

Computing support and Infrastructure for Academic Scientific Research at the Mellon College of Science

OCTOBER 22, 2025

Nuno Chagas

Agenda

- Computing support
- Infrastructure
- Academic Scientific Research
- Mellon College of Science

Agenda

- 1. Mellon College of Science
- 2. Academic Scientific Research
- 3. Computing support
- 4. Infrastructure

Mellon College of Science

4 departments

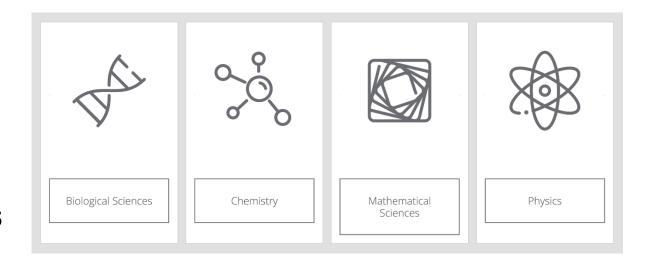
1000+ undergrads

450+ grad students

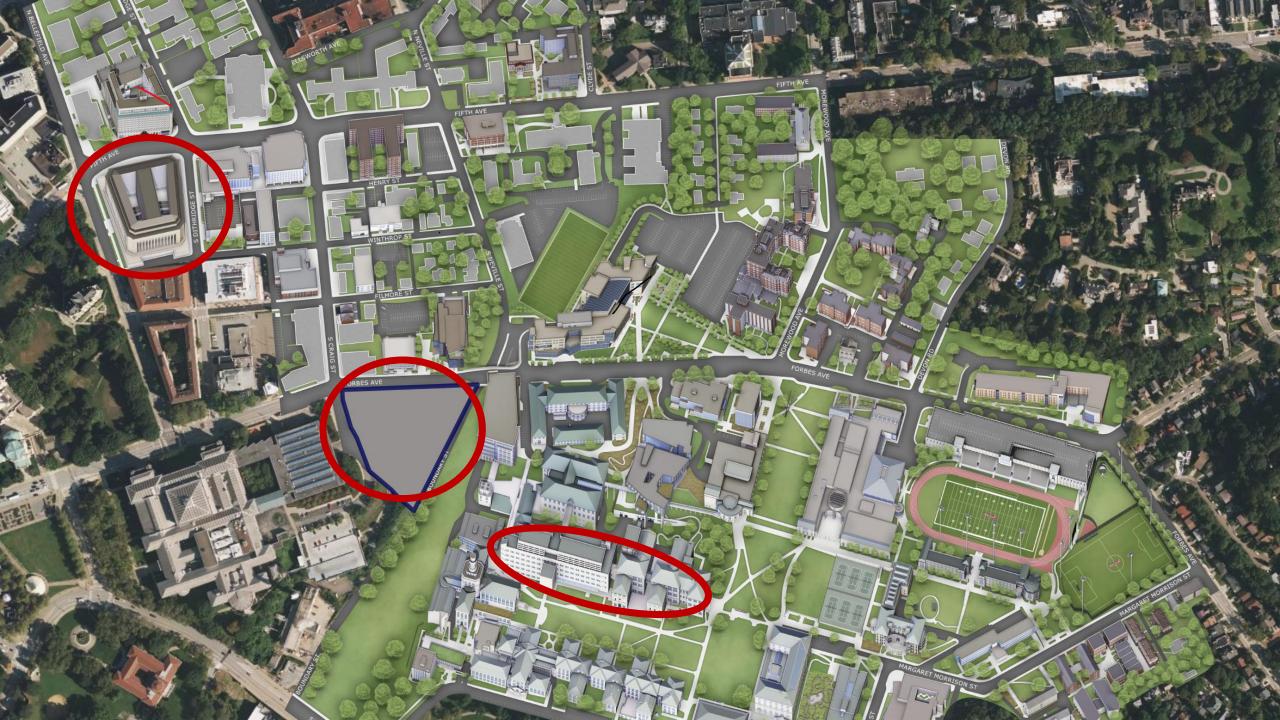
135 faculty

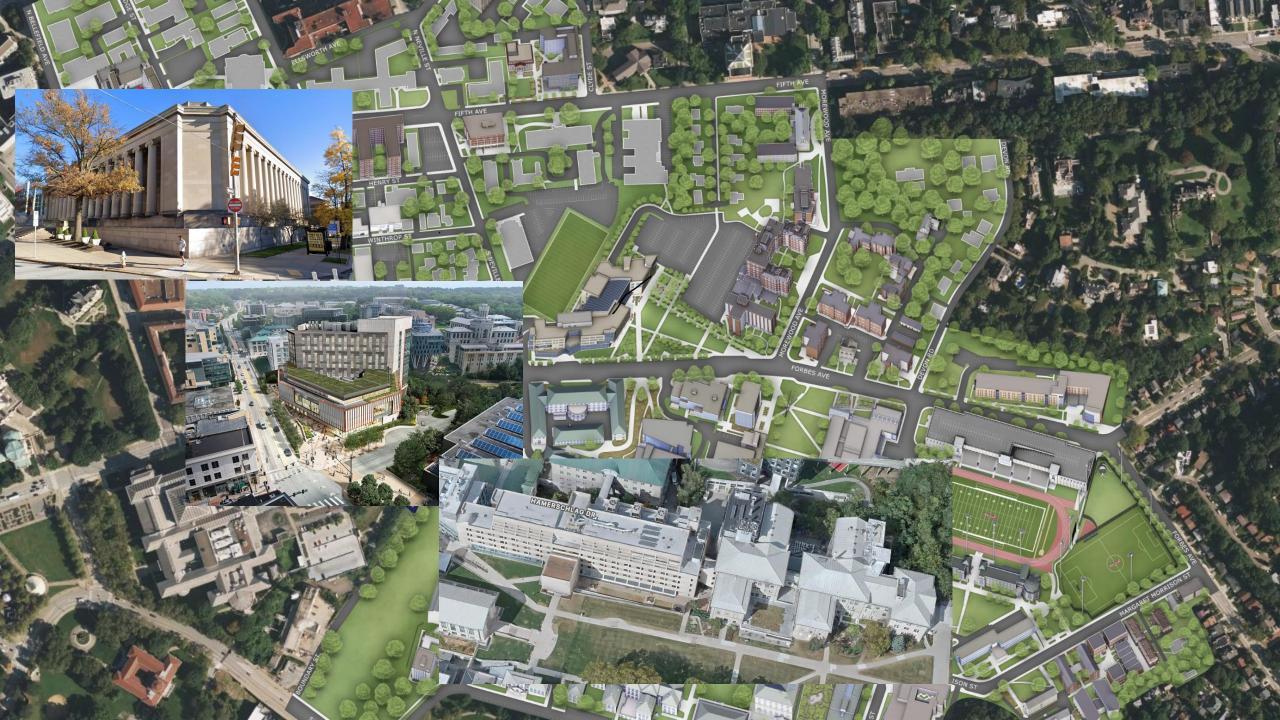
34 staff

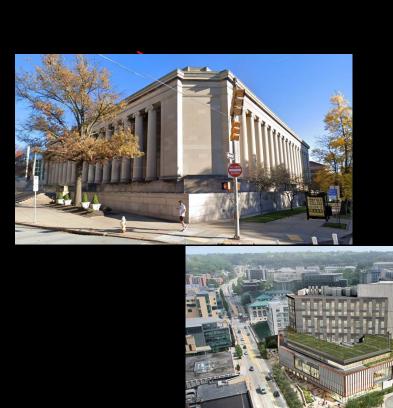
7 IT staff













Mellon College of Science – computing support scope

- college-specific support for research labs, admin, faculty, students
- Mediator, partner and 2nd level support for general computing queries, software-licensing, netOps, facilities management, infosec, physicalsec, mediatech

7 people?

- Backed by self-serve websites, many CMU central services and resources
- This allows us to concentrate on quality subject-matterfocused hardware and application support
- Flexible-work and 24h coverage (ticketing systems)

 WFH pervasive post-covid





- Stakehoders: Undergrad, grad, postgrad, faculty, collaborator
- One size does not fit all
 - Agility
 - Diverse requirements and budgets
 - Value-for-money: Repurposing -- hand-downs thru e-waste
- Their job is "throwing spaghetti at the wall", support's job is to help them to do it efficiently & effectively
 - Trial and error
 - Blameless culture
 - Post-mortem



- Science lingua franca
 - Build trust
- Emphasis on documentation and reproducibility and research integrity&quality
 - Computer architectures
