

Community Engagement: MKI



Josh Borrow (MKI)




subMIT Review: 18 July 2023

About MKI


Build and launch telescopes




Run massive simulations
(100k cores)



Small-scale simulations



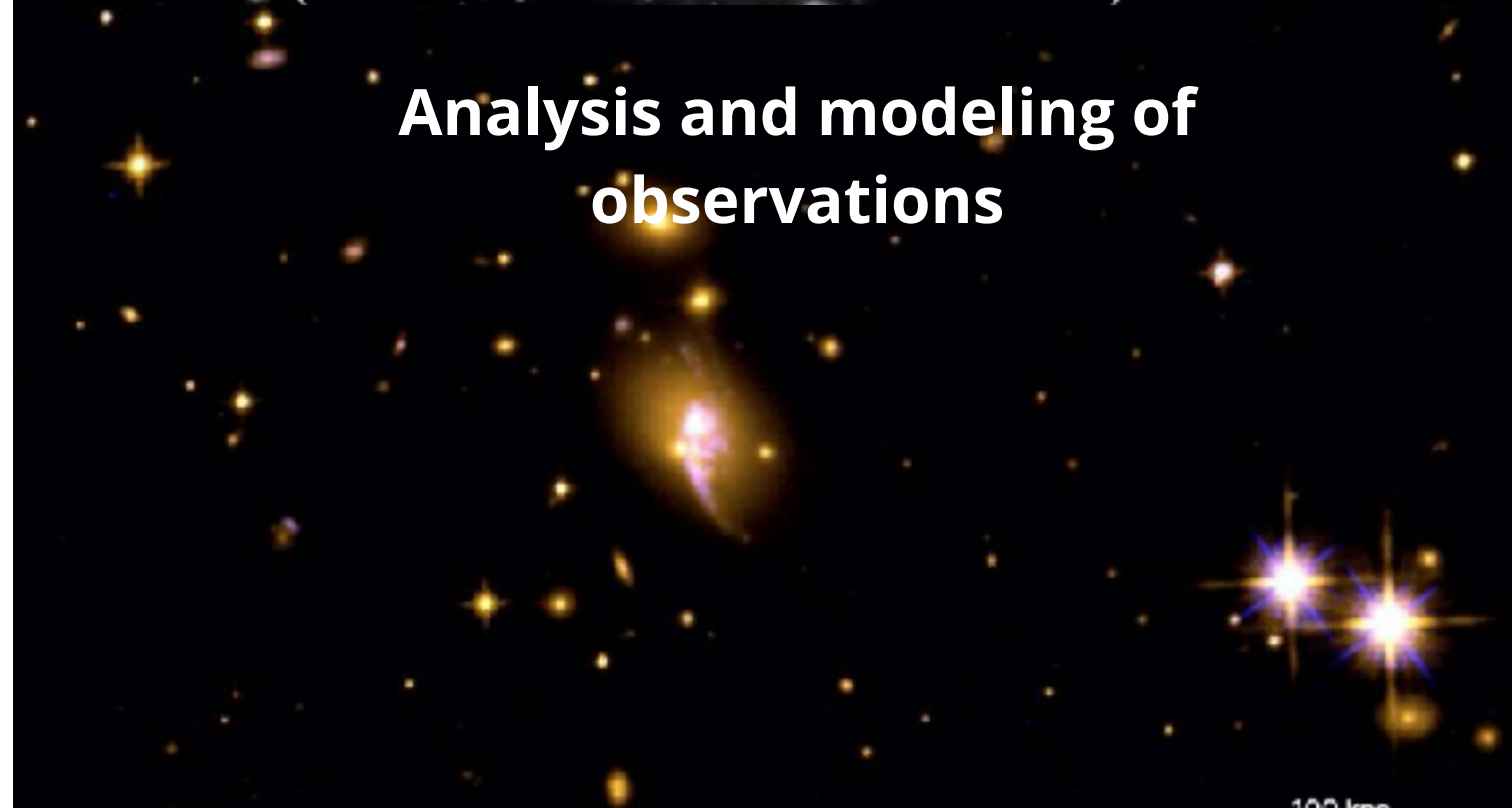
Big data analysis
(Pbs; simulated)




Big data analysis
(Time-series)



