Scalable Patent Search and Analysis Using LLMs and Function Calling

RMACC Symposium 2026

ASU Research Computing
Arizona State University

Juanjo García Mesa
Research Software Engineer
Arizona State University

RMACC
Rocky Mountain Advanced Computing Consortium

Outline

Motivation

The Data

Retrieval-Augmented Generation

Building the Database

Query Workflow

Demo

Collaborators

Jordan Bell

Gil Speyer

Johnathan Lee

Roger Kohler

Faraj Tessilimi

Motivation

SI manages 4,000 ASU granted patents and applications

Skysong Innovations

Advancing the Arizona State University Knowledge Enterprise

SI manages 4,000 ASU granted patents and applications

Manual classification and queries are slow

Skysong Innovations

Advancing the Arizona State University Knowledge Enterprise

I have an investor who invests in the sustainability space. The investor wants to see all ASU technologies that fall into the sustainability space. Please produce a list that includes patents and patent applications by number, title, and includes a one-paragraph summary

I have ar
sustain
to see a
the sus
a list th
applica
include

List all the granted patents and patent applications by inventor Erica Forzani

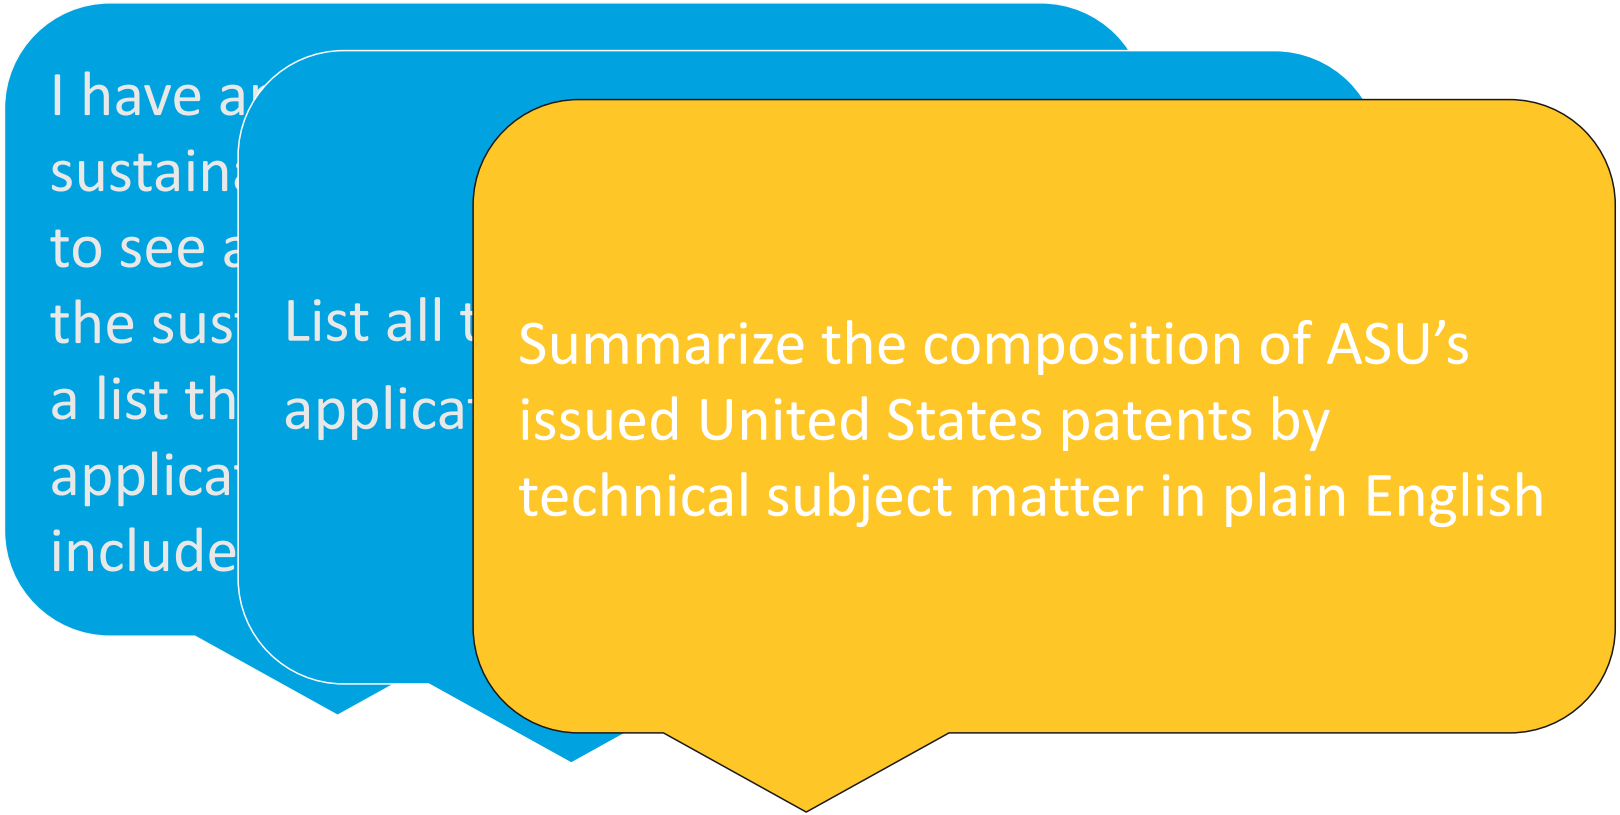

SI manages 4,000 ASU granted patents and applications

Manual classification and queries are slow

Request: Accurate, fast, natural-language-driven tool

Skysong Innovations

Advancing the Arizona State University Knowledge Enterprise

SI manages 4,000 ASU granted patents and applications

Manual classification and queries are slow

Request: Accurate, fast, natural-language-driven tool

“Of course, it’s definitely possible”