 - ECC or scale
 - git
 - Conda environments



- Lab/group epoch management
 - Must survive transient members
- Success defined by
 - Publications
 - Citations
 - Patents
 - Number of students
 - Funding
 - Established/sustained lifecycle







- We are on your side
 - responsiveness
 - in-person in every building
 - Dedicated helpdesk email queues backed by ticketing systems
- Consumables stock, charging(payments), specifying, sourcing
- Workstation provision inc. loaners (mac, windows, linux, ipads, pens, adapters, peripherals)
- Scientific software support and optimization
 - You're not a computer scientist and that's OK, we're not chemists, biologists, astrophysicists, mathematicians...



- Interfacing with lab instruments
- Lab inductions and check-outs (lifecycle)
- external collaborators
 - Logins
 - Data interfacing
- Lab establishment (requirements, specification, phasing, ordering)
 - Liase with pre-award grant proposals, CMU-FMS Power/HVAC, CMU-NetOps
- Lifecycle: Install/maintain/operate/repurpose/e-waste





- Ticketing queues (helpdesks for departments and for admin, e.g expense reimbursements)
- Monitoring systems
- MCS-Linux distro Ubuntu LTS-based preinstalled with commercial scientific software licensed by CMU (Workstation, Server, Laptop variants)
 - Common platform for ease of maintenance
 - Quickly deployable (thousands of machines)
 - Allow sudo access to users if requested
 - Dual –boot, VM and container support



- 7 server rooms, 5 sites --- furthest 20miles
- 16 clusters --- largest with 988 users
- 270 computing nodes
- 43 file servers across 5 physical locations (inc backup)
- 250+ linux workstations
- 2000 HDDs
- 3 dead hard drives per week