Analysis and modeling of
observations




Machine learning models



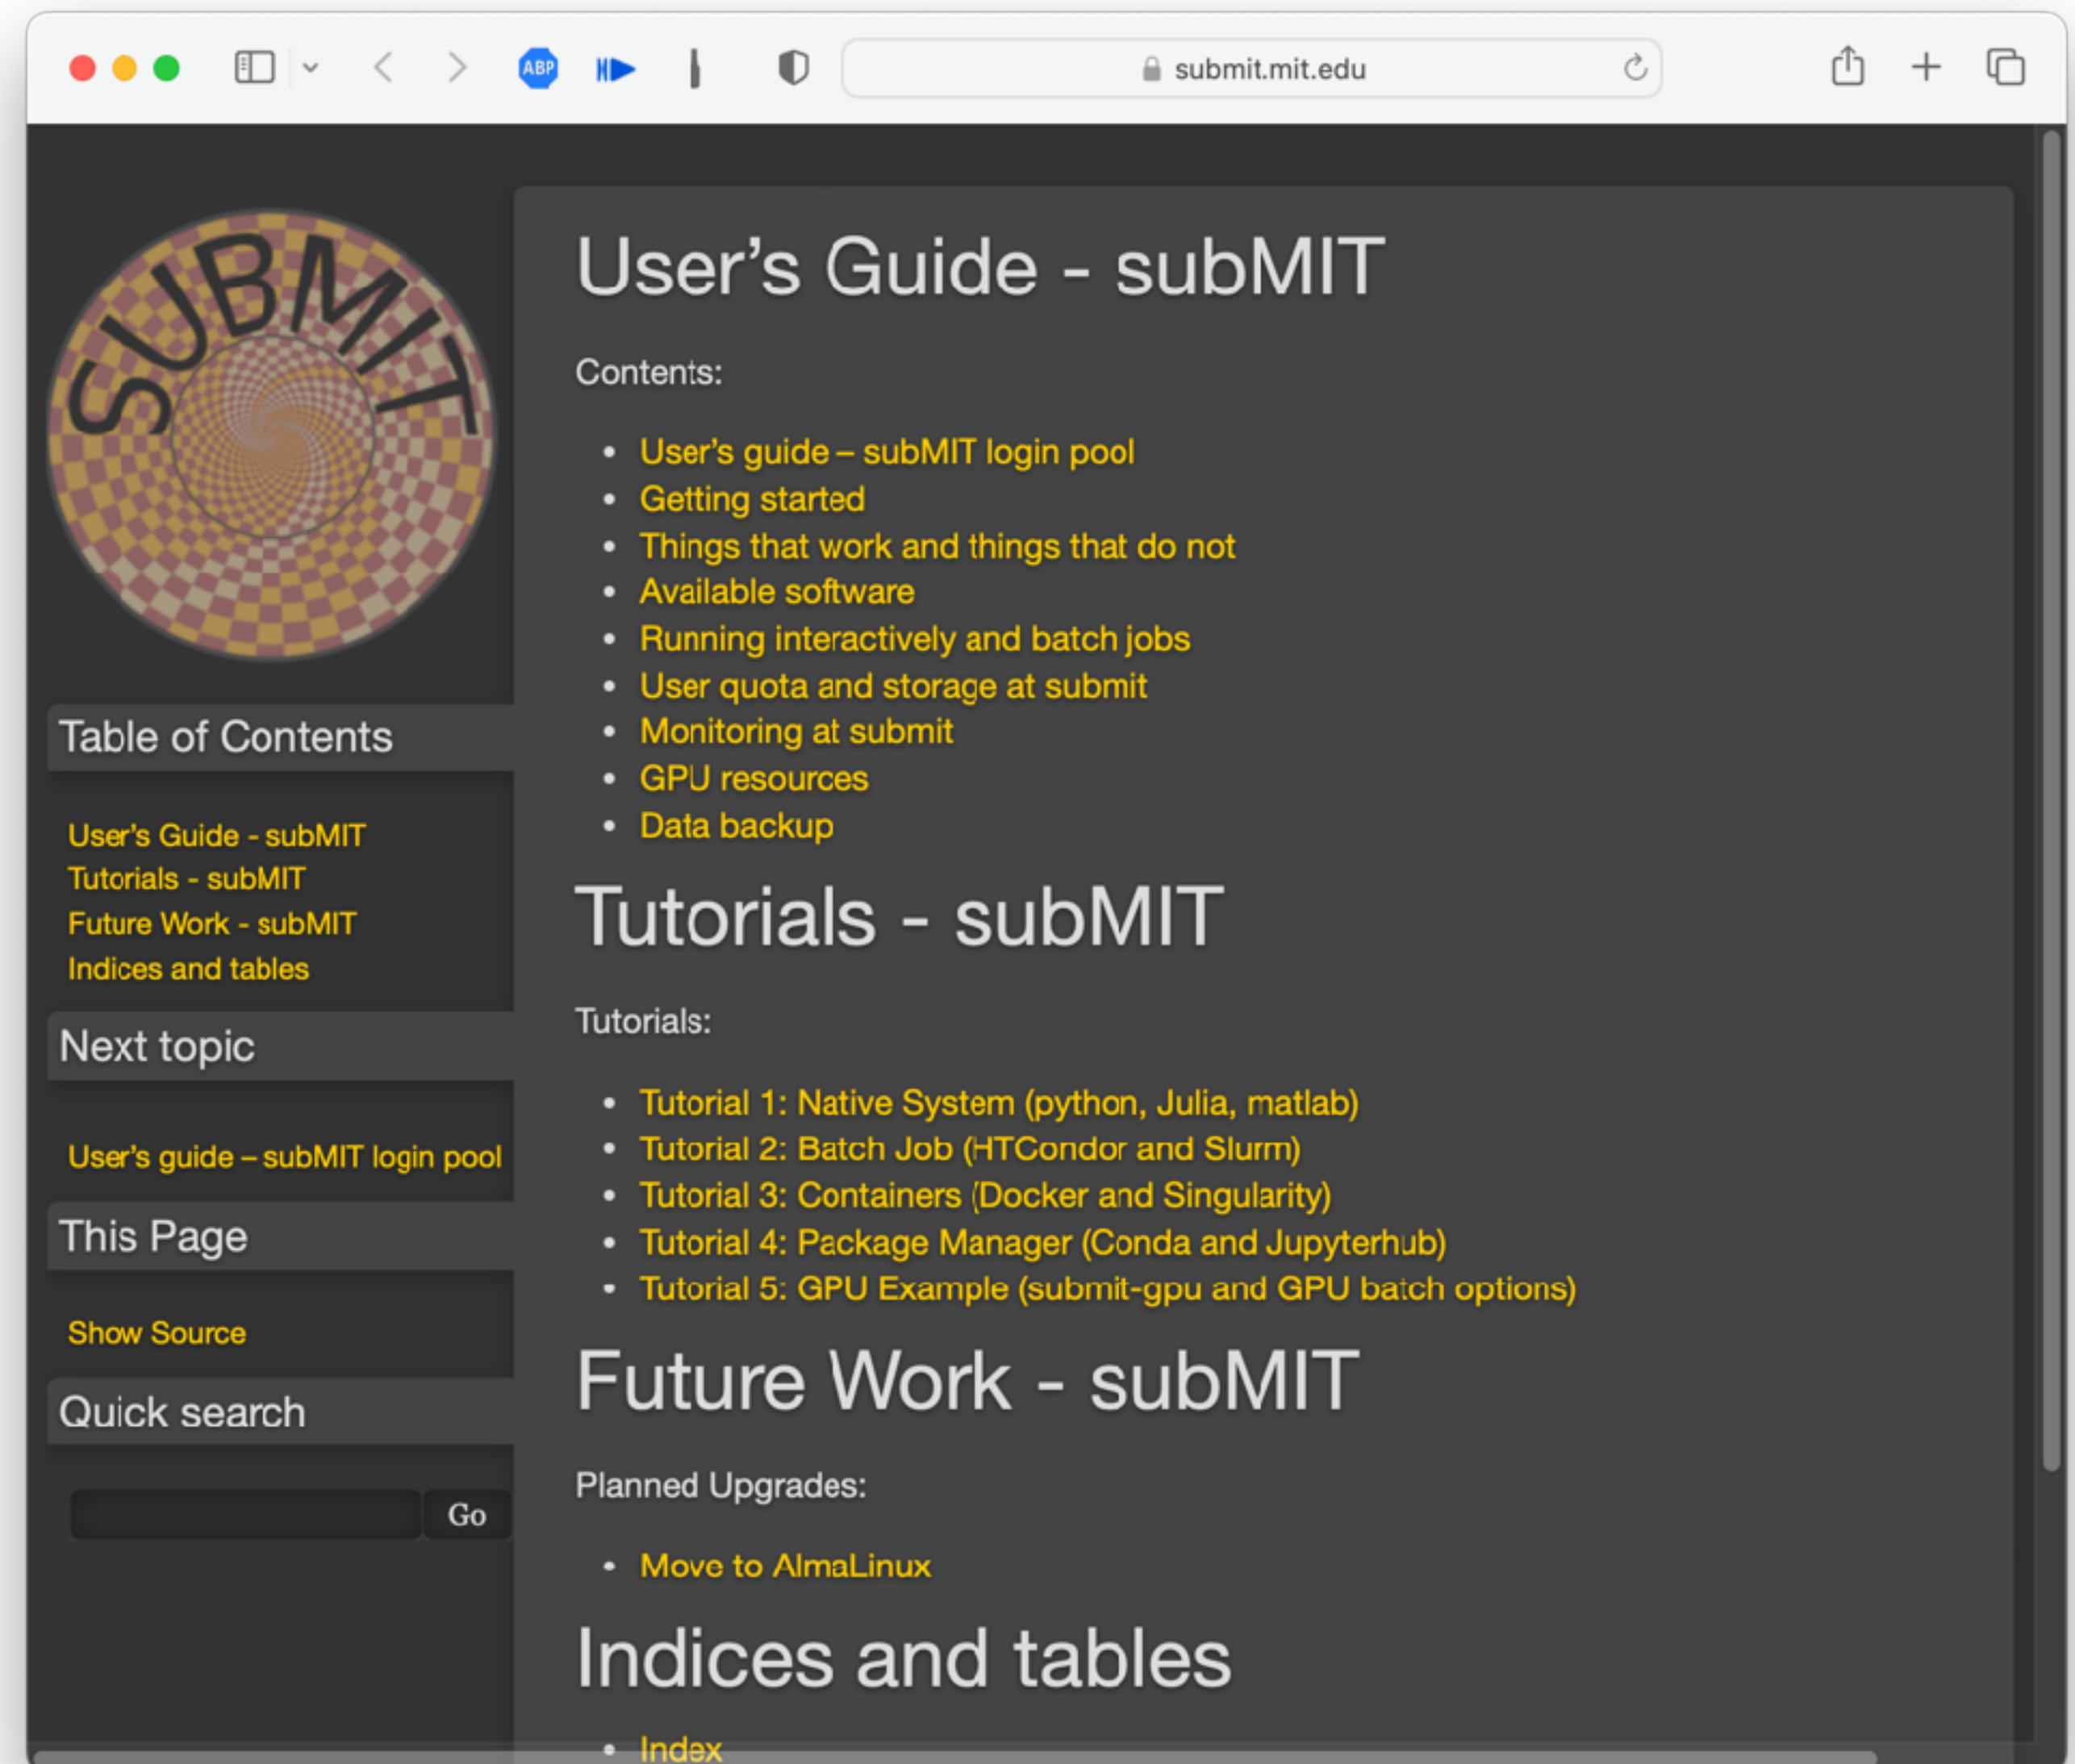
MKI Identified as Under-Served

- MKI was identified as being **under-served by subMIT**, with all team members being outside of the group.
- Myself and Matt were brought on to assist in **integrating the MKI community** with subMIT, and helping tailor subMIT to their workflow.



Towards a comprehensive computing service

