

# ML for Fundamental Physics School 2026



National Energy Research  
Scientific Computing Center

June 1, 2026

<https://indico.global/event/17000/>



Georgia Institute  
of Technology



**Organising committee:**  
Aishik Ghosh (GT), Yifan Chen (SLAC, Stanford), Elham E Khoda (UBC), Dennis Noll (SLAC, Stanford),  
Melissa Quinnan (UCSD)

**Steering committee:**  
Sascha Diefenbacher (Heidelberg University), Steven Farrell (NERSC/LBNL), Aishik Ghosh, (GT), Shih-Chieh Hsu  
(UW), Elham E Khoda (UBC), Benjamin Nachman (SLAC, Stanford), Daniel Whiteson (UCI)



U.S. DEPARTMENT OF  
**ENERGY**

UC San Diego

**SLAC** NATIONAL  
ACCELERATOR  
LABORATORY

**OAK RIDGE**  
National Laboratory

# Welcome to Atlanta

---



Georgia Tech located in midtown Atlanta:

- Subway: **Midtown/North Avenue Marta station**
- Free **stinger bus red route** from accommodation to Marcus Nanotechnology Research Center (**TransLoc app** for live location of busses)
- Coffee / snacks will be served during breaks
- **Lunch / dinner on your own** (student center / delivery bots / Stinger to Tech Square)
- Thursday dinner at nanotech atrium

# Georgia Tech Student Volunteers

---



**Nahreen Shamon Ayala**



**Chuanxin Wang**

They're here to help you navigate!

# Things to do in the city



Beltline



Food



Music, art, history



# Things to do in the city .. if it rains



Shakespeare Tavern



Food



Music, art, history



# ML4FP 2026

---



Georgia Tech College of Sciences  
School of Physics



U.S. DEPARTMENT OF  
**ENERGY**



National Energy Research  
Scientific Computing Center



- Expanded program thanks to financial and institutional support from Georgia Tech School of Physics
- Financial support from DoE, US ATLAS
- Computing support from NERSC
- Institutional support from Oak Ridge National Lab
- Speakers for taking the time to prepare tutorials
- **You all! For participating in ML4FP 2026!**

Special thanks: Teresa Kibler, Robin Finey, Tyercia Nevels, Claudete Richardson

# Introducing the organizers in the room

---


- **Melissa Quinnan (CMS liaison)**
- **Dennis Noll (ATLAS liaison)**
- **Yifan Chen (Neutrino liaison)**
- **Elham E Khoda**
- **Aishik Ghosh**
- **Also Steve Farrell and his Data & AI Services Group at NERSC!**

# Tell us about you

---


Pre-school survey: [link](#)

Help us know you better, tune the school to your background



# Agenda Day 1 – Monday, 6/1

## Introduction of ML Fundamental concepts and hands-on neural network tutorials

MON, JUNE 1		
8:45 AM → 9:00 AM	<b>Welcome</b>	🕒 15m ✎
	<b>Speakers:</b> Aishik Ghosh (Georgia Institute of Technology), Dennis Noll (Stanford University), Elham Khoda (University of British Columbia, Vancouver), Melissa Quinnan (Univ. of California San Diego (US)), Yifan Chen (SLAC National Accelerator Laboratory (US))	
9:00 AM → 10:30 AM	<b>Introduction to Machine Learning: Part I</b>	🕒 1h 30m ✎
	<b>Speaker:</b> Melissa Quinnan (Univ. of California San Diego (US))	
10:30 AM → 11:00 AM	<b>Coffee Break</b>	🕒 30m
11:00 AM → 12:30 PM	<b>Tutorial: Intro to Neural Networks I</b>	🕒 1h 30m ✎
	<b>Speakers:</b> Elham Khoda (University of British Columbia, Vancouver), Roger Huang (Lawrence Berkeley National Lab)	
12:30 PM → 2:00 PM	<b>Lunch</b>	🕒 1h 30m
2:00 PM → 3:30 PM	<b>Introduction to Machine Learning: Part II</b>	🕒 1h 30m ✎
	<b>Speaker:</b> Dennis Noll (Stanford University)	
	 260601_noll-hep_...	
3:30 PM → 4:00 PM	<b>Coffee Break</b>	🕒 30m
4:00 PM → 5:30 PM	<b>Tutorial: Intro to Neural Networks II</b>	🕒 1h 30m ✎
	<b>Speakers:</b> Elham Khoda (University of British Columbia, Vancouver), Roger Huang (Lawrence Berkeley National Lab)	

