

# Bites of FM4S: [2] LLMs for experiments in fundamental physics

Tuesday 3 June 2025 - Tuesday 3 June 2025



*bites of...*  
Foundation Models  
for Science



Themes of discussion:

- Physics-inspired representations
- **LLMs for experiments**
- Applications to Run 3
- Training methods
- Alignment, interpretability
- Datasets and benchmarks

...and more!

## Book of Abstracts



# Contents

|  |   |
|--|---|
| Multi-Agent Research Validator & Enabler Using LLMs (MARVEL): Experiences from LIGO<br>. . . . . | 1 |
| The SpeakYSE: An Agentic LLM for Supernova Science . . . . .                                     | 1 |
| TBD [LLMs for LHCb] . . . . .  | 1 |
| chATLAS: An AI Assistant for the ATLAS Collaboration . . . . .                                   | 1 |
| LLM-based physics analysis assistant at BESIII . . . . .   | 2 |
| AccGPT: A Chatbot for CERN Internal Knowledge . . . . .  | 2 |
| Bridging LLMs and Scientific Infrastructure: A2rchi for Context-Aware Research Support           | 3 |



**Invited talks / 1****Multi-Agent Research Validator & Enabler Using LLMs (MARVEL): Experiences from LIGO****Corresponding Author:** nikhil\_m@mit.edu

Gravitational wave research at the Advanced LIGO observatories integrates complex, interconnected elements of experimental physics, computational simulations, and theoretical astrophysics. However, decades of valuable knowledge remain scattered across unstructured, multi-modal data and fragmented codebases. Efficient dissemination of this knowledge using large language models (LLMs) can significantly accelerate scientific discovery. In this talk, we share experiences from developing MARVEL, a modular, multi-agent research framework designed to provide scientific assistance in highly technical domains. MARVEL leverages open LLMs to ensure data privacy and is designed to be flexible enough to accommodate a broader range of scientific domains. We highlight challenges, including limitations in fine-tuning, pitfalls of naive Retrieval-Augmented Generation (RAG), model hallucinations, context window constraints, and difficulties in processing scientific documents via optical character recognition. To enhance factual accuracy and reasoning capabilities, MARVEL integrates tool usage and leverages test-time computing at the expense of increased latency. Finally, we emphasize the importance of modular workflows and custom benchmarks to adapt to advances in foundational models rapidly.

**Theme of discussion:****Invited talks / 2****The SpeakYSE: An Agentic LLM for Supernova Science****Corresponding Author:** gaglian2@mit.edu

Time-domain astronomy is rapidly entering a data-rich era in which wide-field surveys discover millions of transients per year, overwhelming traditional, hand-driven analysis. In this talk we present SpeakYSE, an agentic and open-source language model that turns natural-language requests into end-to-end analyses for the Young Supernova Experiment. The SpeakYSE links literature retrieval, database reasoning, and low-level tool-calling for on-the-fly exploratory data analysis. We describe the SpeakYSE's architecture, early results from its use within the collaboration, and future directions enabled by next-generation reasoning models. Domain-specific LLM agents like SpeakYSE will be essential for exploiting next-generation surveys such as the Rubin LSST.

**Theme of discussion:****Invited talks / 4****TBD [LLMs for LHCb]****Invited talks / 5****chATLAS: An AI Assistant for the ATLAS Collaboration**

**Corresponding Author:** daniel.thomas.murnane@cern.ch

The ATLAS Collaboration is composed of around 6,000 scientists, engineers, developers, students and administrators, with decades of institutional documentation spread across wikis, code docs, meeting agendas, recommendations, publications, tutorials, and project management systems. With the advent of retrieval augmented generation (RAG) and sophisticated large language models (LLMs) such as GPT-4, there is now an opportunity to produce a “front door” to this intimidatingly large corpus. ChATLAS is an attempt to provide this entrypoint, as ATLAS’ official AI assistant and search system. In this contribution, we review the past year of developments, present the latest updates to the system, and introduce ongoing work to improve back-end performance, agentic information gathering, and science-centric design components.

**Theme of discussion:**

**Invited talks / 7**

## LLM-based physics analysis assistant at BESIII

**Authors:** Beijiang Liu<sup>None</sup>; Changzheng YUAN<sup>None</sup>; Ke Li<sup>1</sup>; Zhengde Zhang<sup>2</sup>

<sup>1</sup> *Chinese Academy of Sciences (CN)*

<sup>2</sup> 中国科学院高能物理研究所

**Corresponding Authors:** liubj@ihep.ac.cn, zdzhang@ihep.ac.cn, ke.li@cern.ch, yuancz@ihep.ac.cn

The data processing and analyzing is one of the main challenges at HEP experiments. To accelerate the physics analysis and drive new physics discovery, the rapidly developing Large Language Model (LLM) is the most promising approach, it have demonstrated astonishing capabilities in recognition and generation of text while most parts of physics analysis can be benefitted. In this talk we will discuss the construction of a dedicated intelligent agent, an AI assistant names Dr.Sai at BESIII based on LLM, the potential usage to boost hadron spectroscopy study, and the future plan towards a AI scientist.

**Theme of discussion:**

LLMs for operations

**Invited talks / 8**

## AccGPT: A Chatbot for CERN Internal Knowledge

**Author:** Florian Rehm<sup>1</sup>

<sup>1</sup> *CERN*

**Corresponding Author:** florian.matthias.rehm@cern.ch

AccGPT is an innovative pilot project utilizing Large Language Models (LLMs) to create a chatbot for interacting with CERN’s extensive internal knowledge base. This initiative is primarily led by the CERN Beams and IT departments, while the objective is to make this chatbot available to the

entire CERN community. AccGPT is designed to provide quick and straightforward answers to queries, similar to ChatGPT, thereby enhancing productivity and decreasing the time experts spend on support tasks. Looking ahead, there are plans to expand AccGPT's functionalities, for example utilizing enhanced Agent features.

**Theme of discussion:**

LLMs for operations

**Invited talks / 9**

## **Bridging LLMs and Scientific Infrastructure: A2rchi for Context-Aware Research Support**

**Authors:** Mariarosaria D'Alfonso<sup>1</sup>; Pietro Lugato<sup>1</sup>

<sup>1</sup> *Massachusetts Inst. of Technology (US)*

**Corresponding Authors:** mariarosaria.d'alfonso@cern.ch, pietro.lugato@cern.ch

We introduce A2rchi (AI-Augmented Research Chat Intelligence), an intelligent, domain-adaptable open source chatbot designed to support research and education workflows.

Beyond its core functionalities, A2rchi also integrates with common communication and workflow tools, including email systems, ticketing platforms, and collaboration platforms like Mattermost, offering seamless assistance across multiple channels.

Leveraging Retrieval-Augmented Generation (RAG), A2rchi combines foundational large language models with custom, project-specific data—such as course materials and documentation—to deliver accurate, context-aware responses.

Originally developed and deployed to support MIT classroom instructions and the Physics Department's analysis facility, A2rchi is now being expanded to serve the CMS experiment at CERN, with applications ranging from Tier-0 operations and data management to end-user physics analysis support.

We present the current implementation, lessons learned from real-world deployments, and the roadmap for scaling A2rchi into a robust, domain-aware assistant for large-scale scientific collaborations.

**Theme of discussion:**

LLMs for operations