- Previous subMIT efforts have focused on:
 - Hardware
 - User support
 - Creation of a users guide
 - A series of users group meetings

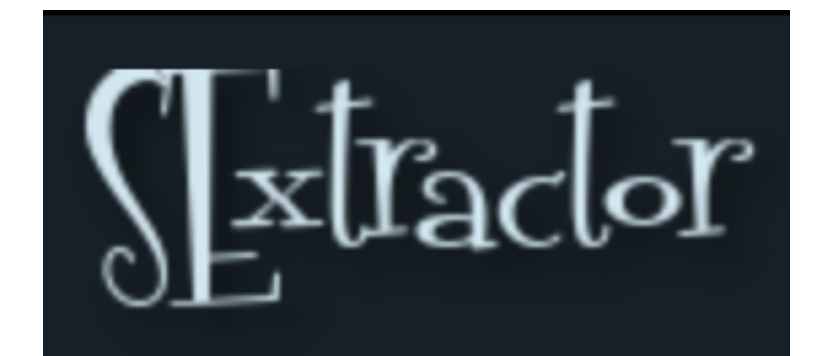


The missing pieces until now

- A few things were missing from a comprehensive computing service:
 - One to one support
 - Understanding of user-base and requirements (see Matt's talk also)
 - User onboarding
 - Community building
 - User training