# Agenda Day 2 – Tuesday, 6/2

## Overview of ML in HEP (collider + neutrino); Differentiable programming; Simulation-based inference

TUE, JUNE 2		
9:00 AM → 9:40 AM	<b>Overview of ML in collider physics</b> Speaker: Aishik Ghosh (Georgia Institute of Technology)	🕒 40m 📄
9:40 AM → 9:45 AM	Break	🕒 5m
9:45 AM → 10:30 AM	<b>Overview of ML in neutrino physics</b> Speaker: Yifan Chen (SLAC National Accelerator Laboratory (US))	🕒 45m 📄
10:30 AM → 11:00 AM	Coffee Break	🕒 30m
11:00 AM → 12:30 PM	<b>Differentiable Programming</b> Speaker: Yifan Chen (SLAC National Accelerator Laboratory (US))	🕒 1h 30m 📄
12:30 PM → 2:00 PM	Lunch	🕒 1h 30m
2:00 PM → 3:30 PM	<b>Simulation-based Inference: Part I</b> Speakers: Aishik Ghosh (Georgia Institute of Technology), Eddie Mcgrady (University of Notre Dame (US))	🕒 1h 30m 📄
3:30 PM → 4:00 PM	Coffee Break	🕒 30m
4:00 PM → 5:30 PM	<b>Simulation-based Inference: Part II</b> Speakers: Aishik Ghosh (Georgia Institute of Technology), Eddie Mcgrady (University of Notre Dame (US))	🕒 1h 30m 📄

# Agenda Day 3 – Wednesday, 6/3






## Generative Models; Anomaly Detection; Neutrino Parallel Session

Free time to explore city for non-neutrino participants (after 3:30 pm)

WED, JUNE 3		
9:00 AM → 10:30 AM	<b>Generative Models: Part I</b> Speaker: Sofia Palacios Schweitzer (Rutgers University)	🕒 1h 30m 📄
10:30 AM → 11:00 AM	Coffee Break	🕒 30m
11:00 AM → 12:30 PM	<b>Generative Models: Part II</b> Speaker: Sofia Palacios Schweitzer (Rutgers University)	🕒 1h 30m 📄
12:30 PM → 2:00 PM	Lunch	🕒 1h 30m
2:00 PM → 3:30 PM	<b>Anomaly Detection (Lecture &amp; Tutorial)</b> Speaker: Melissa Quinnan (Univ. of California San Diego (US))	🕒 1h 30m 📄
3:30 PM → 4:00 PM	Coffee Break	🕒 30m
4:00 PM → 5:30 PM	<b>Neutrino Parallel Session (experiment specific)</b> Speakers: Roger Huang (Lawrence Berkeley National Lab), Yifan Chen (SLAC National Accelerator Laboratory (US))	🕒 1h 30m 📄

# Agenda Day 4 – Thursday, 6/4

## Transformers; Muon Collider; Career Panel; Conference Dinner

THU, JUNE 4		
9:00 AM → 10:30 AM	<b>Transformers: Part I</b> Speaker: Aaron Wang (University of Illinois Chicago (US))	🕒 1h 30m 
10:30 AM → 11:00 AM	Coffee Break	🕒 30m
11:00 AM → 12:00 PM	<b>Transformers: Part II</b> Speaker: Aaron Wang (University of Illinois Chicago (US))	🕒 1h 
12:00 PM → 1:30 PM	Lunch	🕒 1h 30m
1:30 PM → 3:00 PM	<b>Application: Muon Collider BIB (Lecture &amp; Tutorial)</b> Speakers: Angira Rastogi (Lawrence Berkeley National Lab. (US)), Samuel Ferraro (Harvard University)	🕒 1h 30m 
3:00 PM → 3:30 PM	Coffee Break	🕒 30m
3:30 PM → 5:00 PM	<b>Career Panel</b> Speakers: Dr Ben Brubaker (Quanta), Dr Rachel Smith (Xanadu), Dr Jake Daughhete (Oak Ridge National Laboratory), Prof. Mariel Pettee (UW Madison)	🕒 1h 30m 
5:50 PM → 7:50 PM	<b>Conference Dinner: Special message from Chair of Physics</b> Speaker: Prof. Feryal Ozel (Georgia Institute of Technology)	🕒 2h 

# Exciting Career Panel: Send us your questions!

---



[Ben Brubaker](#): Physicist turned journalist at Quanta Magazine



[Mariel Pettee](#): New faculty doing AI and physics at UW Madison



[Rachel Smith](#): Physicist turned AI Specialist at Xanadu



[Jake Daughhete](#): Neutrino physicist turned radiation imaging, detector and simulation scientist at Oak Ridge National Laboratory

**Submit Your Questions (until Thursday):**



[Link](#)

