Engagement with the User Community

Matt Heine

BCS SubMIT Review

June 3, 2026

Outline

Please accept my apologies:
The recent rescheduling of this Review to an earlier time makes it now conflict with a pre-existing (time-sensitive) medical appointment.

- Annual SubMIT Workshop
- Users Group Meetings
- Classroom Usage,
  User-Run Workshops using SubMIT
- System-Level Customizations / User Requests
- Current Limitations & Open Challenges

Annual SubMIT Workshop

Our annual large community engagement event

- introducing potential/new users to SubMIT
- seeking feedback from the community
- sharing knowledge with the community
- meeting users, discussing user needs/challenges
- users showcase their use of SubMIT

- This year added: Hands-on Computing hackathon & help desk session
- Also new format: Project Team tutorials interspersed throughout the day

<u>Project Team Tutorials/Talks</u>

Same Project Team topics as last year + new software (Globus/OpenMPI) + A2rchi talk

- SubMIT Overview; connection/interface options
- Containers for Portable Software
- Virtual environments, conda
- OpenMPI, Globus
- Batch Jobs: SLURM, HTCondor
- A2rchi

<u>User-Contributed Talks</u>

Title	User
MCMCs: Using subMIT for early universe cosmology	Sarah Geller
Modeling dark matter signals in cosmic 21cm radiation	Tracy Slatyer
CMS analyses using SubMIT	Guillelmo Gomez-Ceballos
Machine learning quantum states of matter	Ahmed Abouelkomsan
Hacking on SubMIT	Lina Necib

Attendance ~ 13 non-team members (including speakers)

User-Contributed Talks at User Group Meetings

	Title	Speaker (User)	Speaker Position	
June, 2025	Simulating heavy quarks in quark gluon plasma	Jean Du Plessis	Graduate student	CTP
July, 2025	Is attention all you need to solve the correlated electron problem?	Max Geier	Postdoctoral Associate	CMT
September, 2025	A2rchi – An Open Source LLM Framework	Ivan Paus, **Luca Lavezzo**, Pietro Lugato	Graduate student	LNS
October, 2025	Multimodal training of jet taggers on subMIT	Benedikt Maier	Postdoctoral Fellow	LNS
May, 2026	An integrated neural wavefunction solver for spinful Fermi systems	**Alexander Avdoshkin**	Postdoctoral Associate	CMT

Bold denotes SubMIT Project Team members

- AI / Machine Learning dominated user-given talks this year
- Each meeting opens with updates from the Project Team
- New protocol since last Review: Reach out to User Group Representatives **prior to** User Group meetings in effort to increase engagement
- Attendance ~2-4 User Group Representatives + ~2-3 others (not counting project team)