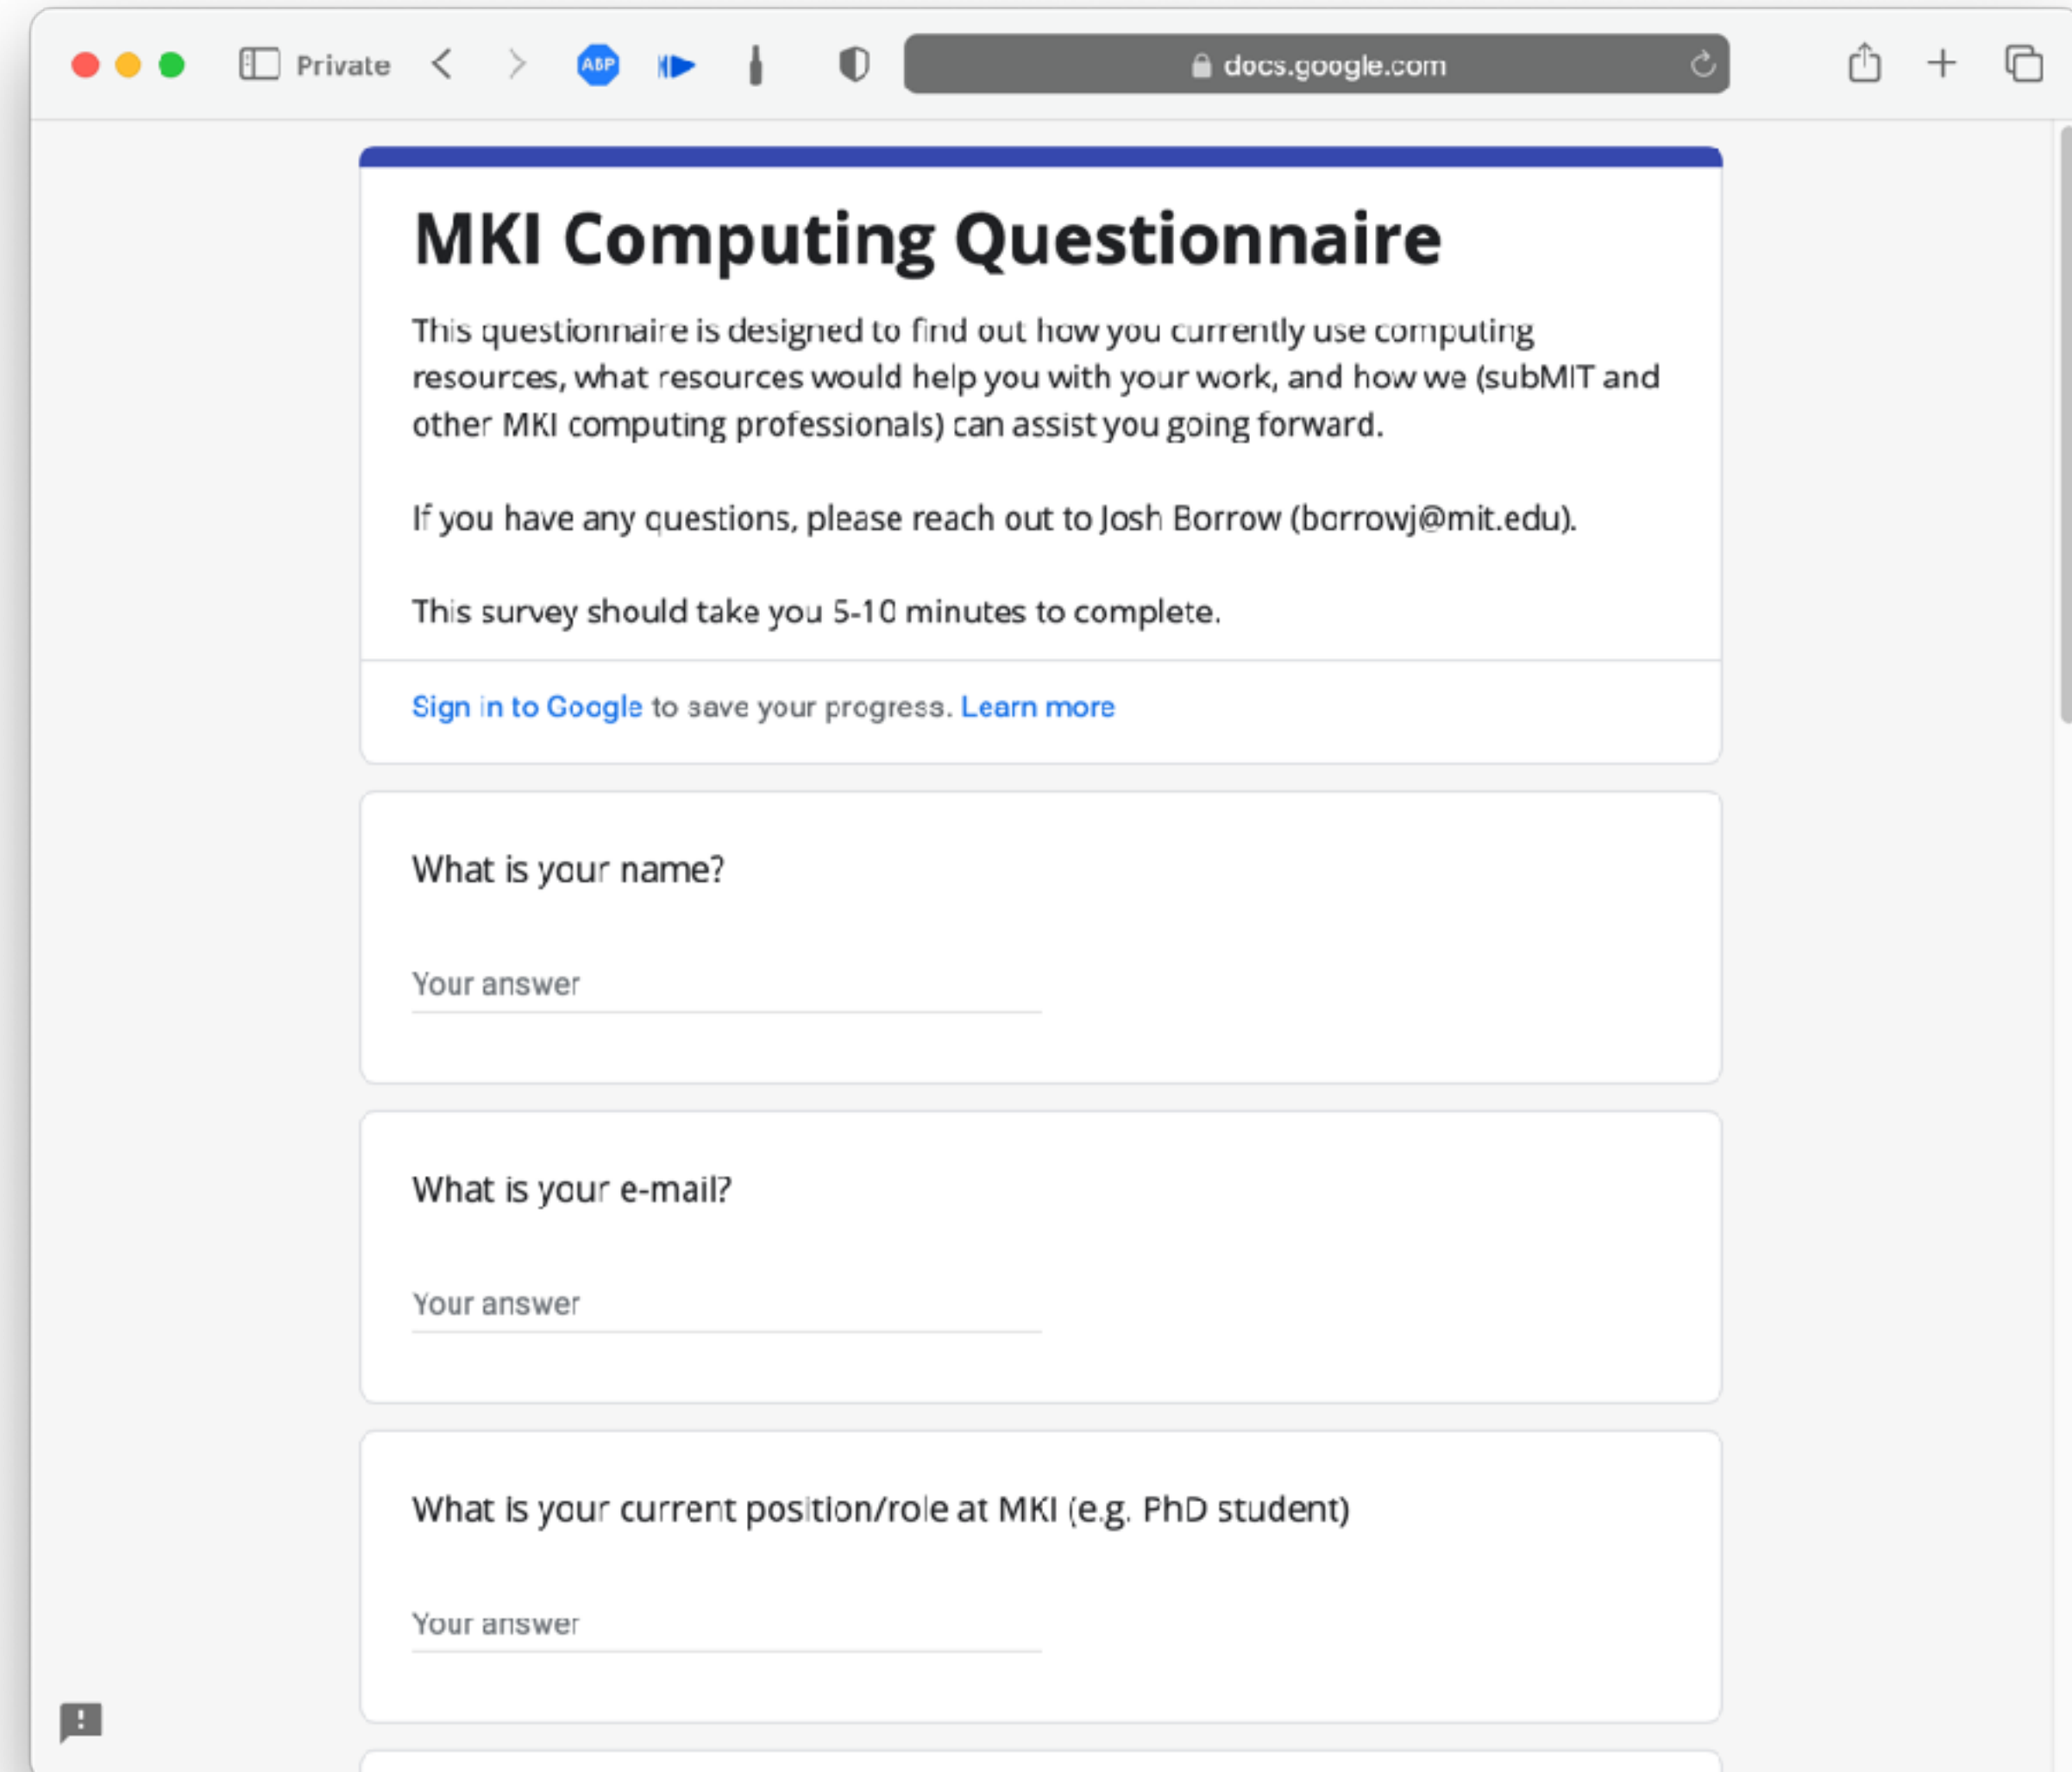
One-to-one user support

- Some users may need more than one-off e-mails.
- To assist with onboarding and to direct users to the most appropriate resources, within MKI I have been providing **longer-term one-to-one support for researchers.**
- Example: Javier Viana & Mariona Agusti in MKI setting up their GPU accelerated ML pipeline.



