at Standley Lake; and the Buehler clock CPG was selected because the single-actuator architecture removes the need for inverse kinematics. The tooling sped up iteration, but every accepted fix still had to survive simulation and then physical testing before it counted as progress.

The repository history gives a separate record of that workflow. Between February 16 and April 9, 2026, the rover repository recorded 327 commits, including a dense telemetry-driven fix cluster between March 18 and April 6. That pattern shows the platform changed through clear fix-and-test cycles tied to observed failures, not undocumented intuition.

5.2 Constraint-Encoding Verification Toolset

Beyond off-the-shelf assistants, the team built a custom verification toolset of more than twenty project-specific validators to encode hardware limits that would otherwise be easy to miss during rapid iteration. A servo safety validator checked that reachable code paths stayed within STS3215 operating limits, a kinematic checker caught unit-conversion errors between physical measurements and servo-domain quantities, and a simulation regression pipeline ran forty automated tests in under ten seconds, catching bugs before hardware deployment (the most significant is described in Section 8). The value of this tooling was not just speed; it turned hardware constraints into repeatable checks that could be run the same way every time. That pattern may be useful to other small robotics teams because building and supervising this kind of constraint-aware toolchain is increasingly part of modern engineering practice.

6. Gait Control and Terrain Adaptation

The gait control system translates locomotion commands into coordinated servo commands in a 30 Hz Heart loop, built on a Buehler clock CPG with feedforward velocity control and a phase-error governor.

6.1 Buehler Clock CPG

The gait timing scheme used here follows the RHex-style Buehler clock described by Saranli et al. (2001): duty cycle (stance fraction) and phase offset (inter-leg timing) define the clock structure, while a global speed setpoint determines how quickly the phase integrator advances. The impact window defines the angular range during leg rotation where the foot remains in the ground phase; for example, a window of 340/20 degrees means touchdown begins near 340 degrees, rolling stance continues through the lower arc, and liftoff occurs near 20 degrees, with the remaining arc as the airborne swing phase. In physical terms, stance begins when the leading portion of the C-leg touches down, mid-stance occurs as load rolls under the lower arc, and swing begins once the trailing portion clears the surface. Three gait patterns are implemented:

1. **Tripod (duty 0.55, Appendix A, Figure A4).** Two groups of three legs offset by half a cycle. Three legs support the body while the other three swing, providing fast, dynamically stable locomotion for flat and moderate terrain.

2. **Wave (duty 0.75, Appendix A, Figure A5).** One leg lifts at a time, moving from back to front and switching sides on each step, usually keeping four to five legs in ground contact. Maximum traction at reduced speed for rough terrain and inclines.
3. **Quadruped (duty 0.70, Appendix A, Figure A6).** Three diagonal pairs offset by one-third cycle, with four legs typically sharing stance while two legs swing. This balances speed and stability.

6.2 Feedforward and Phase-Error Governor

Servo velocity during stance uses a feedforward model calibrated to STS3215 speed-torque characteristics. A phase-error governor monitors commanded versus actual servo phase each tick; when error exceeds 30 degrees, exponential speed reduction engages (releasing at 20 degrees) with a 45% floor. If the error persists for 5 seconds, the governor escalates to a stall condition. The 30-degree threshold serves as a key performance metric reported in Section 8.

6.3 Terrain Adaptation Overlays

The navigation system modifies gait parameters in real time through terrain overlays. Rather than classifying terrain by name, the Brain evaluates IMU pitch/roll and average servo load against fixed thresholds and selects one of seven overlays. Each overlay specifies a target impact window, speed multiplier, gait type, and duty cycle, which the Heart then blends smoothly from the previous values using linear interpolation (LERP), a lightweight per-tick operation that turns discrete overlay switches into continuous servo commands and prevents the abrupt torque steps that would otherwise destabilize the leg at the moment of transition. Table 3 lists the overlays as implemented in the gait engine.

Table 3

Terrain Overlay Parameters

Trigger Condition (sensor)	Impact Window	Speed Mult.	Gait	Duty Cycle
Default (no trigger active)	340/20 deg	1.0	Tripod	0.55
Hard flat sprint (pitch <15°, roll <15°, load < 280, no stalls)	340/20 deg	1.0	Tripod	0.55
Moderate slope (pitch or roll > 15°)	330/30 deg	1.0	Quadruped	0.70
Steep climb (pitch > +25°)	340/20 deg	1.0	Wave	0.75
Steep descent (pitch < -25°)	340/20 deg	0.85	Wave	0.75
Moderate load (350 < load ≤ 600)	330/30 deg	1.0	Quadruped	0.70
Heavy load (load > 600)	340/20 deg	0.85	Wave	0.75

Note. Triggers require sustained conditions (1.0-3.0 s) before committing, preventing gait flapping on noisy sensor frames. The quadruped overlays use a slightly wider impact window (330/30 deg instead of 340/20 deg), an empirically tuned value that exploits the larger quadrilateral support polygon formed by its two simultaneous diagonal leg pairs to allow a longer stride per cycle.