# Agenda Day 5 – Friday, 6/5


## Efficient ML; ATLAS / CMS parallel sessions; Free afternoon for non-collider participants **Closeout**

FRI, JUNE 5		
9:00 AM → 10:30 AM	<b>Efficient Machine Learning (Lecture &amp; Tutorial)</b> Speaker: Elham Khoda (University of British Columbia, Vancouver)	🕒 1h 30m 📄
10:30 AM → 11:00 AM	<b>Coffee Break</b>	🕒 30m
11:00 AM → 12:30 PM	<b>ATLAS Parallel Session: Part I</b>	🕒 1h 30m 📄
11:00 AM → 12:30 PM	<b>CMS Parallel Session: Part I</b> 🔗 Will be held in College of Computing Building (801 Atlantic Dr NW, Atlanta, GA 30332), ground floor seminar room. Short walk from nanotech building	🕒 1h 30m 📍 103 (CCB) 📄
12:30 PM → 2:00 PM	<b>Lunch</b>	🕒 1h 30m
2:00 PM → 3:30 PM	<b>ATLAS Parallel Session: Part II</b>	🕒 1h 30m 📄
2:00 PM → 3:30 PM	<b>CMS Parallel Session: Part II</b> Will be held in College of Computing Building (801 Atlantic Dr NW, Atlanta, GA 30332), ground floor seminar room. Short walk from nanotech building	🕒 1h 30m 📍 103 (CCB) 📄
3:30 PM → 4:00 PM	<b>Coffee Break</b>	🕒 30m
4:00 PM → 4:30 PM	<b>Closeout Session: Recap and Networking</b>	🕒 30m 📄

CMS parallel session  
will be in a different  
building!  
Short walk from here

# Internet Connection

---



# NERSC Computing Accounts



National Energy Research  
Scientific Computing Center

- Make sure you have the NERSC account ready!
- Check [this instruction](#) to set up an interactive session.
- Please manually close your jupyter session after the use.

## Tutorial directions (Day 1)

	Login Node	Shared GPU Node	Exclusive CPU Node	Exclusive GPU Node	Configurable Job
Permutter	<input type="button" value="start"/>	<input type="button" value="start"/>	<input type="button" value="start"/>	<input type="button" value="start"/>	<input type="button" value="start"/>
Resources	Use a login node shared with other users, outside the batch queues.	Use a single GPU on a node within a job allocation using defaults.	Use your own node within a job allocation using defaults.		Use multiple compute nodes with specialized settings.

- Hop into <https://jupyter.nersc.gov/>
- Request a “Configurable Job” and choose:
  - a. Account: **ntrain6** (should be only option)
  - b. Reservation: **m14fp2026\_day1** (dropdown)
  - c. **cpus=32, gpus=1, ntasks-per-node=1**
  - d. **Make sure you use QOS = jupyter\_shared**
- **Hit start**

jupyterhub Home Token Services Documentation

### Server Options

Account (“\_g” suffix will be added as needed):

Constraint:

gpu

QOS:

cpus-per-task (GPU node has 128 cpus, CPU node has 256 cpus):

gpus-per-task (node has 4 GPUs):

nodes (maximum of 4 for jupyter QOS):


ntasks-per-node:

Reservation:

time (time limit in minutes):

# Tutorial Github

---



Tutorial GitHub: <https://github.com/ml4fp/2026-gatech>

- All tutorial and demos will be shared via github
- Please refresh the repo before every new tutorial sessions

# Slack

---



- Join the slack workspace with [this magic link](#)
- Important information in **#announcement**
- Join **#discussion** to discuss the content of the lectures
- Remote participants should ask questions on **#discussion**, we will answer or bring up with speaker
- Join **#social** for 🥨🌮🍜📺🍦🍩🥤🍺🚗🏀🏃🏔️🌉🌉

# Lunch

## 1.5 hours lunch break Please try to be back in time

- Starship robot food delivery app
  - Order well ahead to avoid delays
- Go to [John Lewis Student Center](#) (10 min walk)
- Ask student volunteers for advice!



# Code of Conduct

---

The school intends to be **inclusive and respectful for all**.  
Please abide [the IRIS-HEP code of conduct](#)

- Foster an open, productive, harassment-free environment for everyone.
- Be welcoming and support people of all backgrounds and identities, immigration status, social and economic class, educational level, sex, sexual orientation, gender identity and expression, age, physical appearance, family status, technological choices, academic discipline, political views, religion, mental ability, and physical ability.
- Be considerate. Your work will be used by other people, and you in turn will depend on the work of others. Any decision you take will affect users and colleagues, and you should take those consequences into account when making decisions. Remember that we're a world-wide community. You may be communicating with someone with a different primary language or cultural background.
- Be respectful. Not all of us will agree all the time, but disagreement is no excuse for poor behavior or poor manners. We might all experience some frustration now and then, but we cannot allow that frustration to turn into a personal attack. It's important to remember that a community where people feel uncomfortable or threatened is not a productive one.
- Respect the work of others. We recognize the acknowledgment/citation requests of the original authors. As authors, we are explicit about how we want our own work to be cited or acknowledged.
- Be considerate in the words you choose. Be kind to others. Do not insult or put down other community members. Harassment and other exclusionary behavior are not acceptable.
- When we disagree, try to understand why. Disagreements, both social and technical, happen all the time and the IRIS-HEP community is no exception. Try to understand where others are coming from, as seeing a question from their viewpoint may help find a new path forward. And don't forget that it is human to err: blaming each other doesn't get us anywhere, while we can learn from mistakes to find better solutions.
- A simple apology can go a long way. It can often de-escalate a situation, and telling someone that you are sorry is an act of empathy that doesn't automatically imply an admission of guilt.