Surveying MKI

- Survey was sent out on Monday 3rd April 2023. Results were closed on the 28th April 2023. We received **22 responses**, with a **good spread throughout the hierarchy within MKI**:
 - 5 Graduate Students
 - 4 Postdocs
 - 11 Research Scientist/Other Senior Staff
 - 2 Faculty
- There was a wide range of research areas that responded, from instrumentation development to galaxy formation theory.




The screenshot shows a web browser window displaying a Google Docs document titled "MKI Computing Questionnaire". The document content includes:

- MKI Computing Questionnaire**
- This questionnaire is designed to find out how you currently use computing resources, what resources would help you with your work, and how we (subMIT and other MKI computing professionals) can assist you going forward.
- If you have any questions, please reach out to Josh Borrow (borrowj@mit.edu).
- This survey should take you 5-10 minutes to complete.
- [Sign in to Google](#) to save your progress. [Learn more](#)
- What is your name?
Your answer _____
- What is your e-mail?
Your answer _____
- What is your current position/role at MKI (e.g. PhD student)
Your answer _____

Key survey results

- Up to **90% of respondents work is appropriate for deployment on subMIT**. Despite this, **only 27% of respondents have actually used subMIT**.
- Users request **more training** (e.g. how to use Visual Studio Code with subMIT) and **education** on the system, and changes to the users guide.
- From an **ongoing resource acquisition standpoint**, it appears **many users are CPU and storage limited**.

What was your familiarity with subMIT before taking this survey?
22 responses




- I am already aware of subMIT & use it.
- I am already aware of subMIT & do not use it.
- I was not aware of subMIT

Major barriers to subMIT are mainly 'human', not hardware!

How many cores do you use to perform your current work?

22 responses




- 1-16 cores (e.g. laptop, desktop)
- 16-128 cores (e.g. a workstation, subMIT)
- 128-1024 cores
- 1024+ cores

Major barriers to subMIT are mainly 'human', not hardware!

Which of the following would make subMIT better suit your own needs?

13 responses



User onboarding

- subMIT provided a **user support and onboarding IAP course** for MKI and the wider physics community in early February (around 2 hours).
- This was **very well attended** (40-50 participants in a hybrid mode) and **well received within the department**.
- Led to a number of **follow ups within MKI** for **additional onboarding** assistance.

From Laptops to High
Performance Computing: Low-
Hanging Fruit in Parallelization