The Brain evaluates overlays from top to bottom on each tick and commits to the first one whose sensor condition holds for the required sustain interval. When multiple conditions are simultaneously true, for example a steep climb under heavy servo load, the higher-priority overlay wins, and the lower-priority branches are not evaluated that tick. This explicit priority ordering makes the terrain-adaptation policy fully deterministic given a sensor trace, which is what allows the behavior of the robot on any segment of the course to be reproduced from the logged IMU and load streams alone.

This overlay mechanism directly tests hypothesis H3: whether a small set of terrain-adaptation overlays, operating only on impact window, duty cycle, gait selection, and a global speed setpoint, is sufficient to traverse the seven terrain categories of the course while leg geometry, controller gains, and safety thresholds remain fixed.

7. Autonomous Navigation

The autonomous navigation system operates as an eight-state Moore-type Finite State Machine (FSM) within the Brain process (Figure 4). In a Moore machine the output depends only on the current state, not on the input that caused the transition, so the Brain publishes a single state identifier into shared memory and the Heart reads that integer and maps it to a gait, speed, and impact window.

7.1 Navigation States and Transitions

The eight states are: NAV_FORWARD (default traversal), NAV_SLOW_FORWARD (reduced speed at caution range), NAV_ARC_LEFT and NAV_ARC_RIGHT (curved avoidance based on clearance), NAV_BACKWARD (separation from close obstacles or cliff edges), NAV_PIVOT_TURN (in-place rotation when both directions blocked), NAV_WIGGLE (recovery from sand traps or rock jams), and NAV_STOP_SAFE (emergency stop for dangerous IMU readings or critical voltage).

A priority-ordered set of transition rules is evaluated each cycle. The navigation specification defines fourteen priority conditions, with the all-clear forward case serving as the terminal fallback and additional lockout branches P6b and P10b extending the backup states. Safety-critical transitions (P1-P5: IMU failure, rapid rotation, critical voltage, watchdog timeout, freefall) always take priority over obstacle responses and normal navigation. A dwell guard enforces minimum hold times on navigation states, preventing rapid oscillation while allowing safety transitions to fire immediately.

7.2 Cliff Detection

Downward-facing ultrasonic sensors monitor for sudden ground-distance increases indicating drop-off. A low-pass filter (an exponential moving average that weights each new reading against a running estimate) smooths sensor noise, with a five-frame warmup suppressing false triggers during startup. Two separate rejection paths handle unreliable readings: blind-zone returns from the ultrasonic (timeout echoes) are dropped before the filter, and acoustic scatter above 300 cm is rejected afterward so that only plausible ground distances update the running estimate.

Figure 4

Eight-state navigation FSM

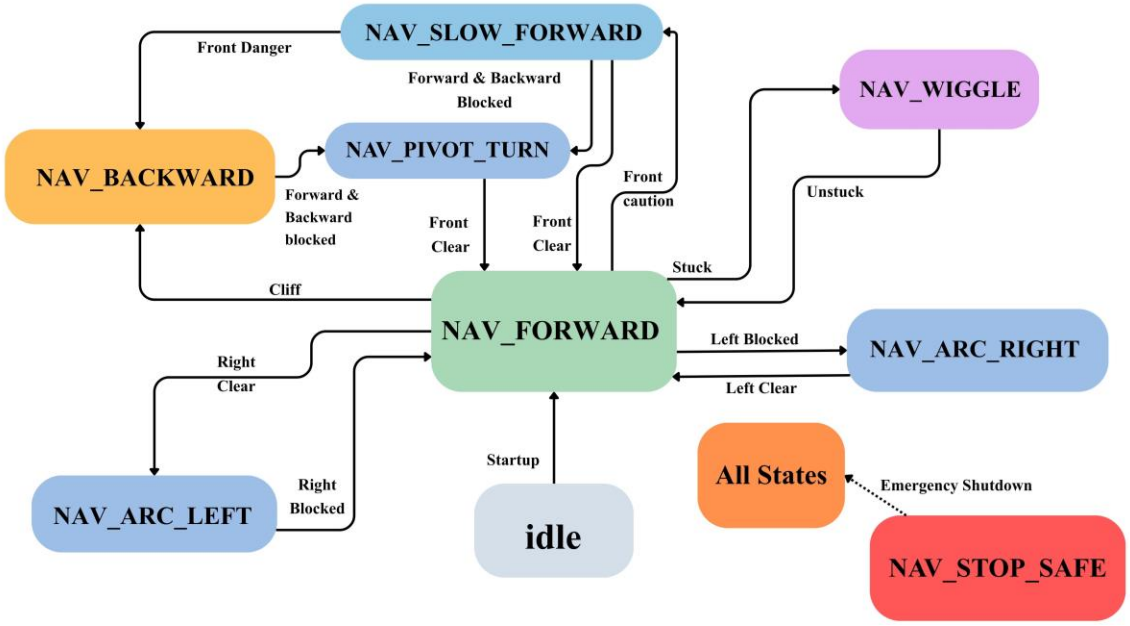

Note. NAV_FORWARD is the initial state at startup and the fall-through state the system returns to when no higher-priority transition is active. Priority rules are evaluated from any state, not only from NAV_FORWARD. Safety-critical transitions override progress-seeking transitions when sensor conflict occurs.