Skysong Innovations

Advancing the Arizona State University Knowledge Enterprise

The Data

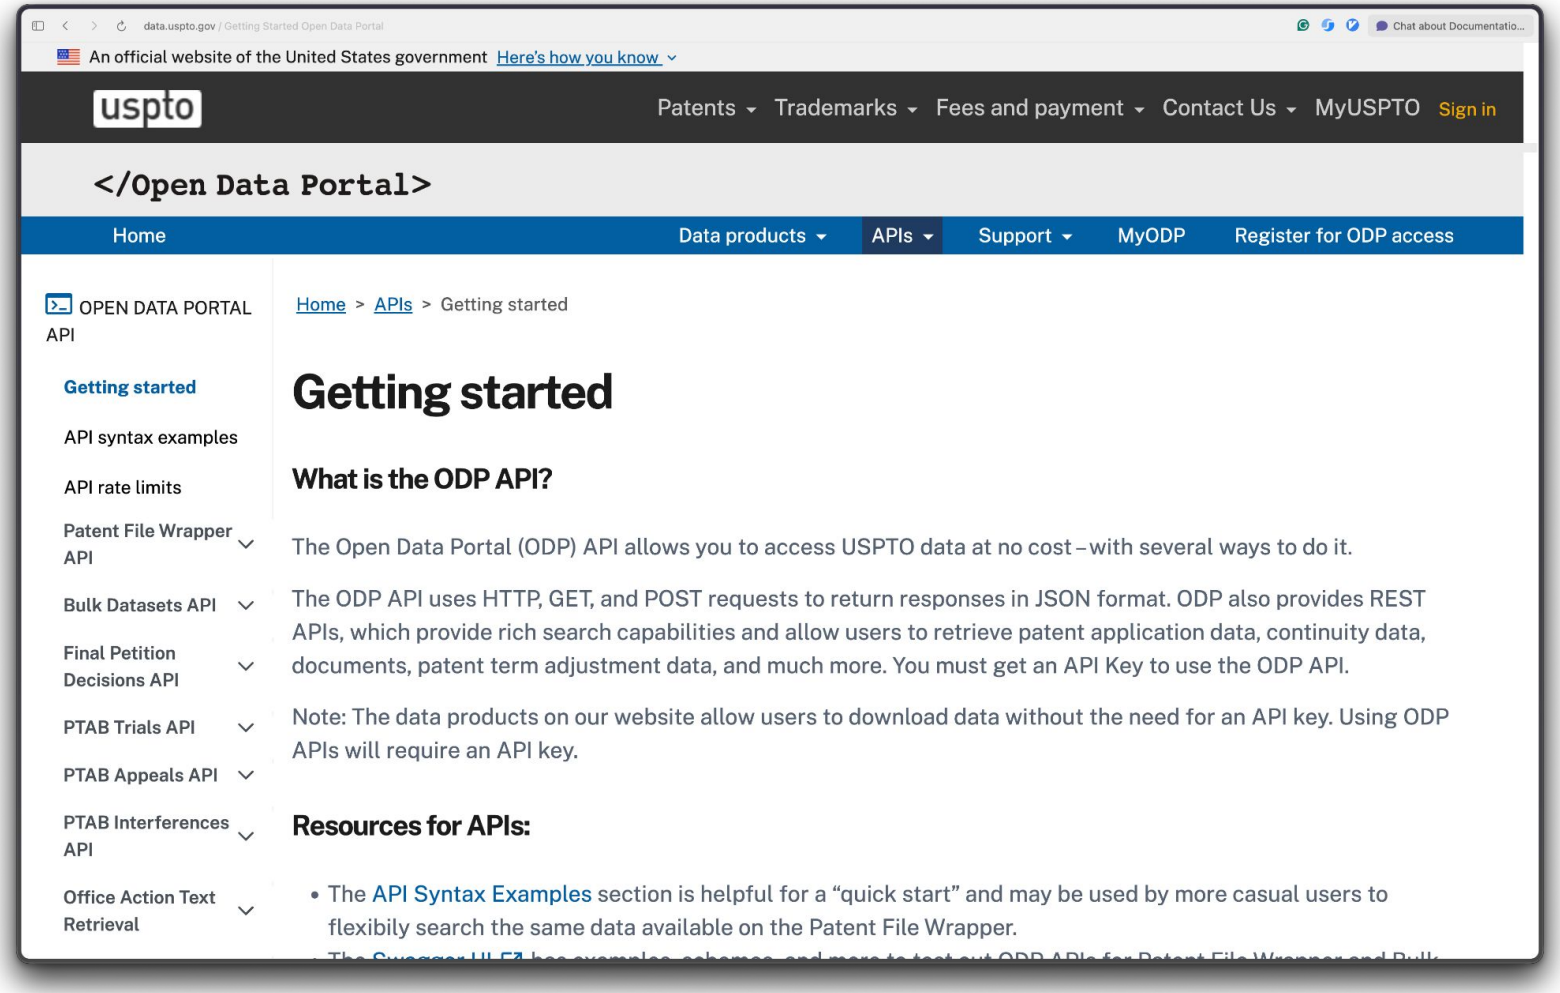

ASU’s Patent Assignees

"Arizona Board of Regents on Behalf of Arizona State University"
"ARIZONA BOARD OF REGENTS ON BEHALF OF ARIZONA STATE UNIVERSITY"
"Arizona Board of Regents on behalf of Arizona State University"

"[The] Arizona Board of Regents"

"Arizone State", "Ariz. State", "ASU", "Arlzona State University", "Arizona Board of Regenta"

Patent Application Full-Text Data (No Images)

Product Identifier: APPXML

Provides the bulk zip files that contain the concatenated full-text of each patent application XML document (non-provisional utility and plant). This page provides an additional feature called “View Patent Records” which allows user to find or discover the patent applications that are bundled in the zip file. However, in this zip file, there may be applications that are not available in ODP and will not be included in the search results. Also, if there are revisions to the applications, the file will only show the original application.

Dates available: March 15, 2001 - July 17, 2025 | **Last Updated:**

Start Date: 01-01-2002 End Date: 07-19-2025 Filter Reset

API Query

January 1, 2002 - July 19, 2025

Showing 1 to 25 of 1266 Results

File Name	Release Date
ipa250717.zip	2025-07-17 / 3:59:45 AM EST
ipa250710.zip	2025-07-10 / 3:59:43 AM EST
ipa250703.zip	2025-07-03 / 4:00:14 AM EST
ipa250626.zip	2025-06-26 / 4:00:03 AM EST
ipa250619.zip	2025-06-19 / 3:01:58 AM EST

Patent Grant Full-Text Data (No Images) - XML

Product Identifier: PTGRXML

Provides the bulk zip files that contains the concatenated full-text of each patent grant document issued weekly. This page provides an additional feature called “View Patent Records” which allows user to find or discover the patent grants that are bundled in the zip file. Even though the zip file may contain grants that belong to the applications which were filed before 2001, ODP will only show the grants that belong to the applications that were filed from 2001.

Dates available: January 1, 2002 - July 15, 2025 | **Last Updated:** 2025-07-15 (Weekly - Tuesday)

Start Date: 01-01-2002 End Date: 07-19-2025 Filter Reset

API Query

January 1, 2002 - July 19, 2025

Showing 1 to 25 of 1238 Results

File Name	Release Date	File Size (Bytes)	Actions
ipg250715.zip	2025-07-15 / 3:57:53 AM EST	171,350,746	View Patent Records
ipg250708.zip	2025-07-08 / 3:57:53 AM EST	169,543,976	View Patent Records
ipg250701.zip	2025-07-01 / 3:57:52 AM EST	158,272,951	View Patent Records
ipg250624.zip	2025-06-24 / 3:01:41 AM EST	138,834,644	View Patent Records
ipg250617.zip	2025-06-17 / 3:01:42 AM EST	173,536,353	View Patent Records

Documents and Resources

us-patent-grant-v47-2022-02-17.dtd
us-sequence-listing-v1_3-2020-10-08.dtd
us-patent-grant-v46-2021-08-30.dtd
us-patent-grant-v45-2014-04-03.dtd
us-patent-grant-v44-2013-05-16.dtd
us-patent-grant-v43-2012-12-04.dtd
us-patent-grant-v42-2006-08-23.dtd
us-patent-grant-v40-2004-12-02.dtd
us-patent-grant-v41-2005-08-25.dtd

Vaults SFTP Sol login

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE PATDOC SYSTEM "ST32-US-Grant-025xml.dtd" [
<!ENTITY USD0452599-20020101-D00000.TIF SYSTEM "USD0452599-20020101-D00000.TIF" NDATA TIF>
<!ENTITY USD0452599-20020101-D00001.TIF SYSTEM "USD0452599-20020101-D00001.TIF" NDATA TIF>
<!ENTITY USD0452599-20020101-D00002.TIF SYSTEM "USD0452599-20020101-D00002.TIF" NDATA TIF>
]>
<PATDOC DTD="2.5" STATUS="Build 20020101">
<SDOBI>
<B100>
<B110><DNUM><PDAT>D0452599</PDAT></DNUM></B110
<B130><PDAT>S1</PDAT></B130>
<B140><DATE><PDAT>20020101</PDAT></DATE></B140
<B190><PDAT>US</PDAT></B190>
</B100>
<B200>
<B210><DNUM><PDAT>29129833</PDAT></DNUM></B210
<B211US><PDAT>29</PDAT></B211US>
<B220><DATE><PDAT>20000922</PDAT></DATE></B220
</B200>
<B300>
<B310><DNUM><PDAT>400 02 946</PDAT></DNUM></B3
<B320><DATE><PDAT>20000322</PDAT></DATE></B320
<B330><CTRY><PDAT>DE</PDAT></CTRY></B330>
</B300>
<B400>
<B472>
<B474><PDAT>14</PDAT></B474>
</B472>
</B400>
```

bulk_ptgrxml_unzip/pg020101.xml

Vaults SFTP Sol login

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE us-patent-grant SYSTEM "us-patent-grant-v40-2004-12-02.dtd" [ ]>
<us-patent-grant lang="EN" dtd-version="v40 2004-12-02" file="USD0500396-20050104.XML" status="SAMPLE-D
ATA" id="us-patent-grant" country="US" date-produced="20041221" date-publ="20050104">
<us-bibliographic-data-grant>
<publication-reference>
<document-id>
<country>US</country>
<doc-number>D0500396</doc-number>
<kind>S1</kind>
<date>20050104</date>
</document-id>
</publication-reference>
<application-reference appl-type="design">
<document-id>
<country>US</country>
<doc-number>29198423</doc-number>
<date>20040129</date>
</document-id>
</application-reference>
<us-application-series-code>29</us-application-series-code>
<priority-claims>
<priority-claim sequence="01" kind="national">
<country>JP</country>
<doc-number>2003-028833</doc-number>
<date>20031001</date>
</priority-claim>
</priority-claims>
<us-term-of-grant>
```