Matt Heine

mheine@mit.edu

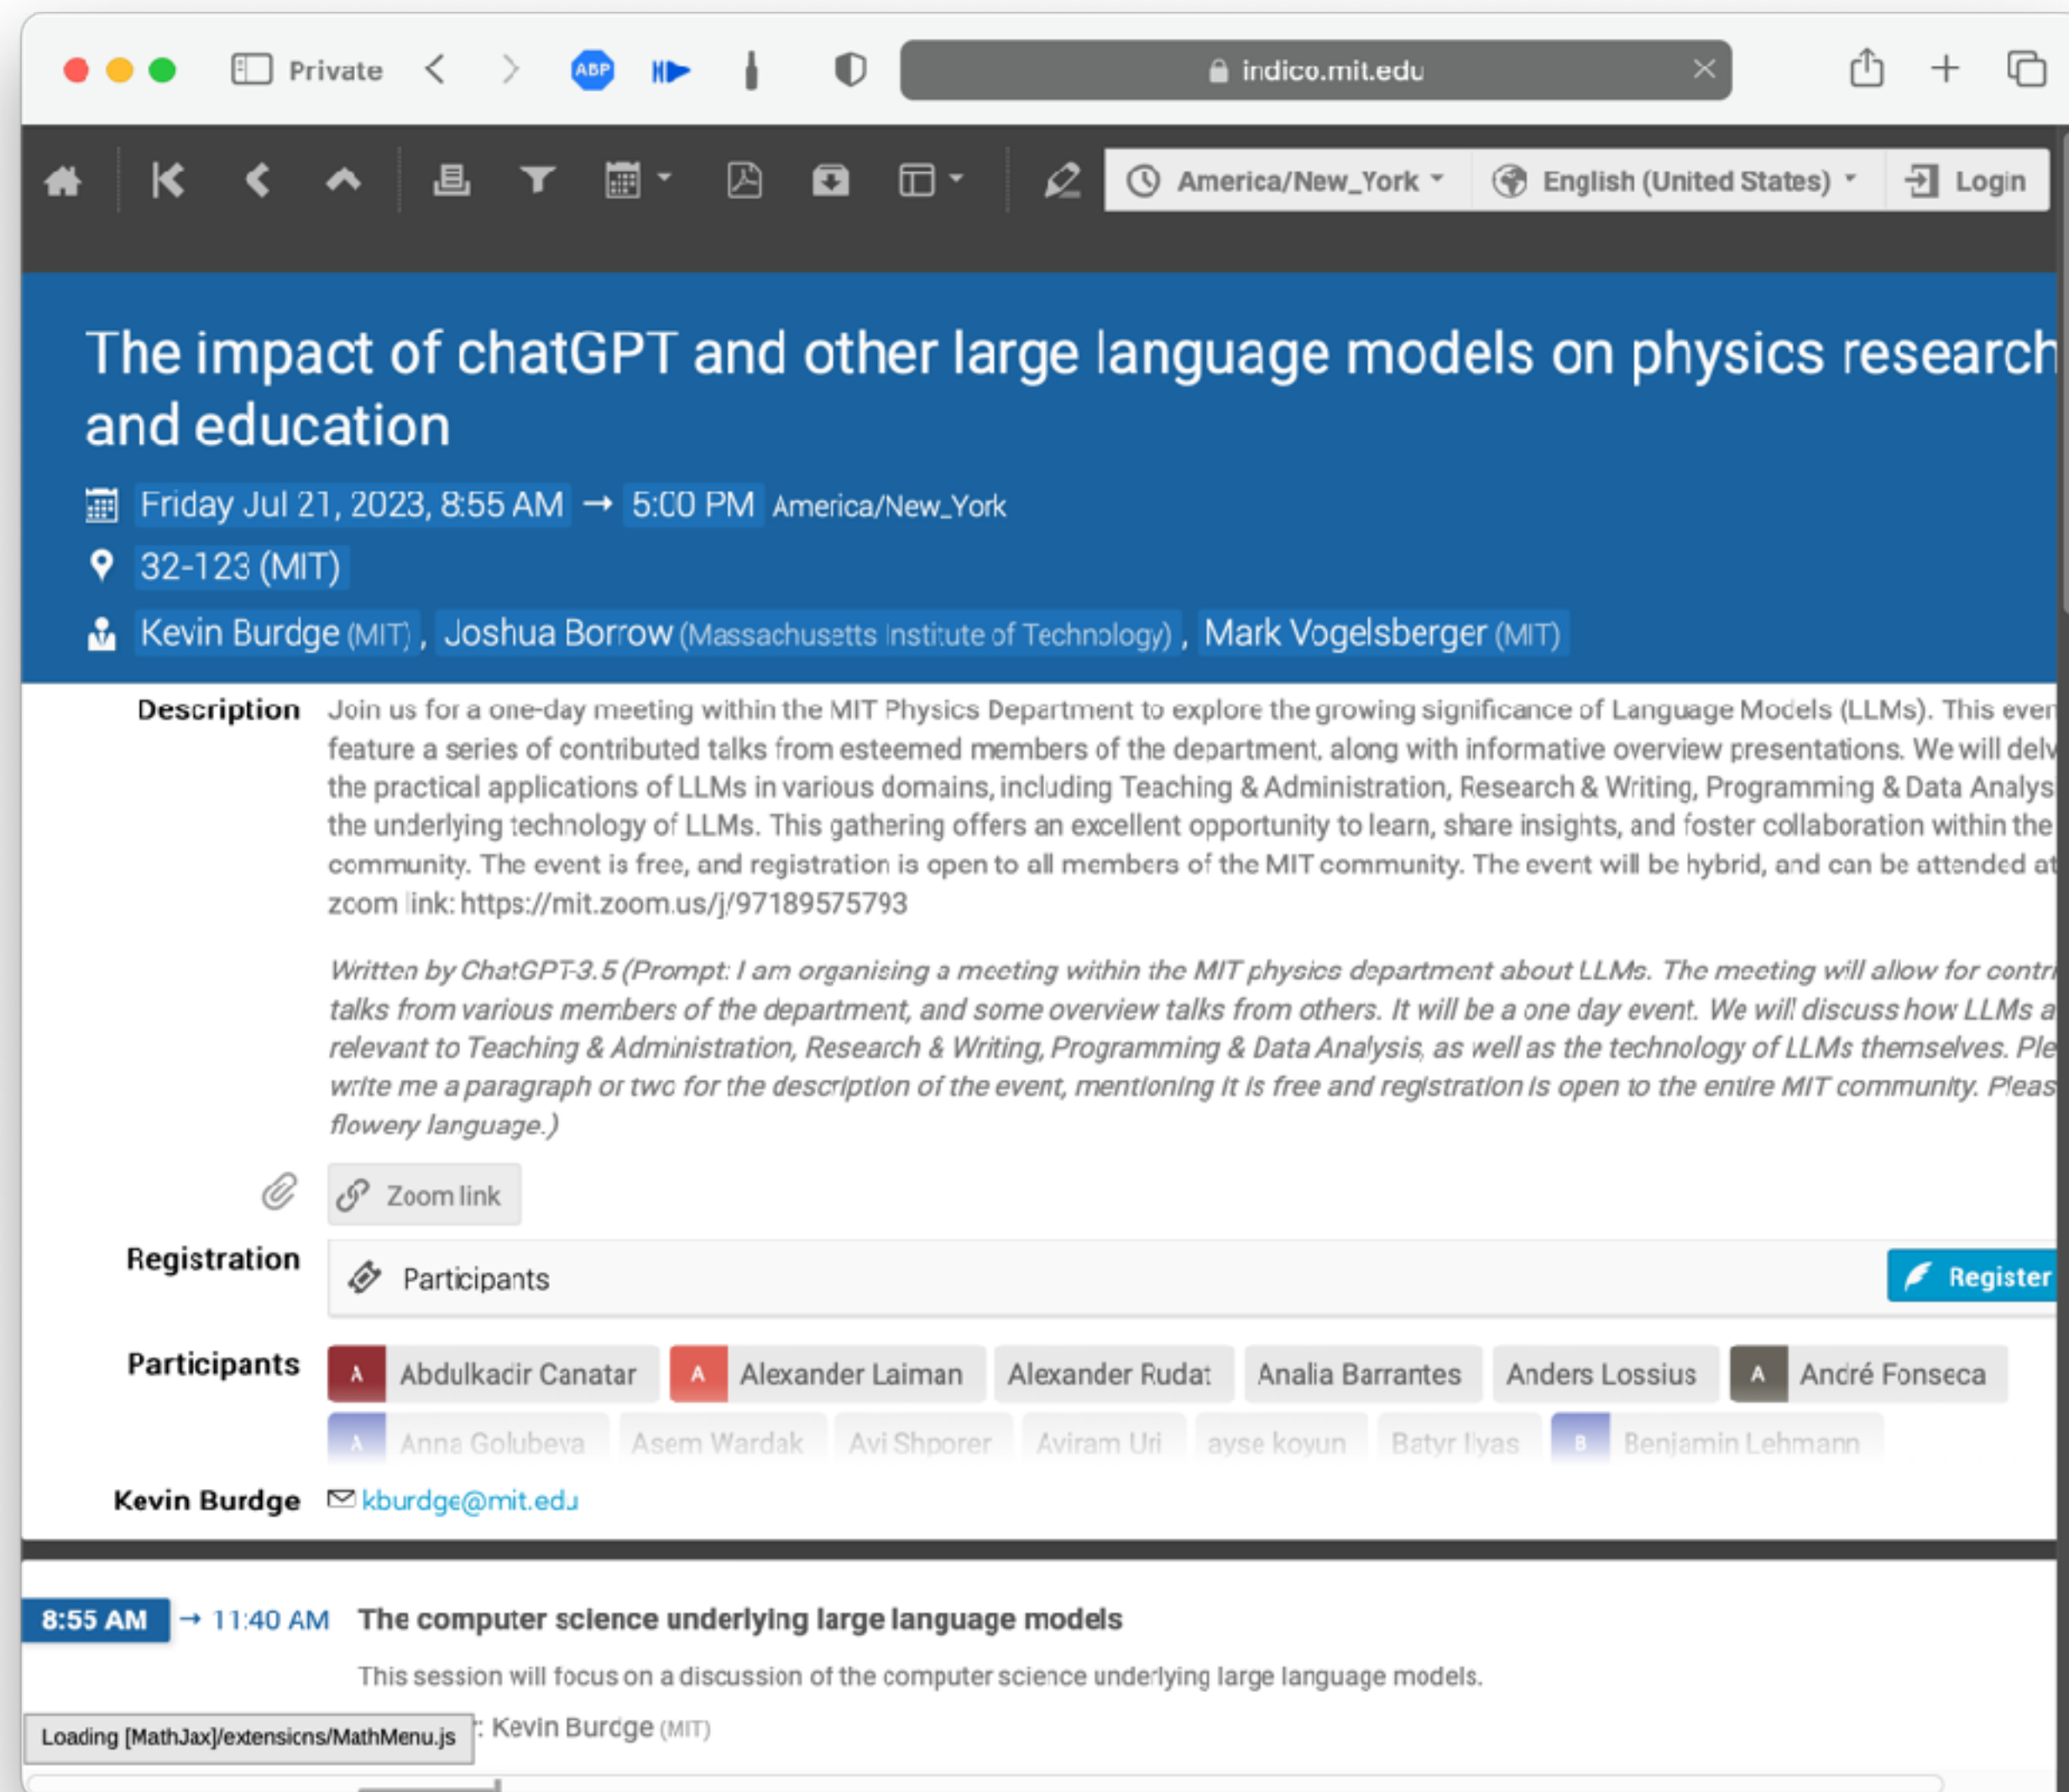
37-410: Wednesdays (all day), Thursdays (afternoon)