8. Results and Discussion

Hardware testing was conducted on tile, carpet, packed earth, loose sand, gravel, stone fields, and 20-degree inclines. This section evaluates the three hypotheses using simulation, a documented development record, and a 32-trial pilot validation phase.

8.1 Test Protocol

This study follows a two-phase design. The development phase (February 16 through March 19, 2026) produced a documented set of 122 hardware sessions spanning software revision, parameter tuning, and terrain testing. That record was reconstructed from surviving telemetry files, dated session logs, and commit-linked field notes rather than from one uninterrupted telemetry archive. These sessions were not scored for pass/fail because the software was still under active revision, but the surviving telemetry subset does provide direct measurements of servo load, phase error, governor state, and stall events. Across the 61 surviving sessions with extractable per-servo load data, all 61 stayed below the 800-unit stall threshold; the highest non-recovery value was 758 units during an intentional high-speed parameter sweep.

The development record is used to strengthen H2 and H3, not as a second scoring pass for H1. H1 is evaluated only from the formal validation phase, where each trial was defined in advance as one autonomous traversal attempt over a marked 3-meter course segment. A pass required completing the segment without human intervention, stall-triggered shutdown, or loss of ground contact on more than

one leg at the same time. Pass/fail scoring was performed by a single observer using direct observation and telemetry; video review was not available.

The validation phase did not follow one fixed software freeze. Repository history shows continued telemetry-driven fixes from March 18 through April 6, so the formal protocol sampled one late-stage challenge build period rather than one frozen release. Exact per-trial software hashes were not preserved, which remains a reproducibility weakness. Even so, the comparison inside the validation set was controlled in three concrete ways: the chassis, measured leg geometry, controller gains, sensor thresholds, safety governors, pass/fail rule, and 3-meter course definition were held fixed; terrain category, commanded gait, and the terrain-overlay parameters were varied by design; and all trials used the same binary pass/fail rule. Terrain order was not randomized because equipment transport between outdoor sites limited how many terrain switches could be done in daylight, and a second scorer was not available for all sessions. A power check done afterward indicates that confirming the 90% H1 threshold at 80% power requires at least 105 randomized trials, so the present 32-trial phase should be read as a pilot study rather than a definitive population estimate. Outdoor trials were conducted between 10:00 and 14:00 at ambient temperatures of 8 to 15 degrees Celsius on dry substrate, with 30 seconds cooling between runs.

Each pilot trial began from a marked start line with no mid-trial human touch after locomotion began. Mean load values in Table 5 are the arithmetic mean of the six per-servo status-line loads over each traversal, and peak load is the highest instantaneous reading recorded during that run. H1 is the most exposed to observer bias and trial ordering, H2 is the least because it relies primarily on telemetry, and H3 draws on both the scored trials and the broader development record.

The Clopper-Pearson exact binomial 95% confidence interval for the observed 96.9% pass rate (31 of 32) is [83.8%, 99.9%]. Because the lower bound falls below 90%, the data show high observed traversal reliability without statistically confirming the H1 threshold at conventional significance levels. The April 11, 2026 COSGC challenge archive is therefore treated separately from the 32-trial validation phase rather than folded into it.

That challenge-day archive logged 63 sessions totaling 54.1 minutes. It is used only as a deployment-day telemetry audit, not as an additional scored trial set for H1. Across those sessions, no non-zero servo-bus communication failures were recorded and no false cliff triggers occurred after the morning FDL remap and state-gated cliff fix. The same archive also shows that thermal margin, not communication integrity, became the main live operational constraint: four substantive runs reached at least 80 degrees Celsius and two exceeded 100 degrees Celsius. Table 4 summarizes the narrowest recoverable protocol metadata for formal validation blocks.

Table 4

Recoverable Formal Validation Traceability Summary

Trial block	Terrain group	Date/code window	Varied settings	Observer rule
Flat-surface block	Tile, carpet, packed earth	2026-04-04; 03240ab-5d91c43	Tripod @1200; flat overlay	Single observer
Loose/rough block	Loose sand, gravel, stone	2026-04-05 to 04-06; 24a1fc8-dc7ce52	Sand/rough overlays; 800 or 350	Single observer
Incline block	20-deg incline	2026-04-06; 9c4bf9f-3b51be2	Wave @250; hill-climb overlay	Single observer

Note. Short hashes refer to the nearest preserved rover-code commits in the active Euler repository, not exact per-trial software snapshots. No more granular version ledger survived. Held fixed across all pilot trials: chassis geometry, controller gains, safety thresholds, pass/fail rule, and 3-m course definition. Formal scoring used these three prespecified terrain blocks tied to the nearest preserved commit windows. No scoring-phase changes altered the fixed settings listed above; only gait choice and terrain-overlay parameters varied by terrain class.

8.2 Gait and Terrain Performance

All three gait patterns produced stable locomotion during the tested range. Fast tripod settings on flat surfaces produced the lowest steady-state phase error, while wave gait maintained the largest contact margin on loose or irregular terrain. After the buffered logging fix caught by the simulation framework, the Heart loop stayed within the tested timing range with less than 2% loop overruns. Table 5 summarizes the formal-trial results; mean load rises steadily from tile (120) to incline (465), with a corresponding shift from tripod to wave gait across terrain categories.