bulk_ptgrxml_unzip/ipg050104.xml

Data Parsing and Formatting

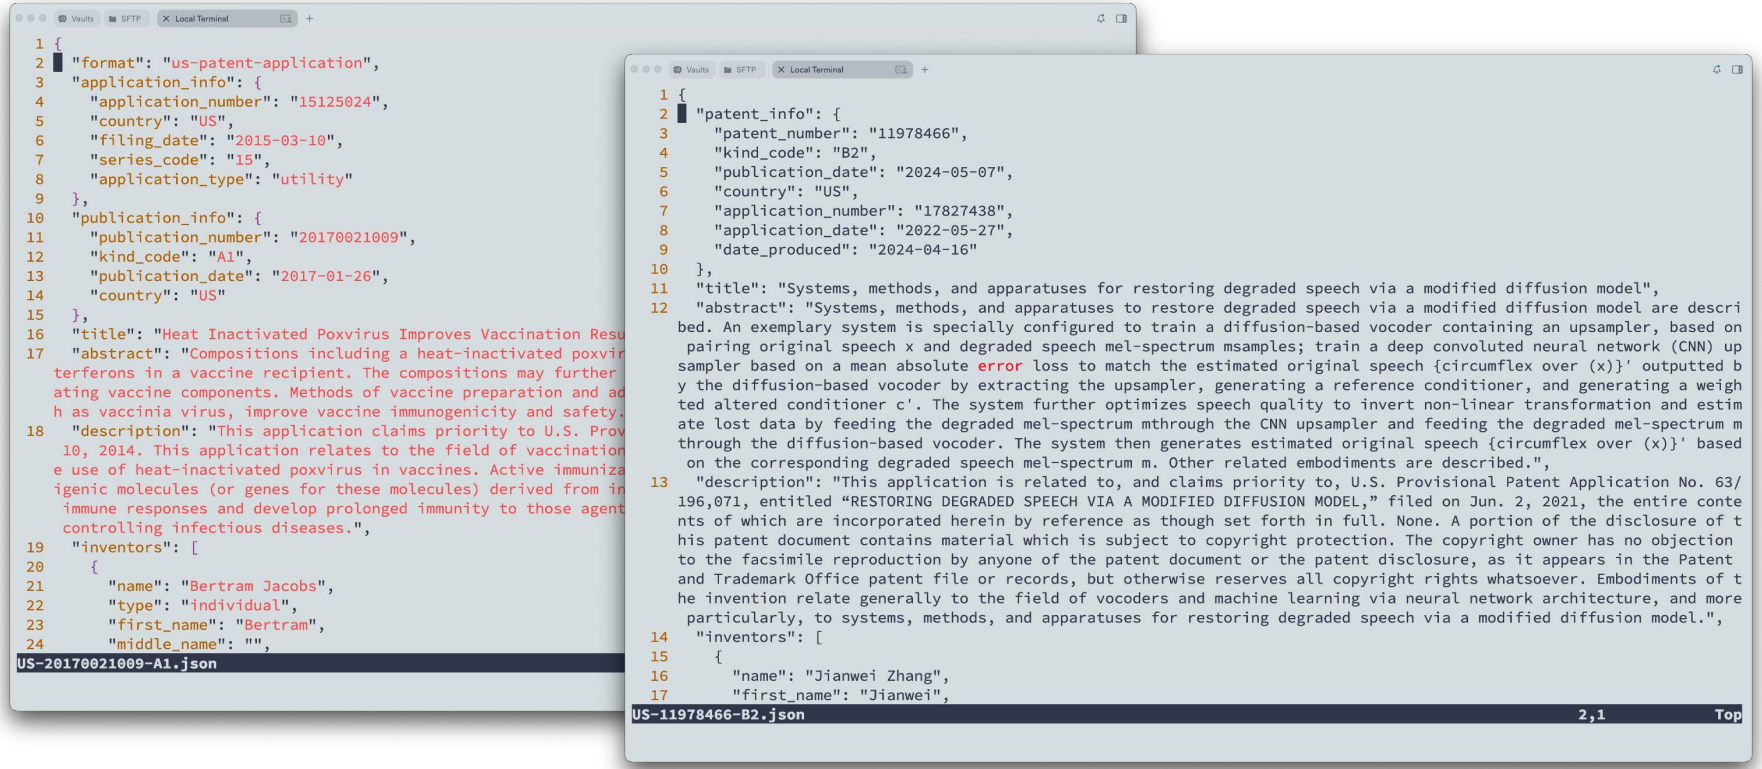

github.com/jgarciamesa/asu-uspto-data

Retrieval-Augmented Generation

Leverage CPC codes

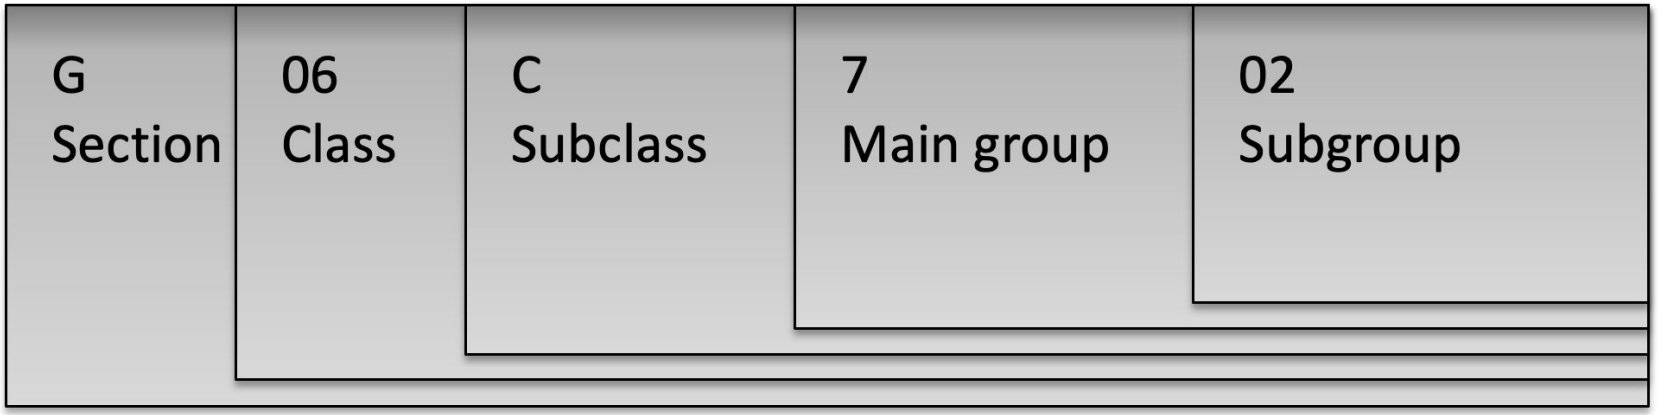

G Physics
06 Computing; calculating or counting
C Digital computers in which all the computation is effected mechanically
7 Input mechanisms
02 Keyboards

