Shared Research Computing Policy Advisory Committee (SRCPAC)

Spring 2023 Meeting

Alexander Urban, *Chair of SRCPAC*

Agenda

- Welcome & Introductions
- RCFC Report
  - Hod Lipson/Darcy Peterka, Co-Chairs of RCFC
- High-Performance Computing Update
  - Alexander Urban, Chair of SRCPAC
- Research Computing Services Update
  - Axinia Radeva, Manager of CUIT Research Services
- Foundations for Research Computing Update
  - Marc Spiegelman, Chair of the FoRC Advisory Committee
  - Anne Cong-Huyen/Jonathan O. Cain CUL
- Discussion of Future Directions
  - Alexander Urban, Chair of SRCPAC
- Other Business & Closing Remarks

RCFC Report

Hod Lipson, *Co-Chair of RCFC*
Darcy Peterka, *Co-Chair of RCFC*

Research Computing

Recommendations draft

2/6/2024

Committee Charge

- To recommend a strategic plan for the University's future computational and data infrastructure for research
  - Compute resources
  - Data resources for analysis, sharing, storage, archiving, privacy and security
  - Technology skills required
  - Policy impacts
  - High level cost implications
- Deliver recommendations by May 31, 2024

nature

Science and the new age of AI

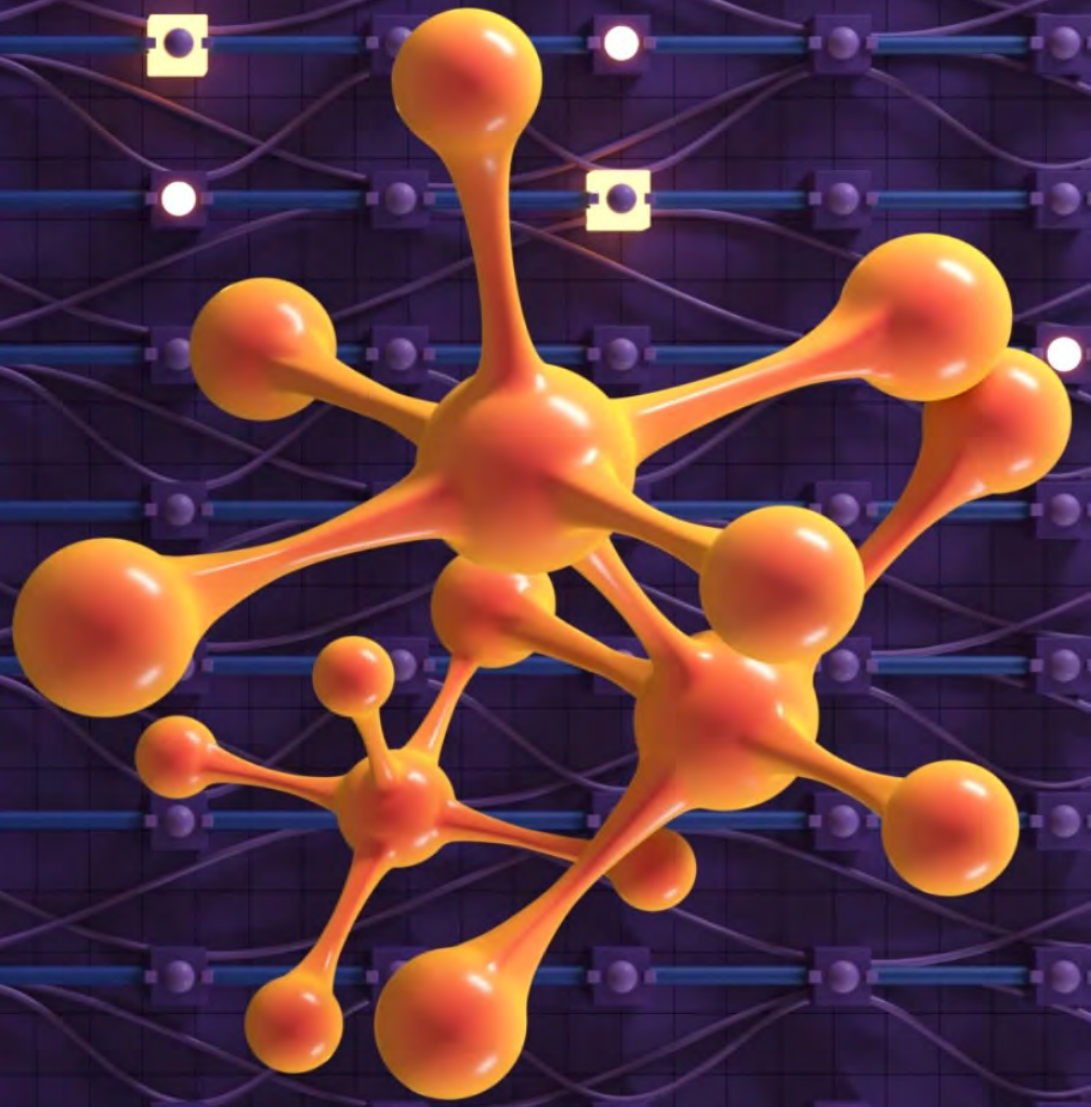

Updated 6 December 2023

Article | Open access | Published: 29 November 2023

Scaling deep learning for materials discovery

Amil Merchant, Simon Batzner, Samuel S. Schoenholz, Muratahan Aykol, Gowoon Cheon & Ekin Dogus Cubuk

Nature (2023) | Cite this article

Article | Published: 26 January 2023

Large language models generate functional protein sequences across diverse families

Ali Madani, Ben Krause, Eric R. Greene, Subu Subramanian, Benjamin P. Mohr, James M. Holton, Jose Luis Olmos Jr., Caiming Xiong, Zachary Z. Sun, Richard Socher, James S. Fraser & Nikhil Naik