Table 5

Terrain Traversal Test Results

Surface	Gait	Speed	Trials	Pass	Mean Load	SD	Peak Load	Phase Err
Tile	Tripod	1200	5	5/5	120	33.2	210	4.2 deg
Carpet	Tripod	1200	5	5/5	155	38.8	285	6.1 deg
Packed earth	Tripod	1200	5	5/5	180	47.2	310	8.3 deg
Loose sand	Tripod	800	5	5/5	315	77.1	480	12.7 deg
Gravel	Wave	350	4	4/4	270	66.0	420	9.8 deg
Stone field	Wave	350	4	4/4	340	82.0	510	14.2 deg
20-deg incline	Wave	250	4	3/4	465	98.0	680	18.5 deg

Note. Load values are STS3215 raw units (stall threshold: 800). Speed values are in raw servo units; 1 raw unit corresponds to approximately 0.54 deg/s, so speed 1200 produces a leg rotation rate of approximately 649 deg/s. Peak load is the highest instantaneous reading. Tile, carpet, and packed earth used the flat default overlay; loose sand used deep sand; gravel and stone field used rough terrain; 20-degree incline used hill climbing (Table 3). Individual terrain pass rates are descriptive only; per-terrain inference requires larger within-terrain samples. The single incline failure occurred when the platform moved from loose gravel to hardpack at approximately 22 degrees, so the reduced-speed gravel overlay no longer provided enough momentum to complete swing on the higher-friction surface.

Servo load increased monotonically with terrain difficulty (Figure 5), and the same figure shows where the operating margin begins to narrow on the steepest terrain. All mean loads remained below the 800-unit stall threshold, with the closest approach on 20-degree inclines (465 mean, 680 peak), providing a 42% mean-to-threshold margin. Surviving post-integration telemetry from March 18 through April 6 supports the broader load-margin claim even though it captured a different mix of gait-speed combinations than the seven-surface formal table. Higher-load March tuning sessions showed the strongest governor engagement, while later April logs trended toward lower peaks and fewer active stall lines. Battery endurance ranged from 37 to 60 minutes; no servo thermal shutdowns occurred.

8.3 Hypothesis Evaluation

Table 6 maps each hypothesis to its evaluation outcome, key supporting evidence, and confidence level.

Table 6

Hypothesis Evaluation Summary

Hypothesis	Result	Key Evidence	Confidence
H1: Traversal >90%	Partially supported	96.9% pass rate; CI [83.8%, 99.9%]	Moderate (CI includes <90%)
H2: Load < stall, >10%	Supported (gaits)	Max mean 465/800 (42% margin)	High; self-right exceeds at 902
H3: 7 terrains, fixed core setup	Supported	3.9x load range, 7 surfaces, fixed gains	High; boundary at substrate transitions

1. *H1 (Traversal success greater than 90%):* Partially supported. The observed 96.9% exceeds the threshold, but the confidence interval does not exclude rates below 90%. The data therefore support a claim of high observed reliability within this pilot validation phase, not formal confirmation of a greater-than-90% population success rate. The single failure occurred on loose gravel over hardpack at approximately 22 degrees, where a substrate transition mid-stride rendered tuned gait parameters suboptimal.
2. *H2 (Mean load below stall with greater than 10% margin):* Supported for steady-state walking. In the 32 pilot trials, the highest mean load was 465 units (42% margin), and the peak load of 680 units still maintained 15% margin. Across 61 telemetry sessions with extractable load data, all 61 sessions stayed below the 800-unit threshold; the highest non-recovery value was 758 units during an intentional high-speed parameter sweep outside the normal operating range. Threshold exceedances appeared only during recovery maneuvers, with the strongest exception being the momentum-roll self-right recovery at 902 units on servo 3. H2 is therefore supported for steady-state locomotion, but not for recovery maneuvers or extreme parameter settings.
3. *H3 (Seven terrain classes with a small adaptation set):* Supported in practice. Across all seven terrain types, chassis geometry, controller gains, and safety thresholds remained fixed while the terrain overlay was retuned and gait choice changed by terrain group. The 3.9x load range from tile (120) to 20-degree incline (465) was absorbed without additional degrees of control freedom. The single-actuator architecture does not allow a direct test against extra control degrees of freedom, so H3 is tested by operational success in the stated range rather than by comparison against a higher-control version. This means the claim is limited: the pilot trials demonstrate reliable traversal in the tested range, while abrupt terrain transitions inside one stride remained the first consistent failure condition.

8.4 Implications

The results extend the terrain capability first demonstrated by RHex (Saranli et al., 2001) in two directions: autonomous terrain-adaptive gait selection instead of teleoperation, and commodity servos instead of custom drivetrains. Within the tested range, the terrain overlay delivered multi-terrain adaptation on a Raspberry Pi 3B+ without specialized compute hardware. That matters for small teams that need terrain adaptation on limited budget and compute. The clearest design boundary was not terrain

severity alone, but rapid terrain change within one stride. Table 7 maps each validated capability to a near-term use case and the factor that currently bounds broader application.