submit.mit.edu



Community building

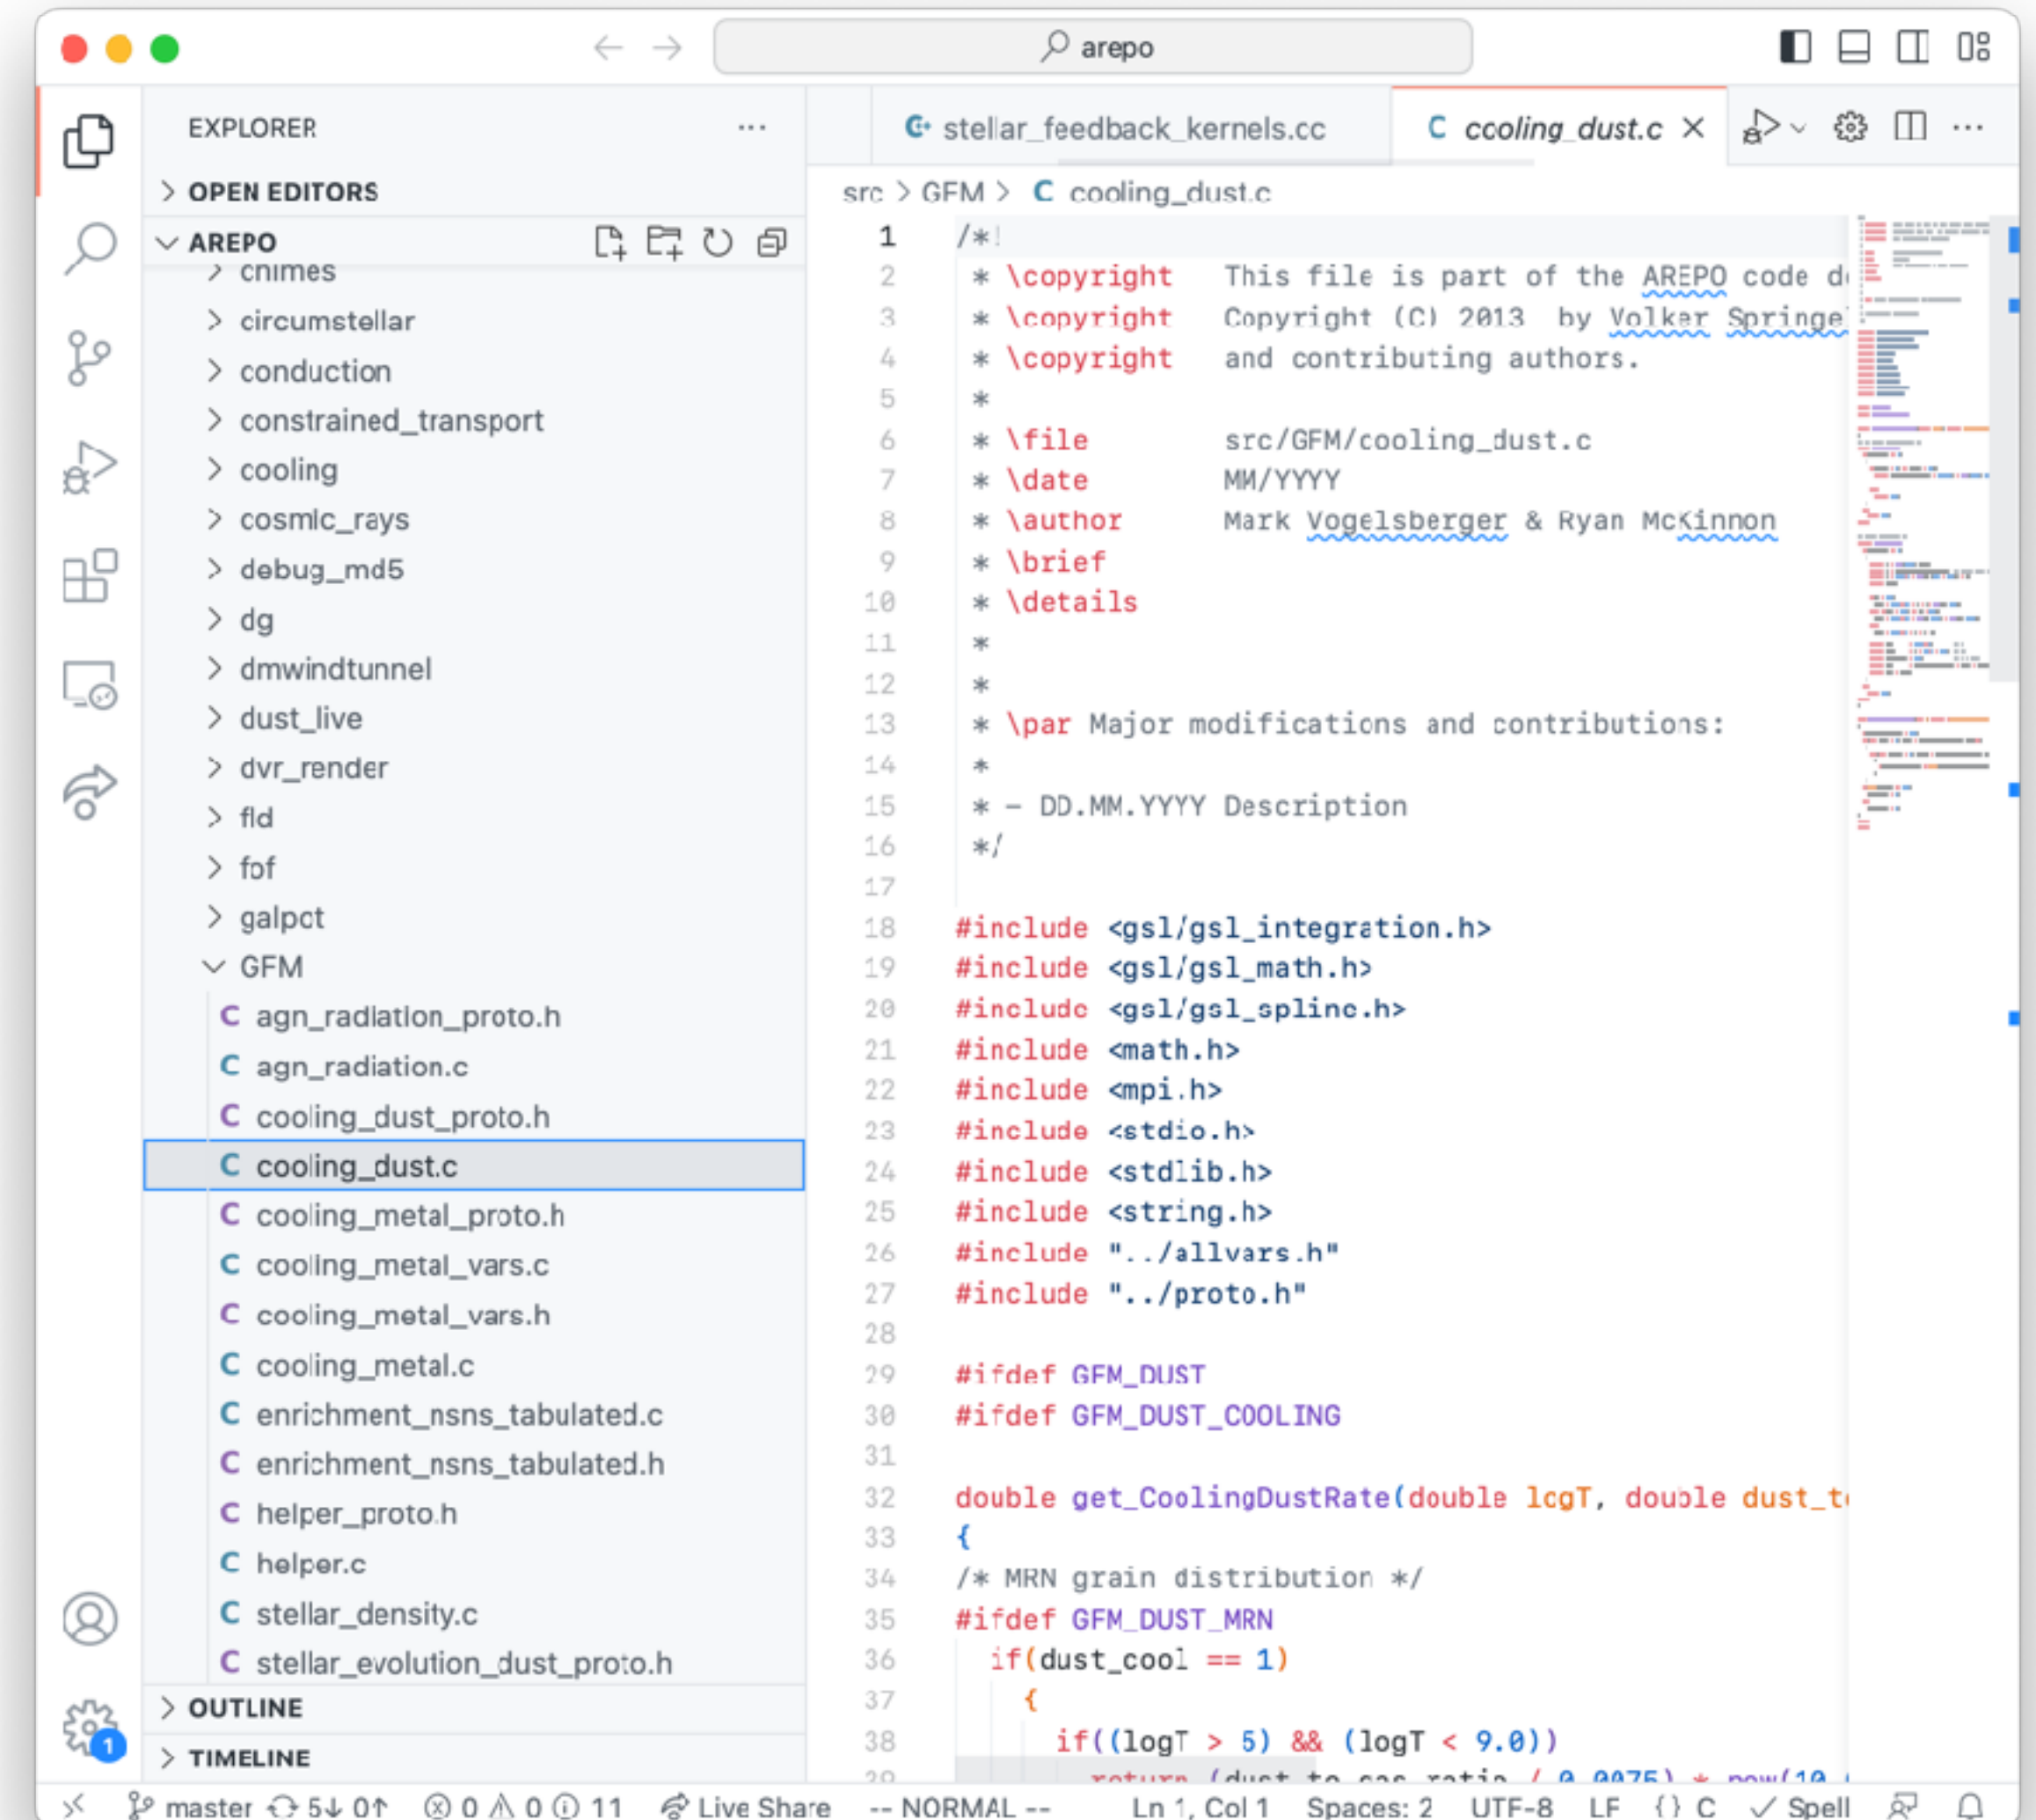
- MKI is leading a department-wide computing meeting on the use of **Large Language Models**.
- Looking forward to developing future user-focused computing meetings through subMIT.