User-Run Workshops & Classroom Use

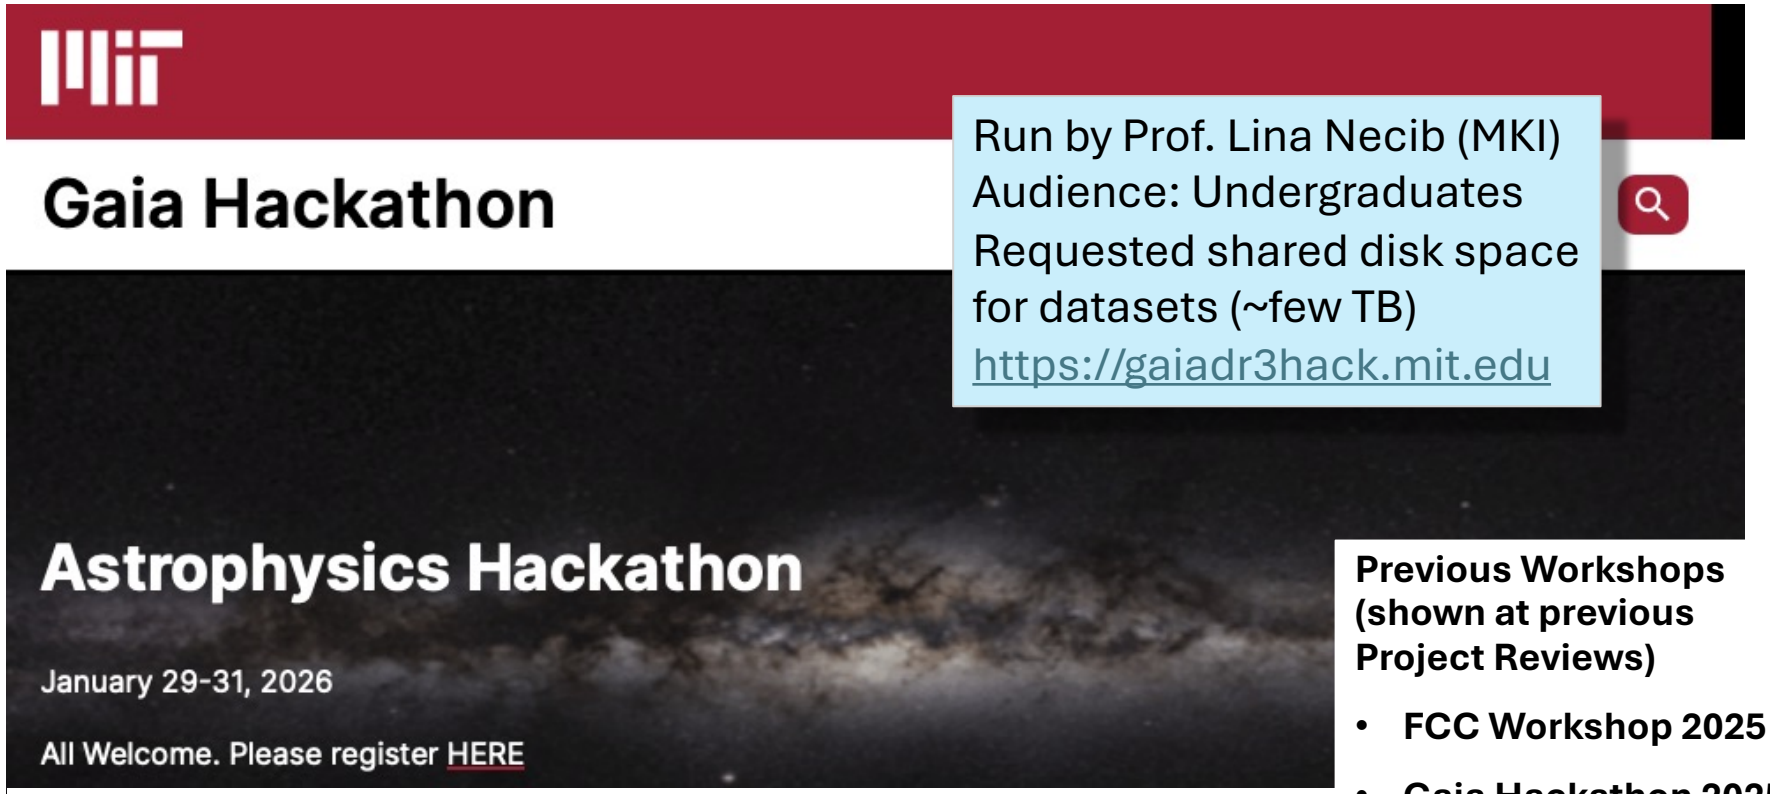

Run by Prof. Lina Necib (MKI)
Audience: Undergraduates
Requested shared disk space for datasets (~few TB)
https://gaiadr3hack.mit.edu

Previous Workshops (shown at previous Project Reviews)

- **FCC Workshop 2025**
- **Gaia Hackathon 2025**

Why was SubMIT used?