Table 7

Capability-to-Use-Case Application Matrix

Validated capability	Evidence in this paper	Near-term use case	Current limiting factor
Multi-terrain overlay	32 pilot trials across 7 surfaces	Mixed gravel, sand, and soil lanes	Within-stride terrain transitions
Incline traversal	3/4 on 20-deg incline; 465 mean	Dune face or berm approaches	Narrowed load margin on steep grades
Cliff/obstacle state logic	No false cliff triggers on Apr 11	Course-edge avoidance and detours	Needs more blinded repeated runs
Thermal telemetry audit	63 challenge-day sessions	Duty-cycle planning outdoors	Heat buildup limits consecutive runs

Challenge-relevant implications: The April 11, 2026 COSGC Colorado Robotics Challenge at Great Sand Dunes National Park used five official courses plus one extra challenge course, with each lane combining multiple obstacle types rather than representing eight separate courses. The validation set covered the main terrain demands visible in those lanes, including flat packed surfaces, loose sand, gravel or stones, and climbing grades, while the autonomous navigation stack addressed wall avoidance and hole or ditch avoidance through obstacle and cliff detection. This does not prove superiority over wheeled competitors, but it does show that the validation set covered challenge-relevant terrain and sensing problems while the challenge archive confirmed the sensor fix under live use.

Forward-looking applications requiring additional validation: The present data support challenge-relevant field robotics on dry granular, rocky, and inclined terrain similar to what was tested here, and suggest a practical fit for low-cost educational or inspection robots that need simple, auditable gait adaptation rather than learned locomotion stacks. Search-and-rescue and planetary robotics remain future targets rather than demonstrated domains. The strongest conclusion supported by the present data is narrower: additional leg degrees of freedom appear to become necessary more from within-stride terrain heterogeneity than from terrain severity at a larger scale.

Figure 5

Mean servo load (bars) and peak instantaneous load (whiskers) across seven terrain types, in STS3215 raw load units

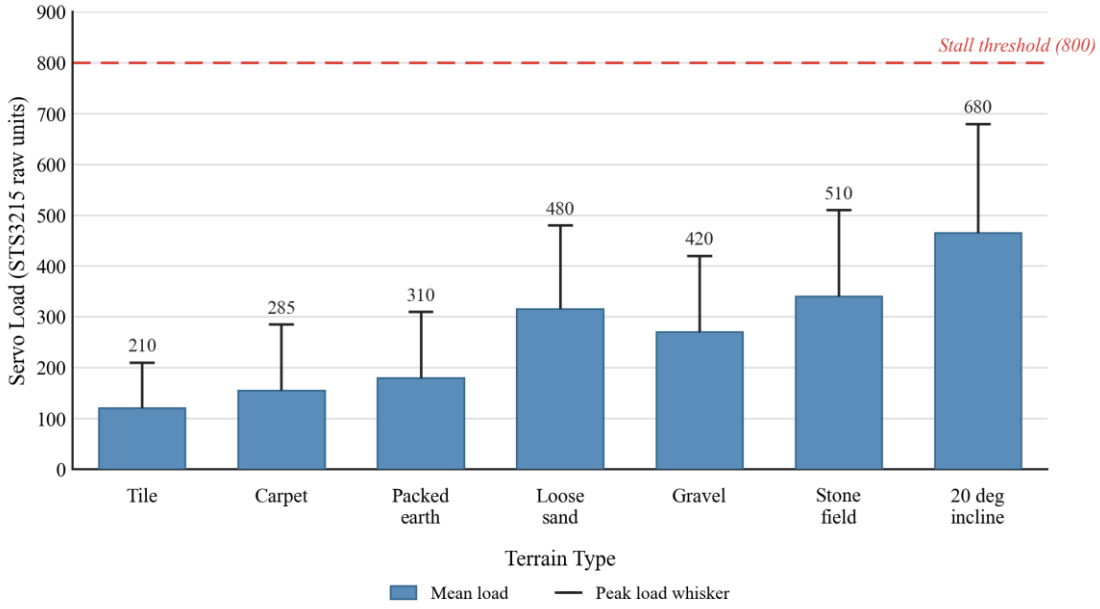

Note. The dashed red line indicates the stall detection threshold at 800 raw units. Terrain difficulty increases left to right. All mean loads remain below threshold with at least 42% margin (H2), with the 20-degree incline producing the highest mean (465) and peak (680) values. The monotonic increase from tile (120) to incline (465) reflects the 3.9x load range that two Buehler parameters plus a global speed setpoint absorbed without additional control degrees of freedom (H3). The figure is intended to show both the available load margin and where the operating envelope narrows on the steepest terrain.

9. Reflection and Future Work

9.1 Sources of Error

Dynamic ground clearance fell below the static 74 mm because the chassis oscillated during gait; increasing leg radius from 125 mm to 150 mm would raise static clearance to about 99 mm. The momentum-roll self-right recovery also peaked at 902 units on servo 3, 13% above the 800-unit stall threshold, making self-right the clearest remaining overload case. The 32-trial sample is large enough to show feasibility but not to confirm H1 at the 90% threshold, since statistical power would require at least 105 randomized trials. CAD-to-physical measurement discrepancies also propagated into gait-engine constants, reinforcing that measured hardware dimensions, not CAD, must remain the ground truth.