Leverage CPC codes

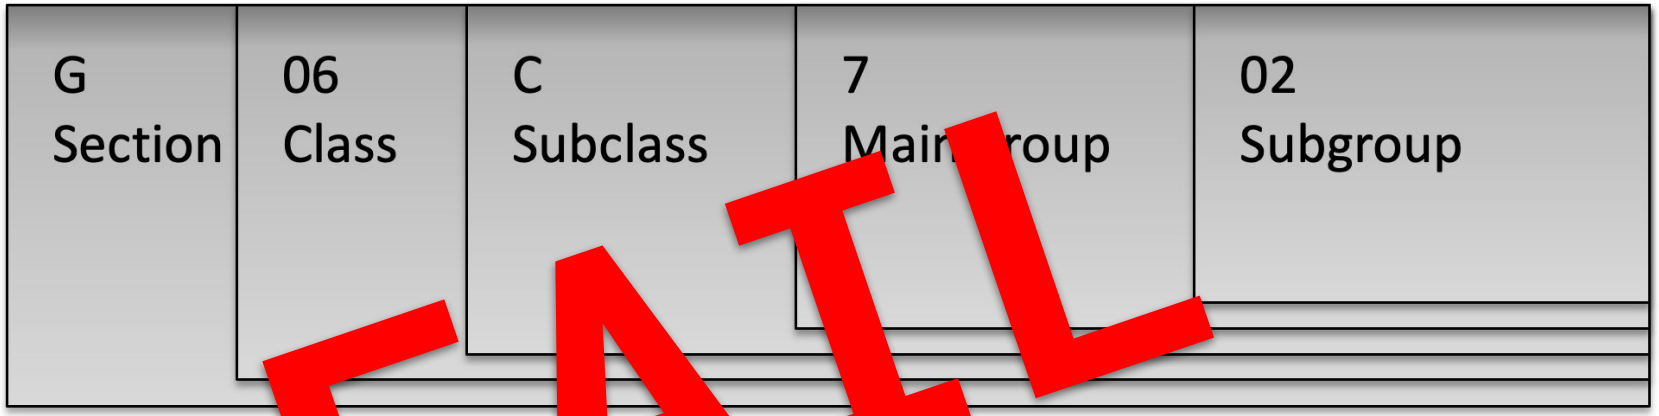

G Physics
06 Computing; calculating or counting
C Digital computers in which all the computation is effected mechanically
7 Input mechanisms
02 Keyboards