- Easy access to few-TB datasets for many students simultaneously (Hackathon)
- Easy student access to pre-installed software (ESG 8.022, 8.02, 8.284)
  - Included custom JupyterHub server extension to accommodate visualization package (will not work on Google Colab)
- Introduce students to clusters (Hackathon Learning Goal) & Jupyter Notebooks (ESG 8.022 Learning Goal)
- A2rchi hosting (8.01)

Cumulative Classroom Use

Bold denotes usage this past year

- **ESG 8.022** (Jupyter notebooks)
- **8.13/8.14 Jr Lab**
- 8.01 (A2rchi)
- 8.02
- 8.284 Modern Astrophysics
- Educators may use w/out notifying us
- Discussion towards future classroom use (during Users meeting)

User Requests & System-Level Accommodations

(beyond ordinary support)

New System-Level User Accommodations

- gpu-express queue: Users said GPU code development was choked by unavailability of GPUs for debugging/testing (busy running production jobs) → implemented interactive queue (dedicated node, very short wall-time)

Previous System-Level User Accommodations

- Globus, OpenMPI, group websites, JupyterHub server customizations, MLOps server, LIGO cvmfs mounts, central classroom software, user-requested tutorials …

These were discussed at previous Project Reviews

Declined User Requests

- Request for centrally-installed common software/libraries (eg: cfitsio, fortran libraries)
- Request to accommodate Windows software
- Request for SubMIT documentation of inefficiency of native Python loops compared to compiled code [or vectorized Python operations, e.g. NumPy]
- Installation of Mathematica on more than 1 node. Barrier: licensing

* See backup slide for more details on Windows request

**Some requests came from audience members at meetings/workshop; these may not have been users but researchers shopping for a system / solution

Open Challenges

Balancing Restrictions/Rules

- Rules/Restrictions can prevent a few users from dominating the system
  - … but they also necessarily limit users. → A careful balance must be struck

  Example: short wall-time, whole-node jobs can help throughput + turnaround time mandated by some clusters) but that would require some user workflows to be rewritten

Priority Access of User-Purchased Hardware

- Desire to allow shared use of user-purchased nodes when not in use by the purchasing user … but purchasing user prefers not to wait when they need it
- Currently handled via manual intervention from SubMIT team in exceptional cases (e.g. temporary reservations of purchased nodes)
- SubMIT policy: no job preemption (no killing running jobs) & minimize nodes being reserved exclusively for purchasing users (reduce idle nodes)

Backup Slides

Windows Software Request Details

- SubMIT servers run AlmaLinux 9. Workarounds to support Windows on AlmaLinux (i.e. virtual machines) are too resource-intensive for the shared SubMIT resources.
- Windows system admin would require significant additional time & specialized Windows expertise (e.g. security, etc.), so SubMIT does not support Microsoft Windows
- Most likely most efficient for user to run on their own workstation or attempt to use WineHQ on SubMIT
- Asked user to let us know if this is not a sufficient option

Open Challenges --verbose

- **Balancing Restrictions/Rules vs Open/Fairshare Access**
  - There is tension between the desires to:
    - Have machines available immediately when users want to them
    - Not have idle resources / Not reserve resources for interactive use, besides JuypterHub
    - Have long wall time
    - Allow reservation of less than a full node
  - If left unchecked, a few users can dominate a rule-free open system … but adding global rules to prevent this necessarily limits users. A careful balance must be struck.
  - Long wall time + few restrictions on job characteristics lowers the barrier to cluster use for our diverse userbase … but can limit turnaround time (waiting in queue) and even throughput when the system is busy, depending on the types of jobs submitted.

    Example: short wall-time, whole-node jobs can help throughput + turnaround time (& are mandated by some clusters) but user workflows likely need to be rewritten for this
- **Priority Access for User-Purchased Hardware**
  - We strive to make use of idle machines … but hardware-purchasing users prefer not to wait to use their hardware. SubMIT policy has always been not to kill running jobs (no job preemption)
  - Currently handled via manual intervention from SubMIT team in exceptional cases (e.g. temporary reservations of purchased nodes)