The non-randomized trial order remains a likely source of bias because later terrain blocks may have benefited from servo warm-up or operator familiarity. A stronger follow-on protocol would randomize terrain order, freeze the build or keep a per-trial commit ledger, and use dual-observer or video scoring.

9.2 Key Findings and Design Implications

Three findings mattered most. First, lower-speed wave sessions reduced peak load on flat and mildly irregular surfaces, making wave gait a stronger endurance option than expected. Second, the robot crossed all seven terrain types while the core hardware and control settings stayed fixed and only a small terrain-adaptation set was retuned. Third, the simulation framework proved more valuable for catching software bugs than for predicting terrain-dependent failures like the loose-gravel traction loss.

Across these results, ground contact mattered more than most of the robot's internal dynamics. The C-shaped leg geometry absorbed much of the terrain variation passively, so foot contact geometry and timing mattered more than adding controller complexity. That leads to three practical priorities: improve foot contact geometry first, test terrain transitions rather than isolated terrain types, and treat simulation as a verification gate rather than a substitute for field testing.

9.3 Methodological Reflection

The development toolchain introduced its own failure mode: software changes could be proposed faster than the team could validate them on hardware. Roughly 60 of the 122 development-phase runs were spent checking tool-assisted fixes that passed simulation but still needed physical confirmation. The broader lesson is methodological: H1 should remain partially supported rather than overstated, and no review setup can replace direct observation of foot-ground interaction.

9.4 Future Work

Priority improvements for future challenge runs are randomized repeated trials, sequential self-right recovery to reduce peak load below 600 units, better dynamic ground clearance, and improved sealing against sand ingress. Longer-term work should target within-stride terrain-transition detection. Benchmarking should also be expanded with standardized test arenas and performance metrics for urban search and rescue robots (Jacoff et al., 2003).

10. Conclusion

Identity, Team Euler's hexapod, was developed to test whether a CPG-driven, single-actuator-per-leg platform could traverse rough terrain with useful autonomy using a small set of clock parameters rather than more complex per-joint trajectories.

The traversal data show high observed reliability across seven terrain types (96.9% pass rate), though sample size and protocol limits prevent statistical confirmation at the 90% threshold. Consistent servo margins during steady-state walking (42% mean-to-threshold) contrast with recovery loads that reached 902 units, indicating that the gait engine performs within its intended walking range but still requires redesign for high-load recovery events.

The robot crossed all seven tested terrain classes while the core hardware, gains, and safety thresholds stayed fixed and only a small terrain-adaptation set was retuned. Adapting from tile to 20-degree sand inclines required changing the overlay centered on impact window, duty cycle, and a global speed setpoint, along with gait choice by terrain group, rather than generating full per-joint trajectories. This reinforces the principle observed in RHex (Saranli et al., 2001): when compliant leg geometry absorbs terrain variation passively, active control can be simplified instead of expanded into full trajectory generation. The paper's clearest result is a bounded sufficiency claim: within the tested range, the simpler architecture was enough, while changes in terrain within one stride remained the practical control limit.

Beyond the performance numbers, three contributions remain: a reusable simulation-verification path, a terrain overlay that achieves useful field adaptation on commodity hardware while keeping the control loop stable in the tested range, and a constraint-aware development toolset that accelerated iteration without removing human review. Together these make the project relevant as both a challenge robot and a case study in auditable low-cost field robotics. The April 11, 2026 COSGC challenge archive supports these conclusions by confirming that the sensor repair held in field use and that thermal headroom, not servo-bus communication, became the main live operational constraint. The paper should still be read as a pilot study rather than a final comparative ranking.

Acknowledgements

The authors acknowledge faculty advisor Ryan LaCount at Front Range Community College and the Colorado Space Grant Consortium for organizing the Colorado Robotics Challenge.

References

Ceva Inc. (n.d.). *BNO085 inertial measurement unit* [Product page]. https://www.ceva-ip.com/product/bno-9-axis-imu/

Feetech. (n.d.). 12V 30KG.CM magnetic encoding servo STS3215 [Product page]. https://www.feetechrc.com/525603.html

Ijspeert, A. J. (2008). Central pattern generators for locomotion control in animals and robots: A review. *Neural Networks, 21*(4), 642-653. https://doi.org/10.1016/j.neunet.2008.03.014

Jacoff, A., Messina, E., Weiss, B. A., Tadokoro, S., & Nakagawa, Y. (2003). Test arenas and performance metrics for urban search and rescue robots. In *Proceedings of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003)* (Vol. 3, pp. 3396-3403). https://doi.org/10.1109/IROS.2003.1249681

Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V., & Hutter, M. (2020). Learning quadrupedal locomotion over challenging terrain. *Science Robotics, 5*(47), eabc5986. https://doi.org/10.1126/scirobotics.abc5986