The screenshot shows a web browser window displaying an Indico event page. The browser's address bar shows 'indico.mit.edu'. The event title is 'The impact of chatGPT and other large language models on physics research and education'. The event is scheduled for Friday, July 21, 2023, from 8:55 AM to 5:00 PM in the America/New_York time zone, at room 32-123 (MIT). The organizers listed are Kevin Burdge (MIT), Joshua Borrow (Massachusetts Institute of Technology), and Mark Vogelsberger (MIT). The description states that the event is a one-day meeting for the MIT Physics Department to explore the significance of LLMs, featuring contributed talks and overview presentations. It is free and open to all members of the MIT community. A Zoom link is provided: https://mit.zoom.us/j/97189575793. A note mentions that the description was written by ChatGPT-3.5. The registration section shows a 'Participants' list with a 'Register' button. The participants listed include Abdulkadir Canatar, Alexander Laiman, Alexander Rudat, Analia Barrantes, Anders Lossius, Anré Fonseca, Anna Golubeva, Asem Wardak, Avi Shporer, Aviram Uri, ayse koyun, Batyr Iyas, and Benjamin Lehmann. The event is organized by Kevin Burdge (kburdge@mit.edu). At the bottom, a session titled 'The computer science underlying large language models' is shown, scheduled for 8:55 AM to 11:40 AM, with a description: 'This session will focus on a discussion of the computer science underlying large language models.' A loading message at the bottom indicates 'Loading [MathJax]/extensions/MathMenu.js' by Kevin Burdge (MIT).

User training

- Based upon **user feedback**, we are developing a **user training session** for the use of **Visual Studio Code (VSC)**.
- The focus of the session will be both using **VSC locally** and **remotely through subMIT** and other MIT machines.
- **Training planned** for in **Sept/Oct** once new graduate students join.



```
src > GFM > C cooling_dust.c
1  /*
2  * \copyright This file is part of the AREPO code d
3  * \copyright Copyright (C) 2013 by Volker Springe
4  * \copyright and contributing authors.
5  *
6  * \file      src/GFM/cooling_dust.c
7  * \date      MM/YYYY
8  * \author    Mark Vogelsberger & Ryan McKinnon
9  * \brief
10 * \details
11 *
12 *
13 * \par Major modifications and contributions:
14 *
15 * - DD.MM.YYYY Description
16 */
17
18 #include <gsl/gsl_integration.h>
19 #include <gsl/gsl_math.h>
20 #include <gsl/gsl_spline.h>
21 #include <math.h>
22 #include <mpi.h>
23 #include <stdio.h>
24 #include <stdlib.h>
25 #include <string.h>
26 #include "../allvars.h"
27 #include "../proto.h"
28
29 #ifdef GFM_DUST
30 #ifdef GFM_DUST_COOLING
31
32 double get_CoolingDustRate(double logT, double dust_t
33 {
34 /* MRN grain distribution */
35 #ifdef GFM_DUST_MRN
36     if(dust_cool == 1)
37     {
38         if((logT > 5) && (logT < 9.0))
39             return (dust_to_gas_ratio / 0.0075) * pow(10,
```

Conclusions

- subMIT has been actively engaging with the MKI & wider Physics community, providing **direct user support, onboarding of users, and targeted training sessions.**
- **MKI's research activities are very suitable for subMIT's architecture.**
- **More community engagement is needed, but the future looks very bright for MKI research on the subMIT facility.**