Traditional Retrieval-Augmented Generation

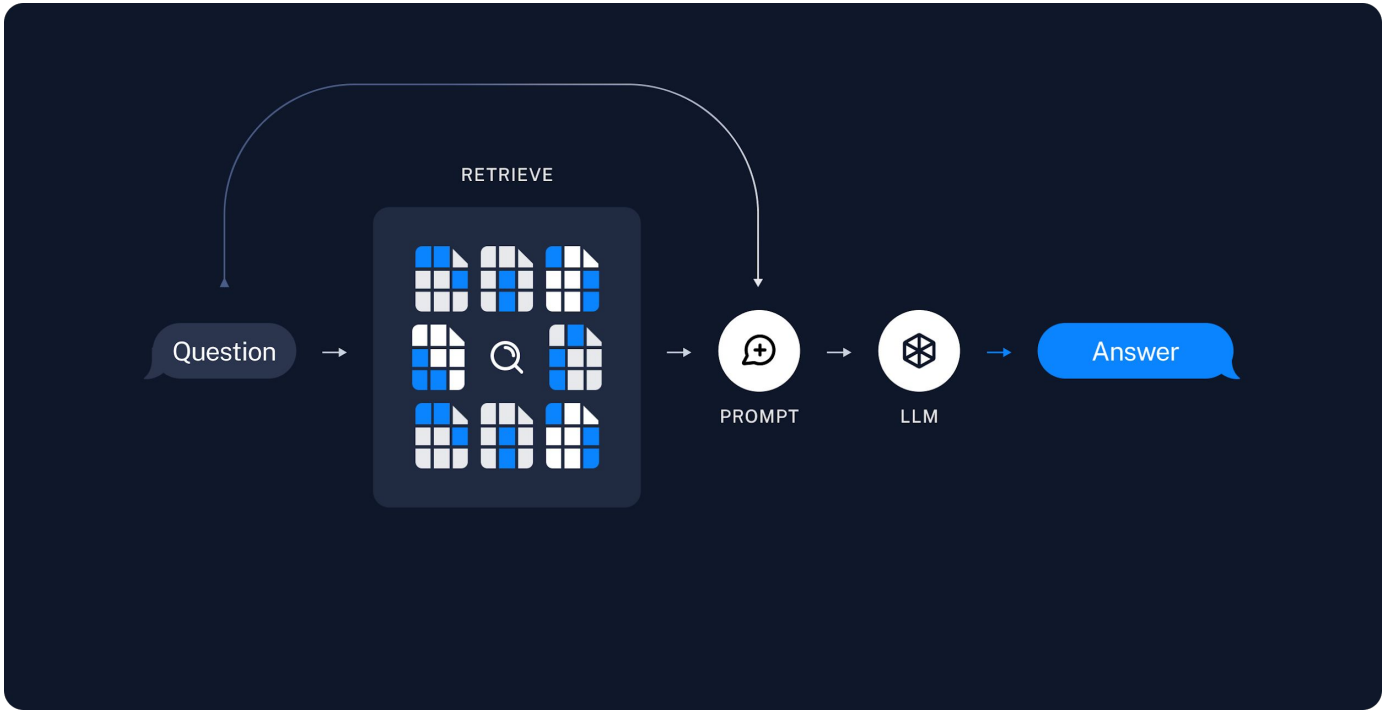

Traditional Retrieval-Augmented Generation

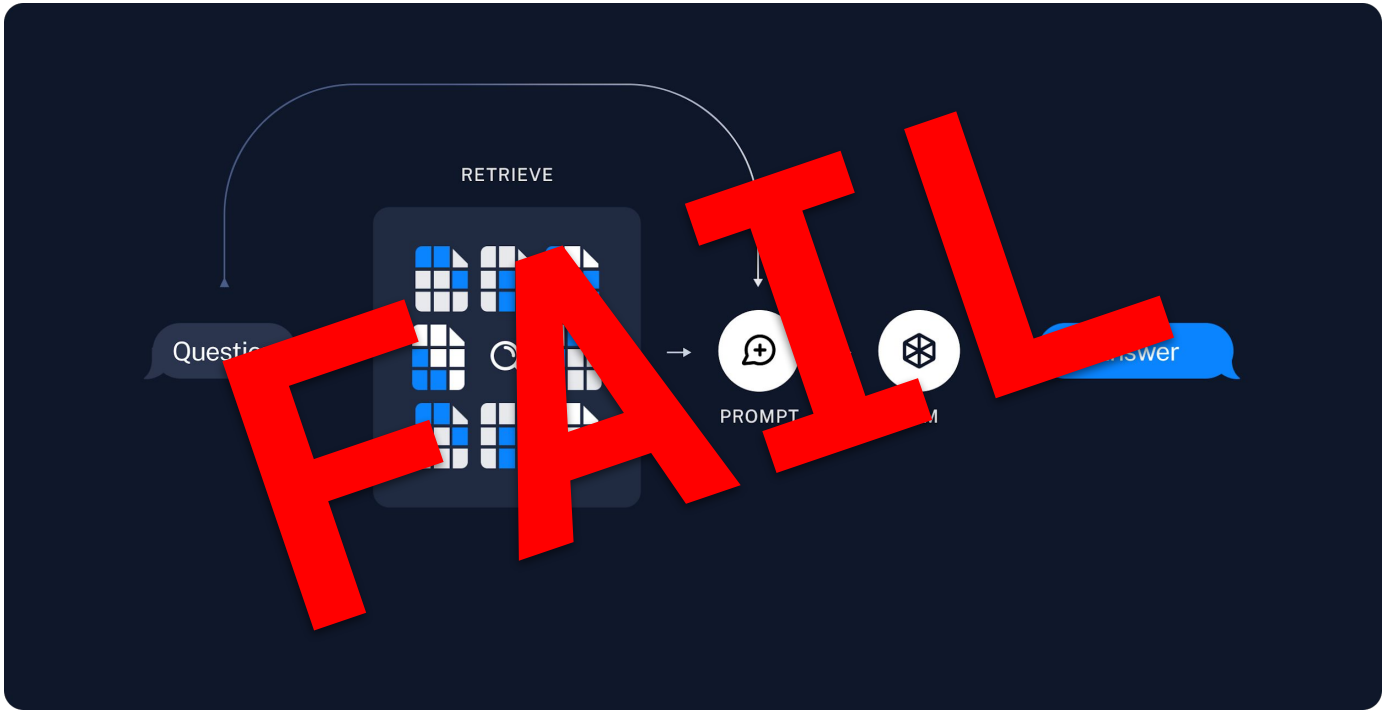

SOME TIME LATER...

Multi-stage Pipeline

https://github.com/jgarciamesa/asu-uspto

Build the Database

PostgreSQL + pgvector Embeddings

Each JSON file is added to the database

- title, abstract, claims, inventors, assignees
- classifications: a dict of CPC code
- searchable text: title + abstract + claims text + CPC codes
- technical areas: list of 1-3 LLM pre-generated taxonomy slugs
- embedding: pgvector embedding using searchable text
  https://github.com/pgvector/pgvector/

Query Workflow

Function Calling

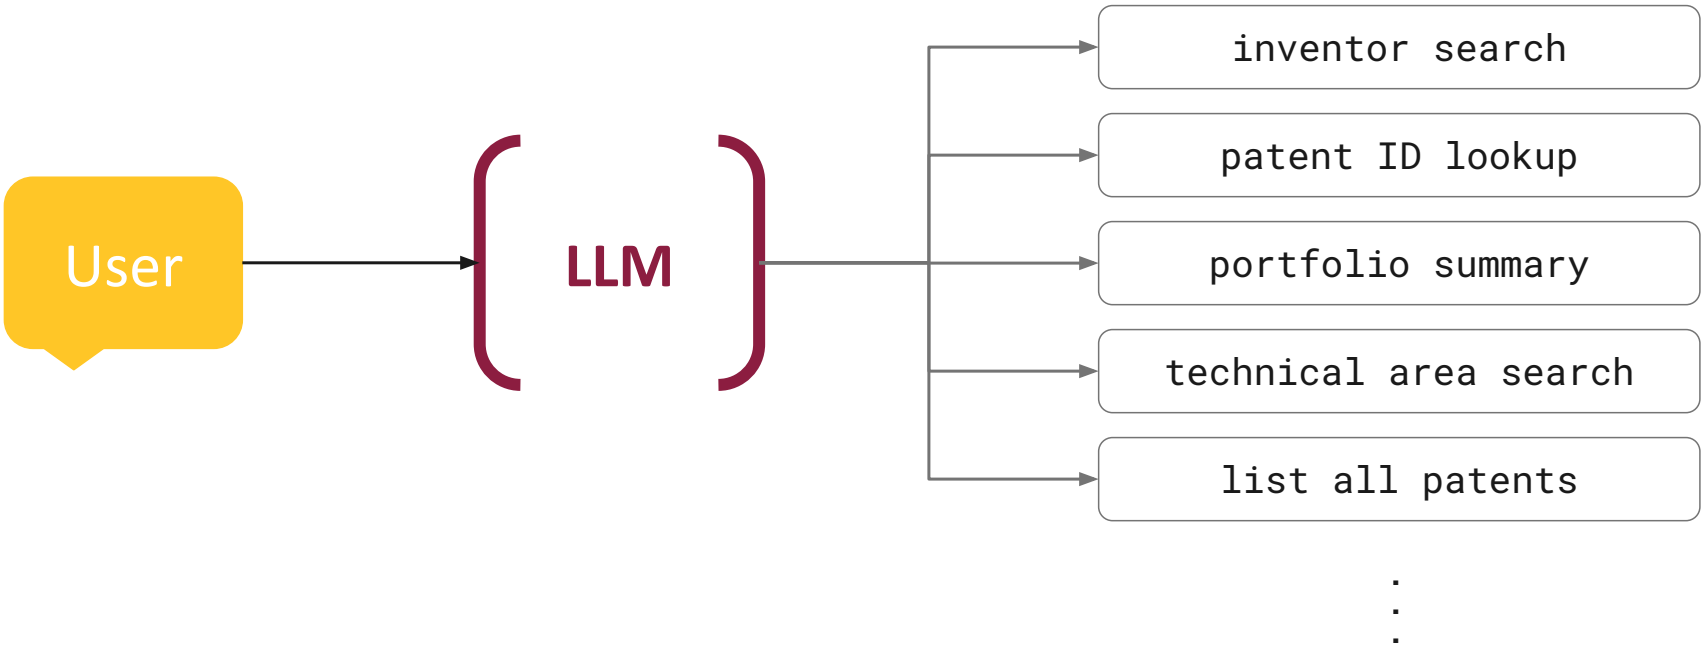

https://github.com/jgarciamesa/asu-uspto

```
{
    "type": "function",
    "function": {
        "name": "technical_area_search",
        "description": "Find patents that fall under a specific category.",
        "parameters: {
            "type": "object",
            "properties": {
                "categories": {
                    "type": "array",
                    "items": {"type": "string"},
                    "description": "List of categories to search for patents."
,
                },
            },
            "required": ["categories"],
        },
    },
}
```

[No Name] [+]

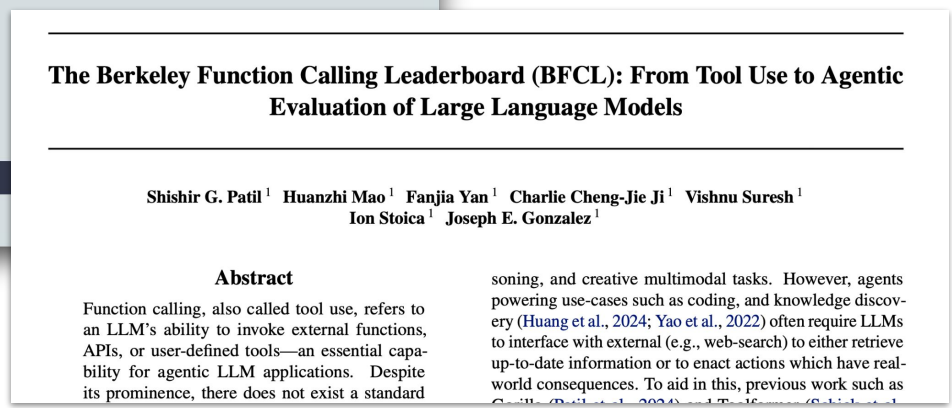

The Berkeley Function Calling Leaderboard (BFCL): From Tool Use to Agentic Evaluation of Large Language Models

Shishir G. Patil [1] **Huanzhi Mao** [1] **Fanjia Yan** [1] **Charlie Cheng-Jie Ji** [1] **Vishnu Suresh** [1]
Ion Stoica [1] **Joseph E. Gonzalez** [1]