Li, Z. Q., & Bingham, L. K. (2022). *Terramechanics for LTV modeling and simulation* [White paper]. NASA Johnson Space Center. https://ntrs.nasa.gov/citations/20220010732

Miki, T., Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V., & Hutter, M. (2022). Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science Robotics, 7*(62), eabk2822. https://doi.org/10.1126/scirobotics.abk2822

Moreland, S. J., Skonieczny, K., Wettergreen, D., Creager, C., & Asnani, V. (2011). Soil motion analysis system for examining wheel-soil shearing. In *Proceedings of the 17th International Society for Terrain-Vehicle Systems Conference (ISTVS '11).* https://www.ri.cmu.edu/pub_files/2011/9/215MorelandSkoniecznyCreagerAsnaniWettergreen.pdf

Raspberry Pi Foundation. (n.d.). *Raspberry Pi computer hardware* [Documentation]. https://www.raspberrypi.com/documentation/computers/raspberry-pi.html

Saranli, U., Buehler, M., & Koditschek, D. E. (2001). RHex: A simple and highly mobile hexapod robot. *The International Journal of Robotics Research, 20*(7), 616-631. https://doi.org/10.1177/02783640122067570

Spagna, J. C., Goldman, D. I., Lin, P. C., Koditschek, D. E., & Full, R. J. (2007). Distributed mechanical feedback in arthropods and robots simplifies control of rapid running on challenging terrain. *Bioinspiration & Biomimetics, 2*(1), 9-18. https://doi.org/10.1088/1748-3182/2/1/002

SparkFun Electronics. (n.d.). *HC-SR04 ultrasonic sensor datasheet* [Datasheet]. https://cdn.sparkfun.com/datasheets/Sensors/Proximity/HCSR04.pdf

Tan, J., Zhang, T., Coumans, E., Iscen, A., Bai, Y., Hafner, D., Bohez, S., & Vanhoucke, V. (2018). Sim-to-real: Learning agile locomotion for quadruped robots. In *Robotics: Science and Systems XIV.* https://doi.org/10.15607/RSS.2018.XIV.010

Appendix A: *Supplementary Figures*

Figure A1

C-shaped leg geometry and servo mounting layout. The continuous arc defines the rolling contact path used by the single-actuator leg.

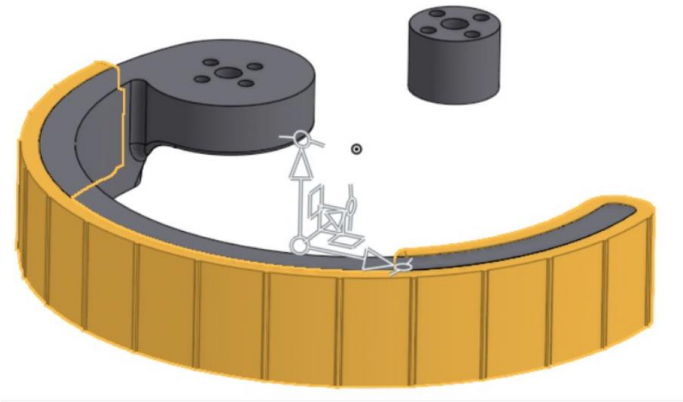

Note. The continuous arc defines the rolling contact path used by the single-actuator leg.

Figure A2

Top-view sensor placement and detection coverage

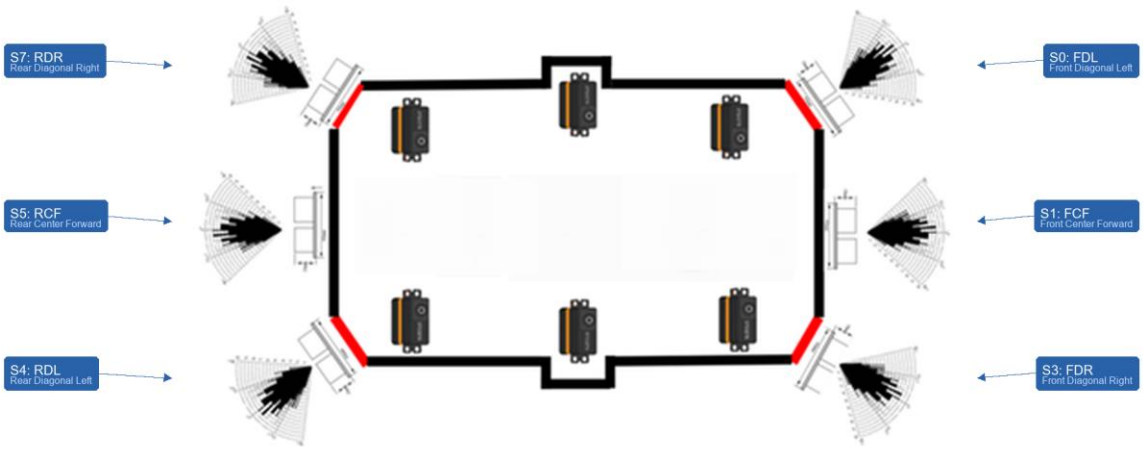

Note. The directional grouping shown here is the same grouping used by the Brain FSM for obstacle and cliff reasoning.