- Student cluster + research cluster boostable to cloud
- Firedrake
- Matlab
- Mathematica
- General purpose nodes
- Single-processor optimized nodes
- Large memory nodes (4TB)
- GPU nodes



- slurm private clusters
- CERN xrootd (cosmology/physics) P2P ipv6
- REANA (Reusable and Reproducible Data Analysis) is a platform designed to facilitate the execution of computational workflows in scientific research using containerization and cloud computing
 - help to structure research data analysis in reusable manner
 - instantiate computational workflows on remote clouds
 - rerun analyses with modified input data, parameters or code
 - support for several compute clouds (Kubernetes/OpenStack)
 - support for several workflow specifications (CWL, Yadage)
 - support for several shared storage systems (Ceph, EOS)
 - support for several container technologies (Docker, singularity)



- Lab clusters (dedicated funding)
- cp2k
- gromacs
- matlab
- mathematica
- Pmemd.cuda (ambermd)
- Python / conda
- 140 nodes over 3 research groups with 2PB storage
- ECC and non ECC nodes (VRAM)



- Lab clusters (dedicated funding)
- matlab
- mathematica
- Kilosort (spike sorting electrophysiological data)
- Video processing
- non ECC nodes (VRAM)
- Storage bandwidth and data reduction

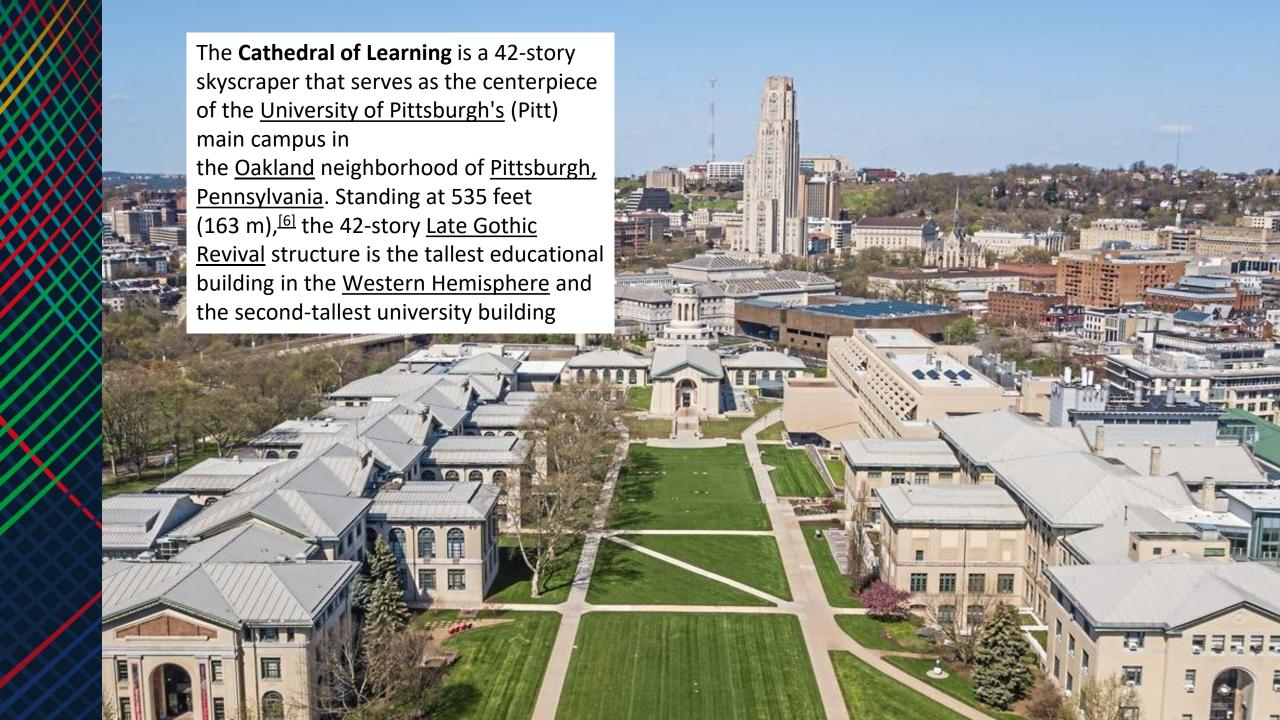


Infrastructure MCS operated – humanities and social sciences

- LLM
- slurm
- 3x 8x Nvidia L40 Infiniband RDMA 200gpbs
- 1PB fileserver



- PSC.EDU lower the error bars
 - 70 staff
 - Joint centre between MCS & University of Pittsburgh
 - Co-location
 - Offsite backups
 - 4th supercomputing center in the US





- Attitude / culture
- Self-insuring
- Self-build vs integrator
- Supplier relations
- Purchasing agility / Bureaucracy
- Pre-deploy stress-testing
- Pocicies / Cluster-policing
- Donations
- Decision Discretion
- Scale



Computing support and Infrastructure for Academic Scientific Research at the Mellon College of Science

OCTOBER 22, 2025

Nuno Chagas