Abstract

Function calling, also called tool use, refers to an LLM's ability to invoke external functions, APIs, or user-defined tools—an essential capability for agentic LLM applications. Despite its prominence, there does not exist a standard

soning, and creative multimodal tasks. However, agents powering use-cases such as coding, and knowledge discovery (Huang et al., 2024; Yao et al., 2022) often require LLMs to interface with external (e.g., web-search) to either retrieve up-to-date information or to enact actions which have real-world consequences. To aid in this, previous work such as

https://gorilla.cs.berkeley.edu/leaderboard.html

Patents Related to a Technical Area

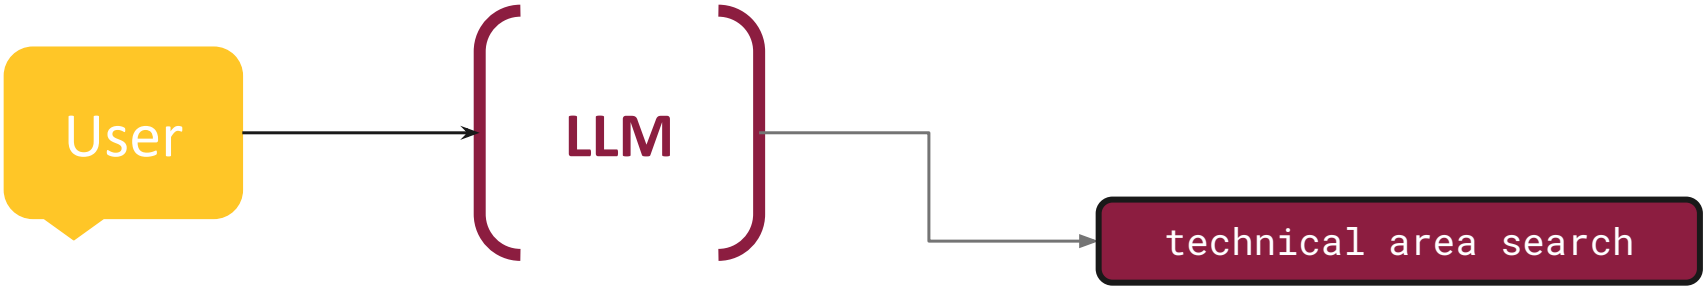

https://github.com/jgarciamesa/asu-uspto

Hybrid Retrieval Score (more art than science)

```
score = lexical * 0.4 + semantic * 0.4 + taxonomy * 0.15 + cpc * 0.05
```

- lexical: overlap between augmented query keywords and searchable text
- semantic: cosine similarity between query embedding and record embedding
- taxonomy: predicted query taxonomy overlap with technical areas and taxonomy keywords that appear in searchable text
- cpc: record CPC codes matching CPC prefixes from predicted taxonomy areas

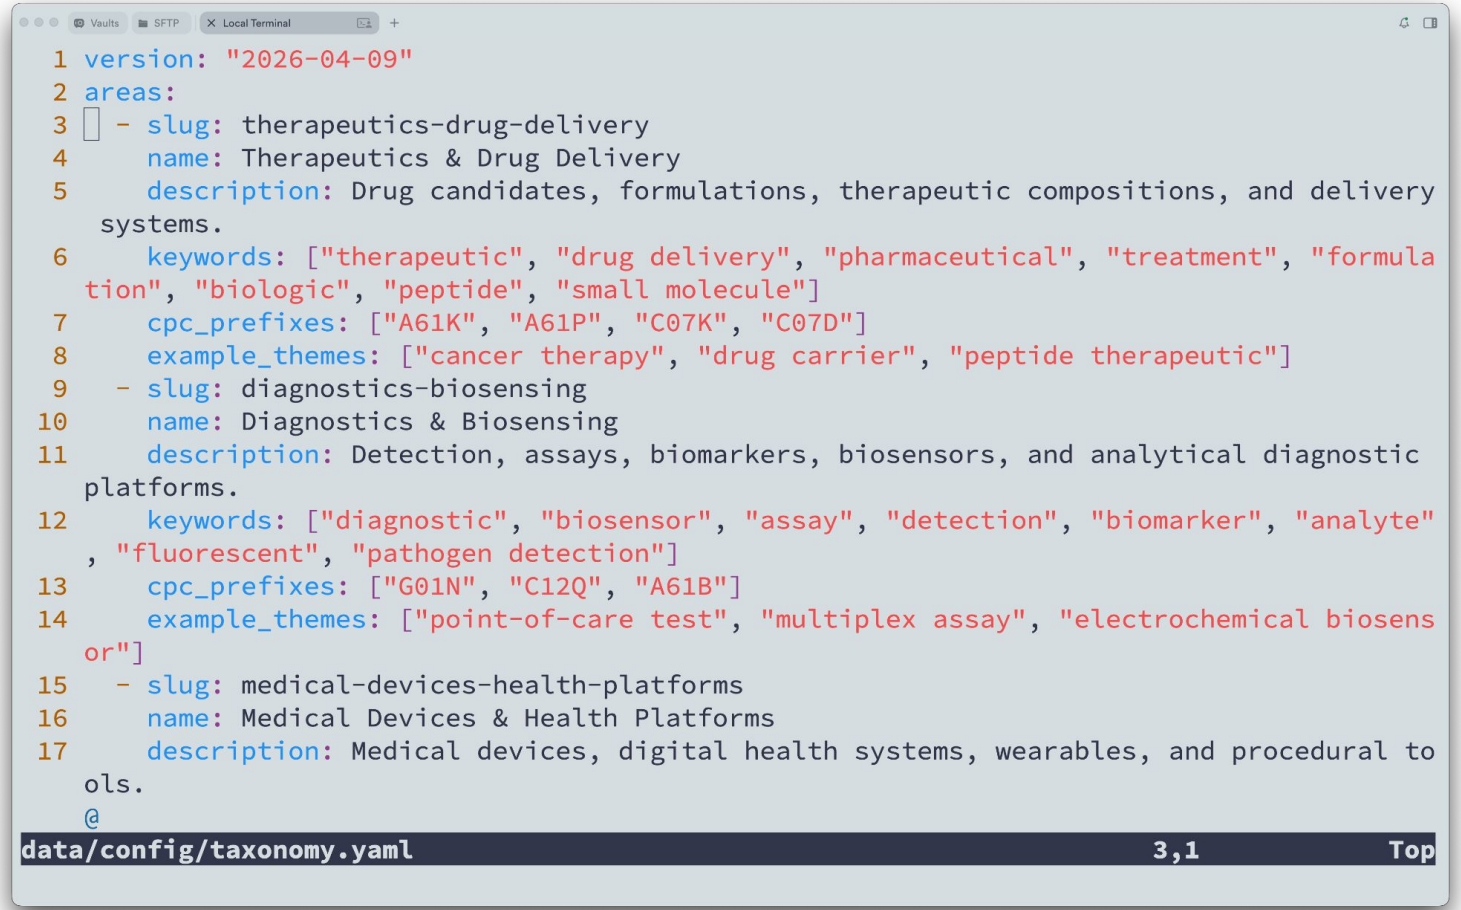

https://github.com/jgarciamesa/asu-uspto

Deployment

- Python FastAPI backend serving query/report endpoints
- Streamlit frontend for the chat
- Dockerized PostgreSQL database using the pgvector extension