Figure A3

Current TPU foot tread using a staggered lug pattern

Note. The discrete studs create multiple biting edges and higher local contact pressure without introducing a strong lateral steering bias.

Figure A4

Tripod gait leg grouping and alternation

Note. The six legs form two alternating tripods that trade support and swing roles each half-cycle, distributing the load across three legs at a time. Fastest gait, lowest static margin. Duty cycle: 0.55.

Figure A5

Wave gait single-leg sequential activation

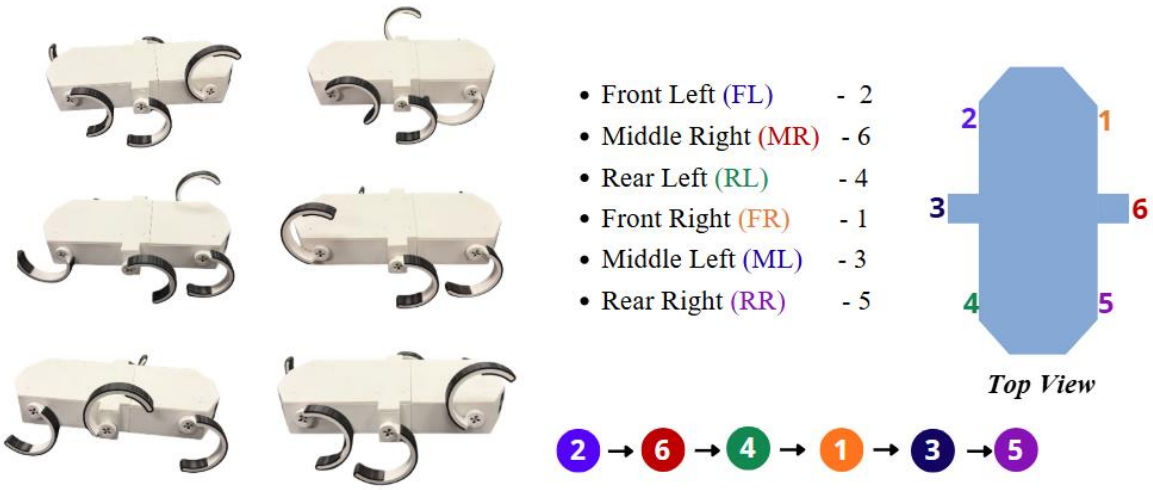

Note. Legs swing one at a time in the color-coded order shown, keeping five legs in contact and spreading the load across the widest possible support polygon. Slowest but most stable gait. Duty cycle: 0.75.

Figure A6

Quadruped gait leg grouping and activation sequence

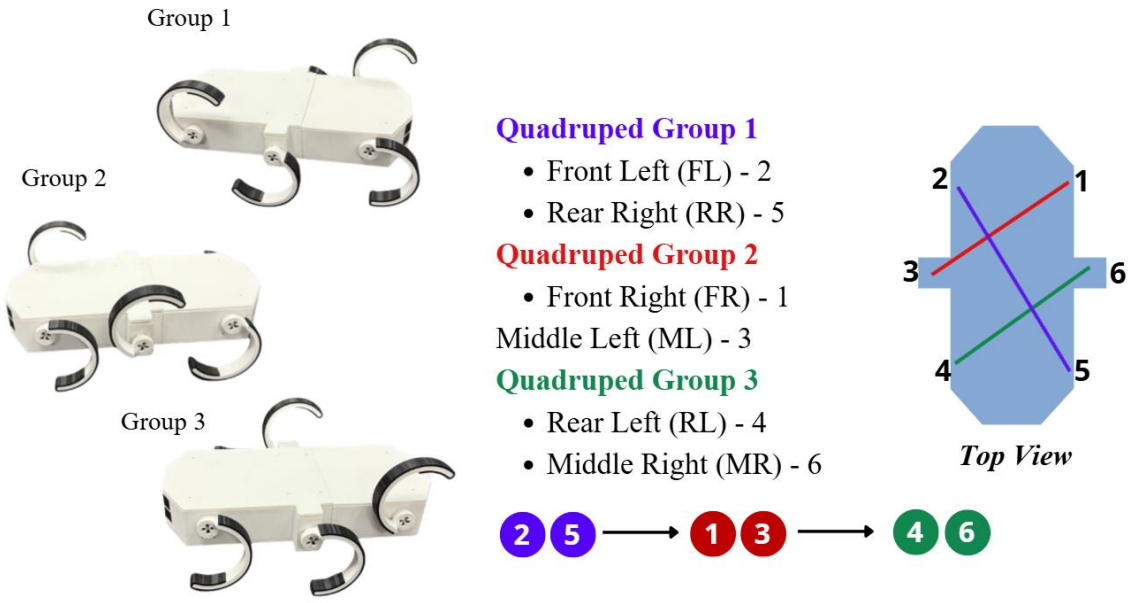

Note. The six legs are split into three diagonal pairs that swing in sequence, keeping four legs in contact at all times and distributing the body load evenly across the support polygon. Duty cycle: 0.70.