- Currently exploring permanent deployment solutions

https://github.com/jgarciamesa/asu-uspto

Demo

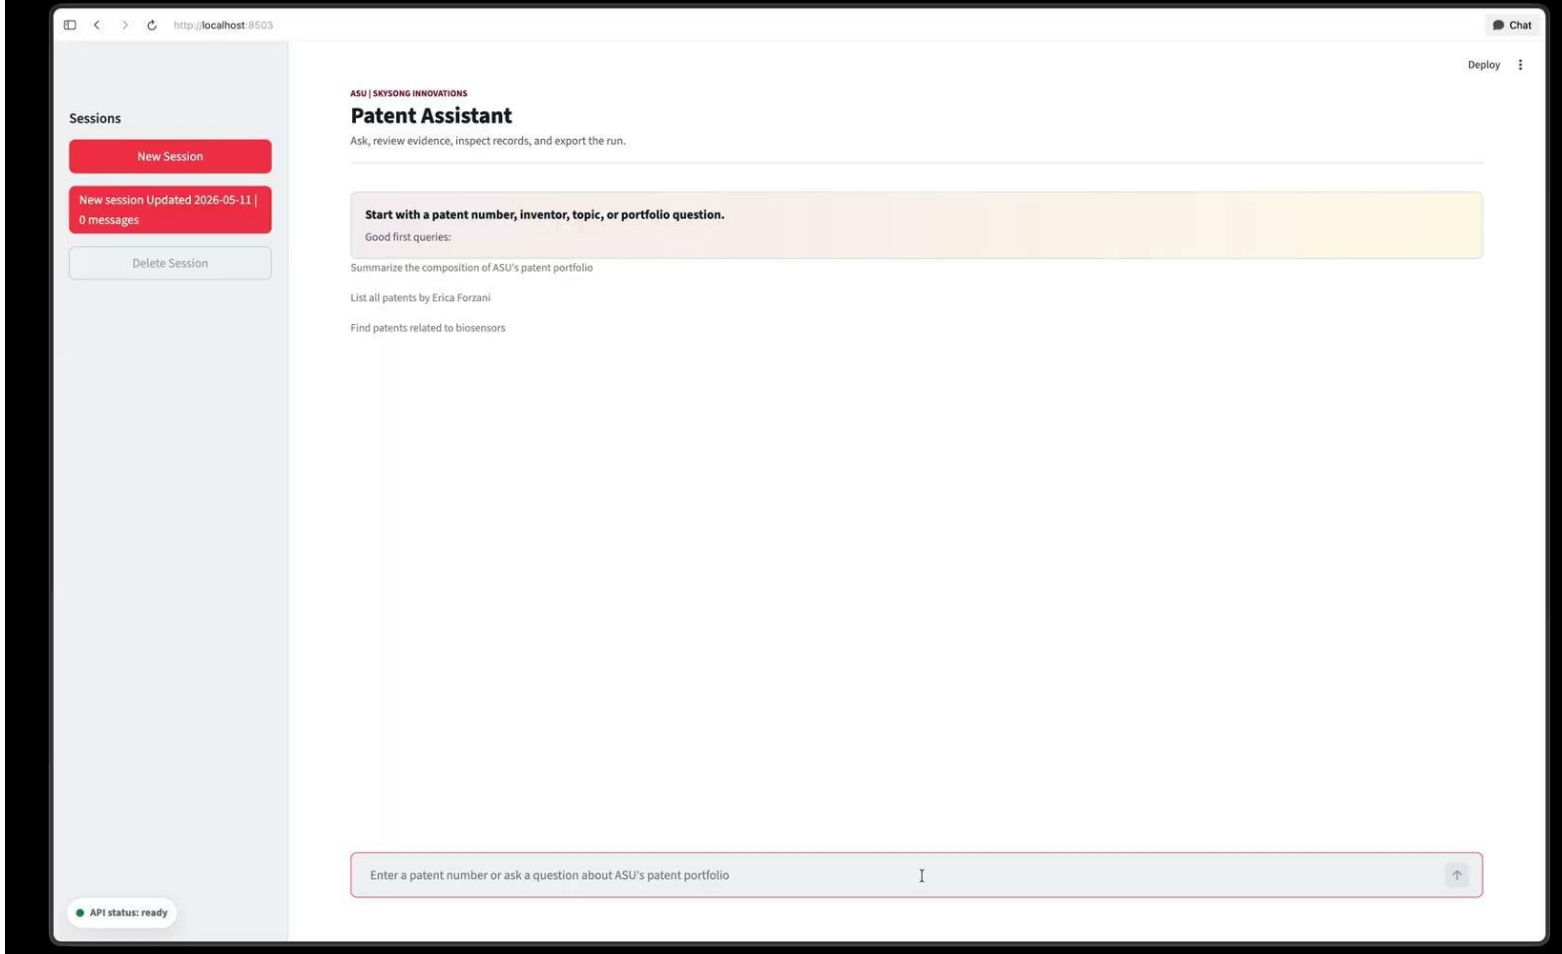

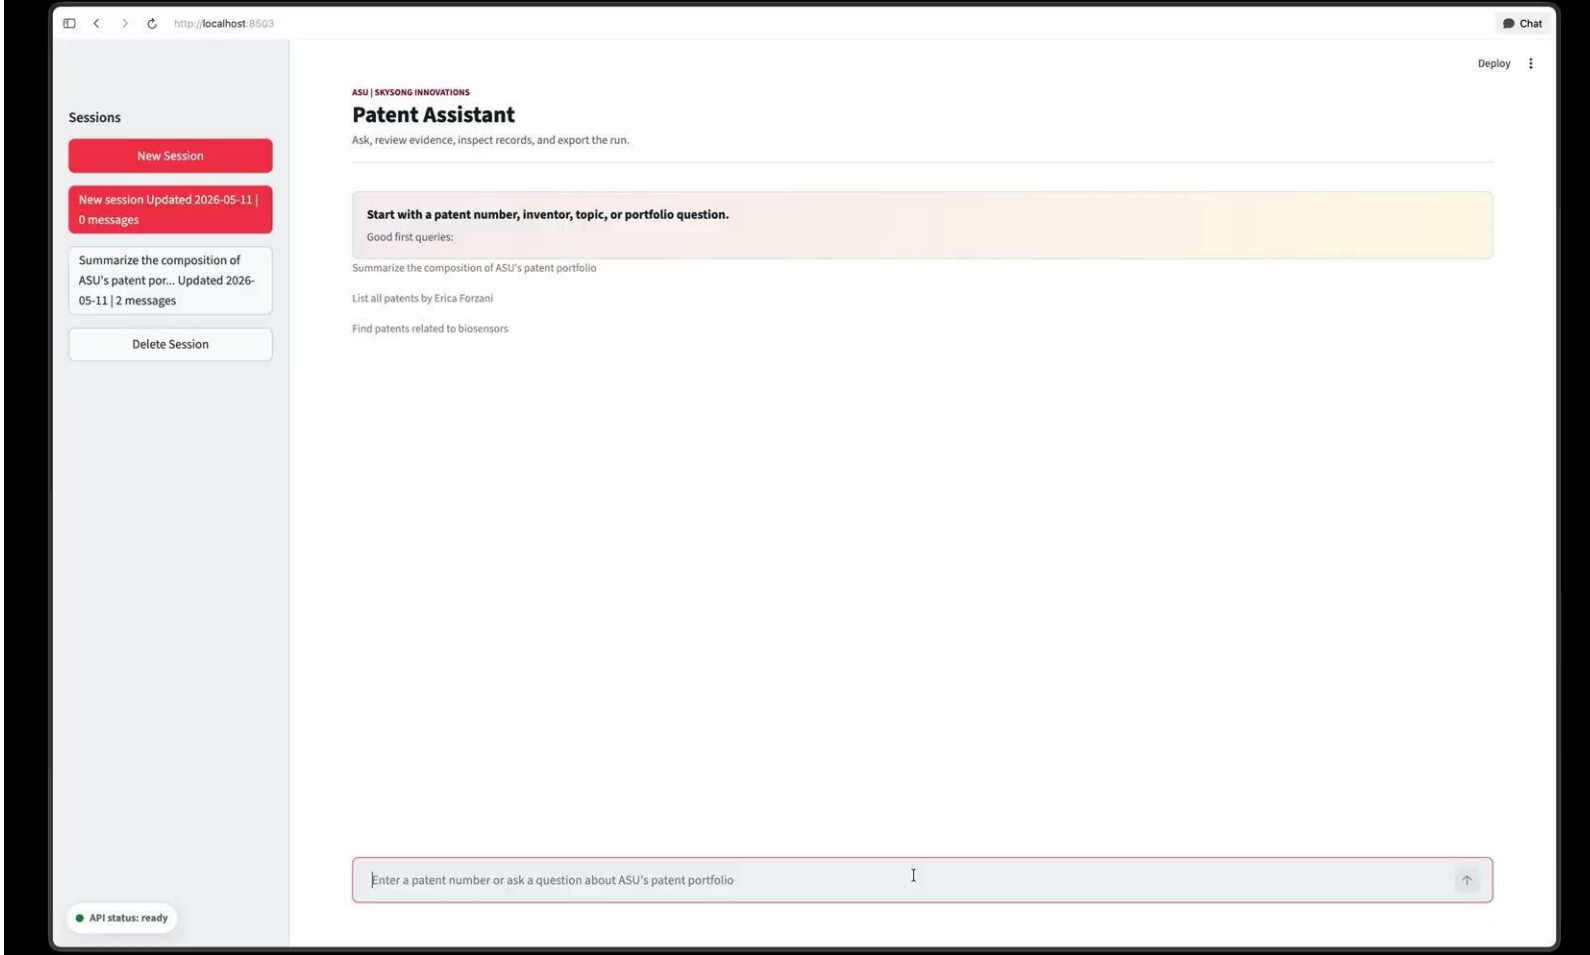
http://localhost:8503
Chat
Deploy
Sessions
New Session
New session Updated 2026-05-11 | 0 messages
Summarize the composition of ASU's patent por... Updated 2026-05-11 | 2 messages
Delete Session
ASU | SKYSONG INNOVATIONS
Patent Assistant
Ask, review evidence, inspect records, and export the run.
Start with a patent number, inventor, topic, or portfolio question.
Good first queries:
Summarize the composition of ASU's patent portfolio
List all patents by Erica Forzani
Find patents related to biosensors
Enter a patent number or ask a question about ASU's patent portfolio
API status: ready

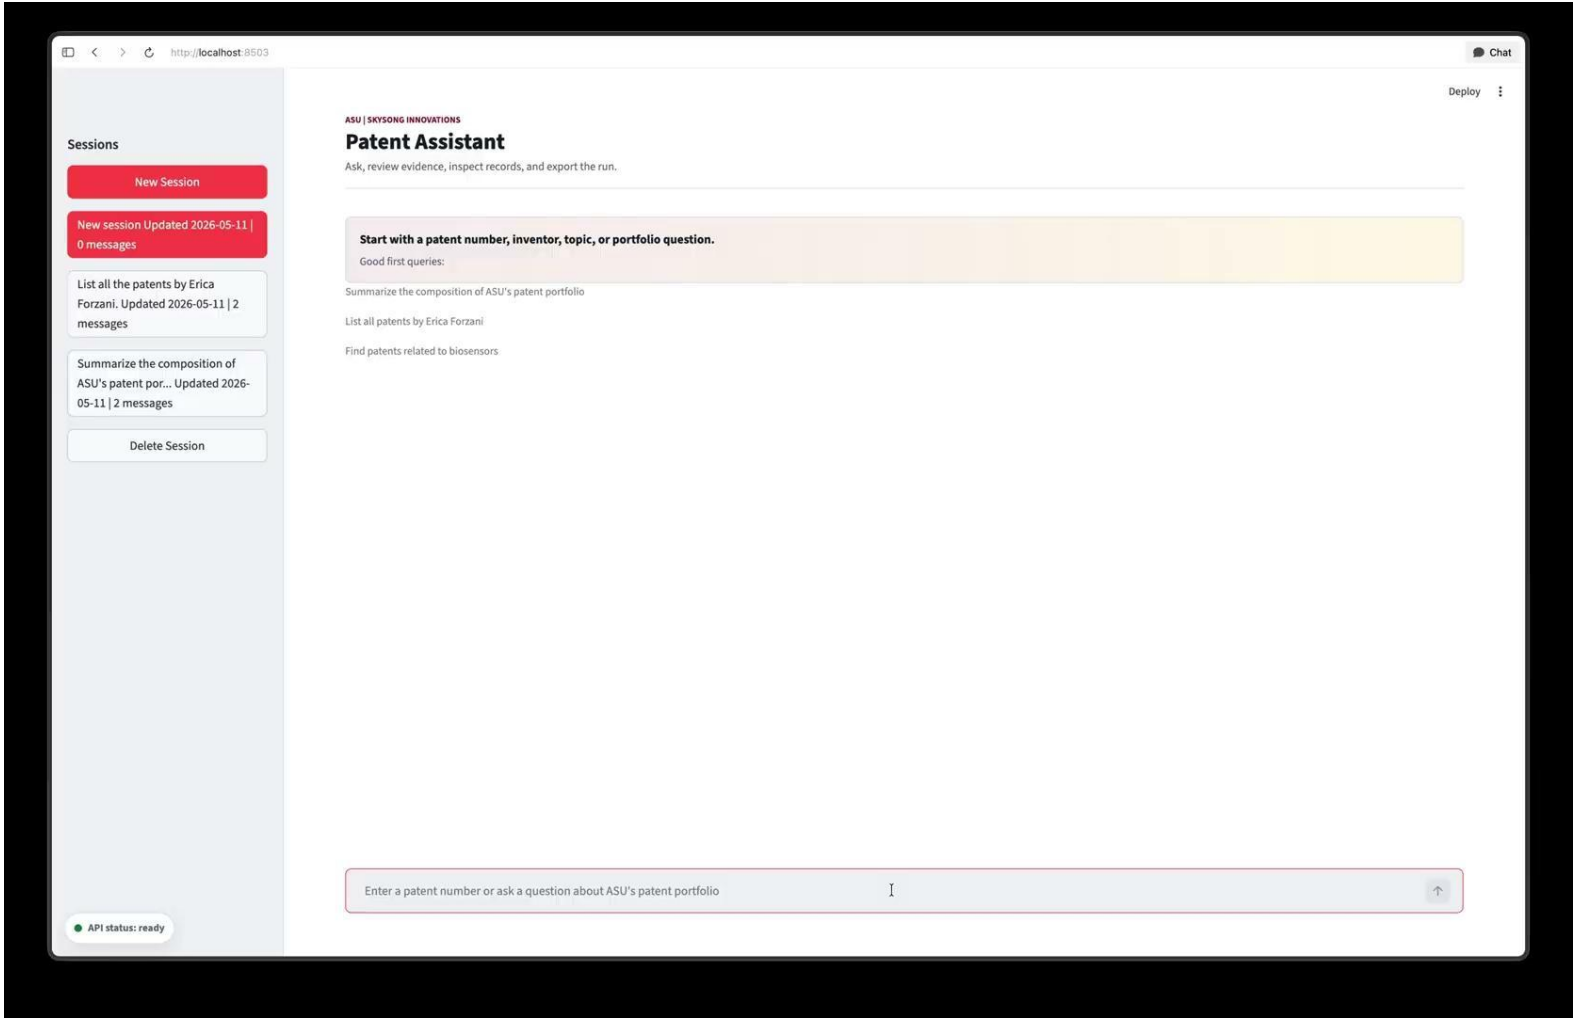

Future Work

- Internal deployment to SI for testing
- Public access for university partners
- Polish report generation
- Add general conversation skills
- Identify missing functionalities
- Decouple the tool from the data and create a blueprint for other applications

https://github.com/jgarciamesa/asu-uspto

Acknowledgements

Jordan Bell

Gil Speyer

Johnathan Lee

Roger Kohler

Faraj Tessilimi

Skysong Innovations
Advancing the Arizona State University Knowledge Enterprise

ASU Research Computing
Arizona State University

RMACC
Rocky Mountain Advanced Computing Consortium

https://github.com/jgarciamesa/asu-uspto