Nature Biotechnology (2023) | Cite this article

b

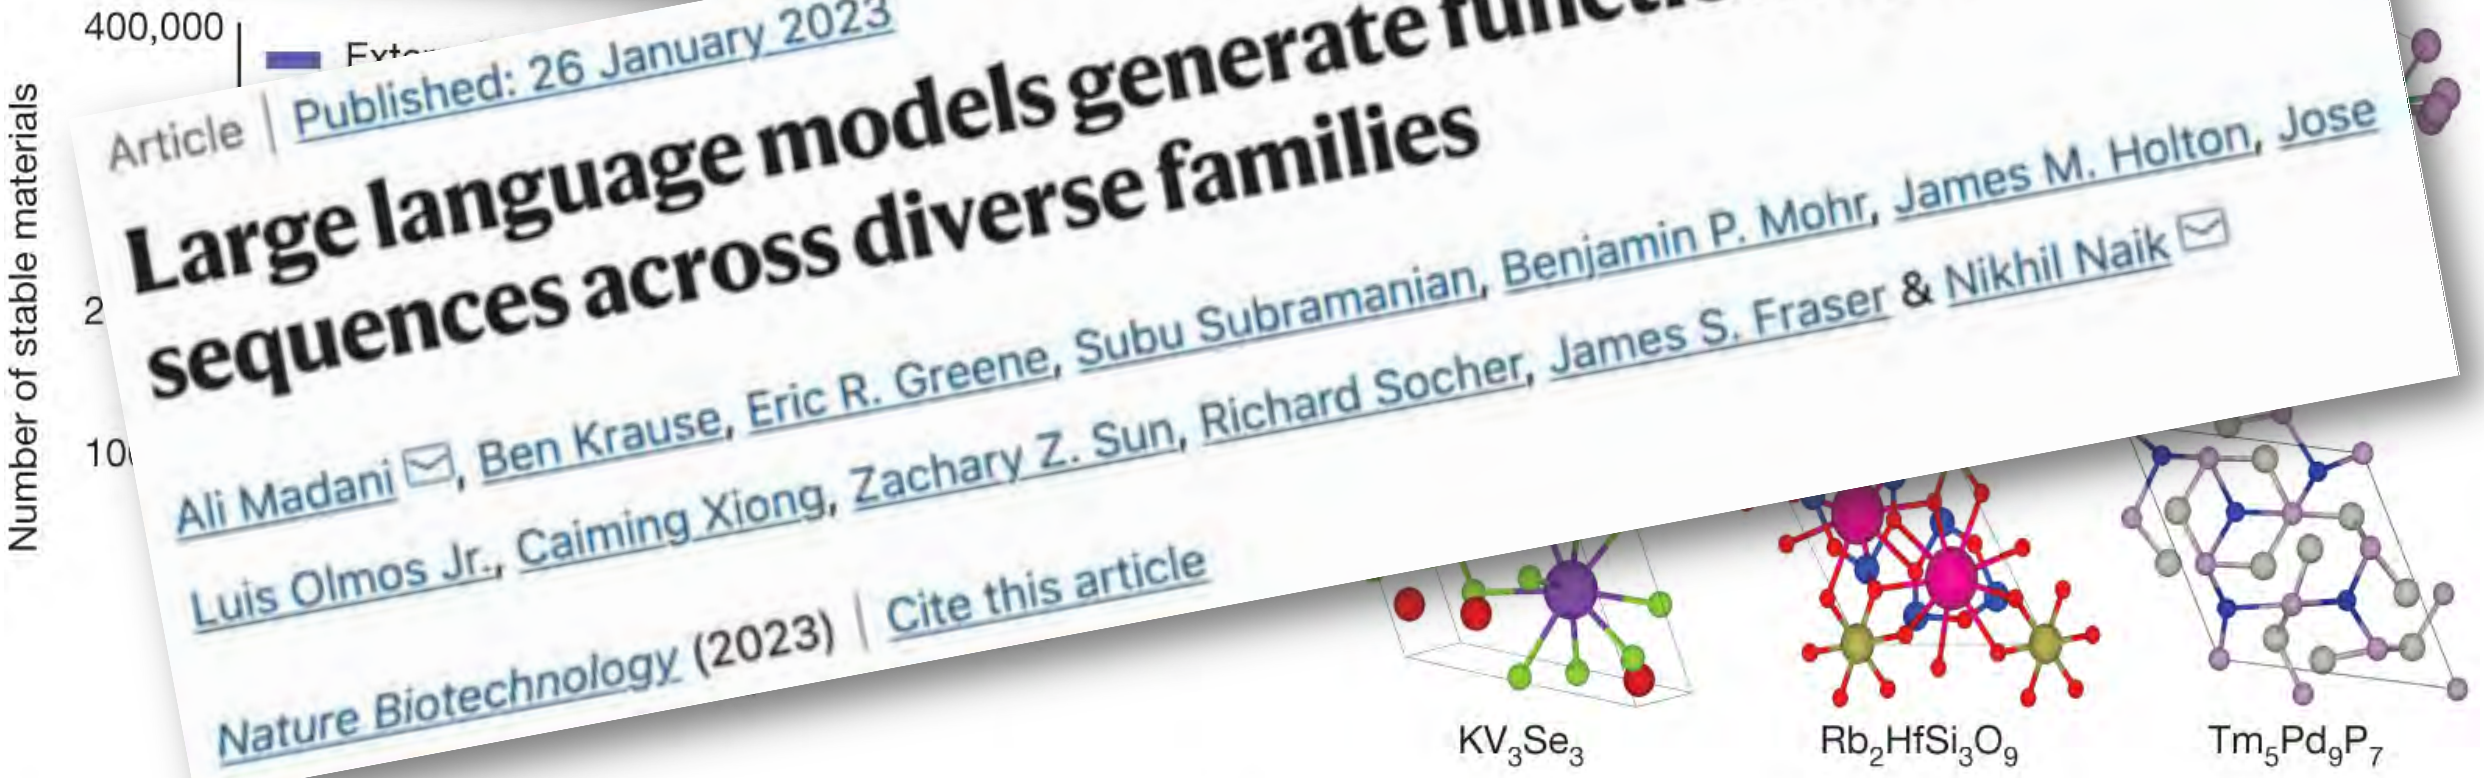

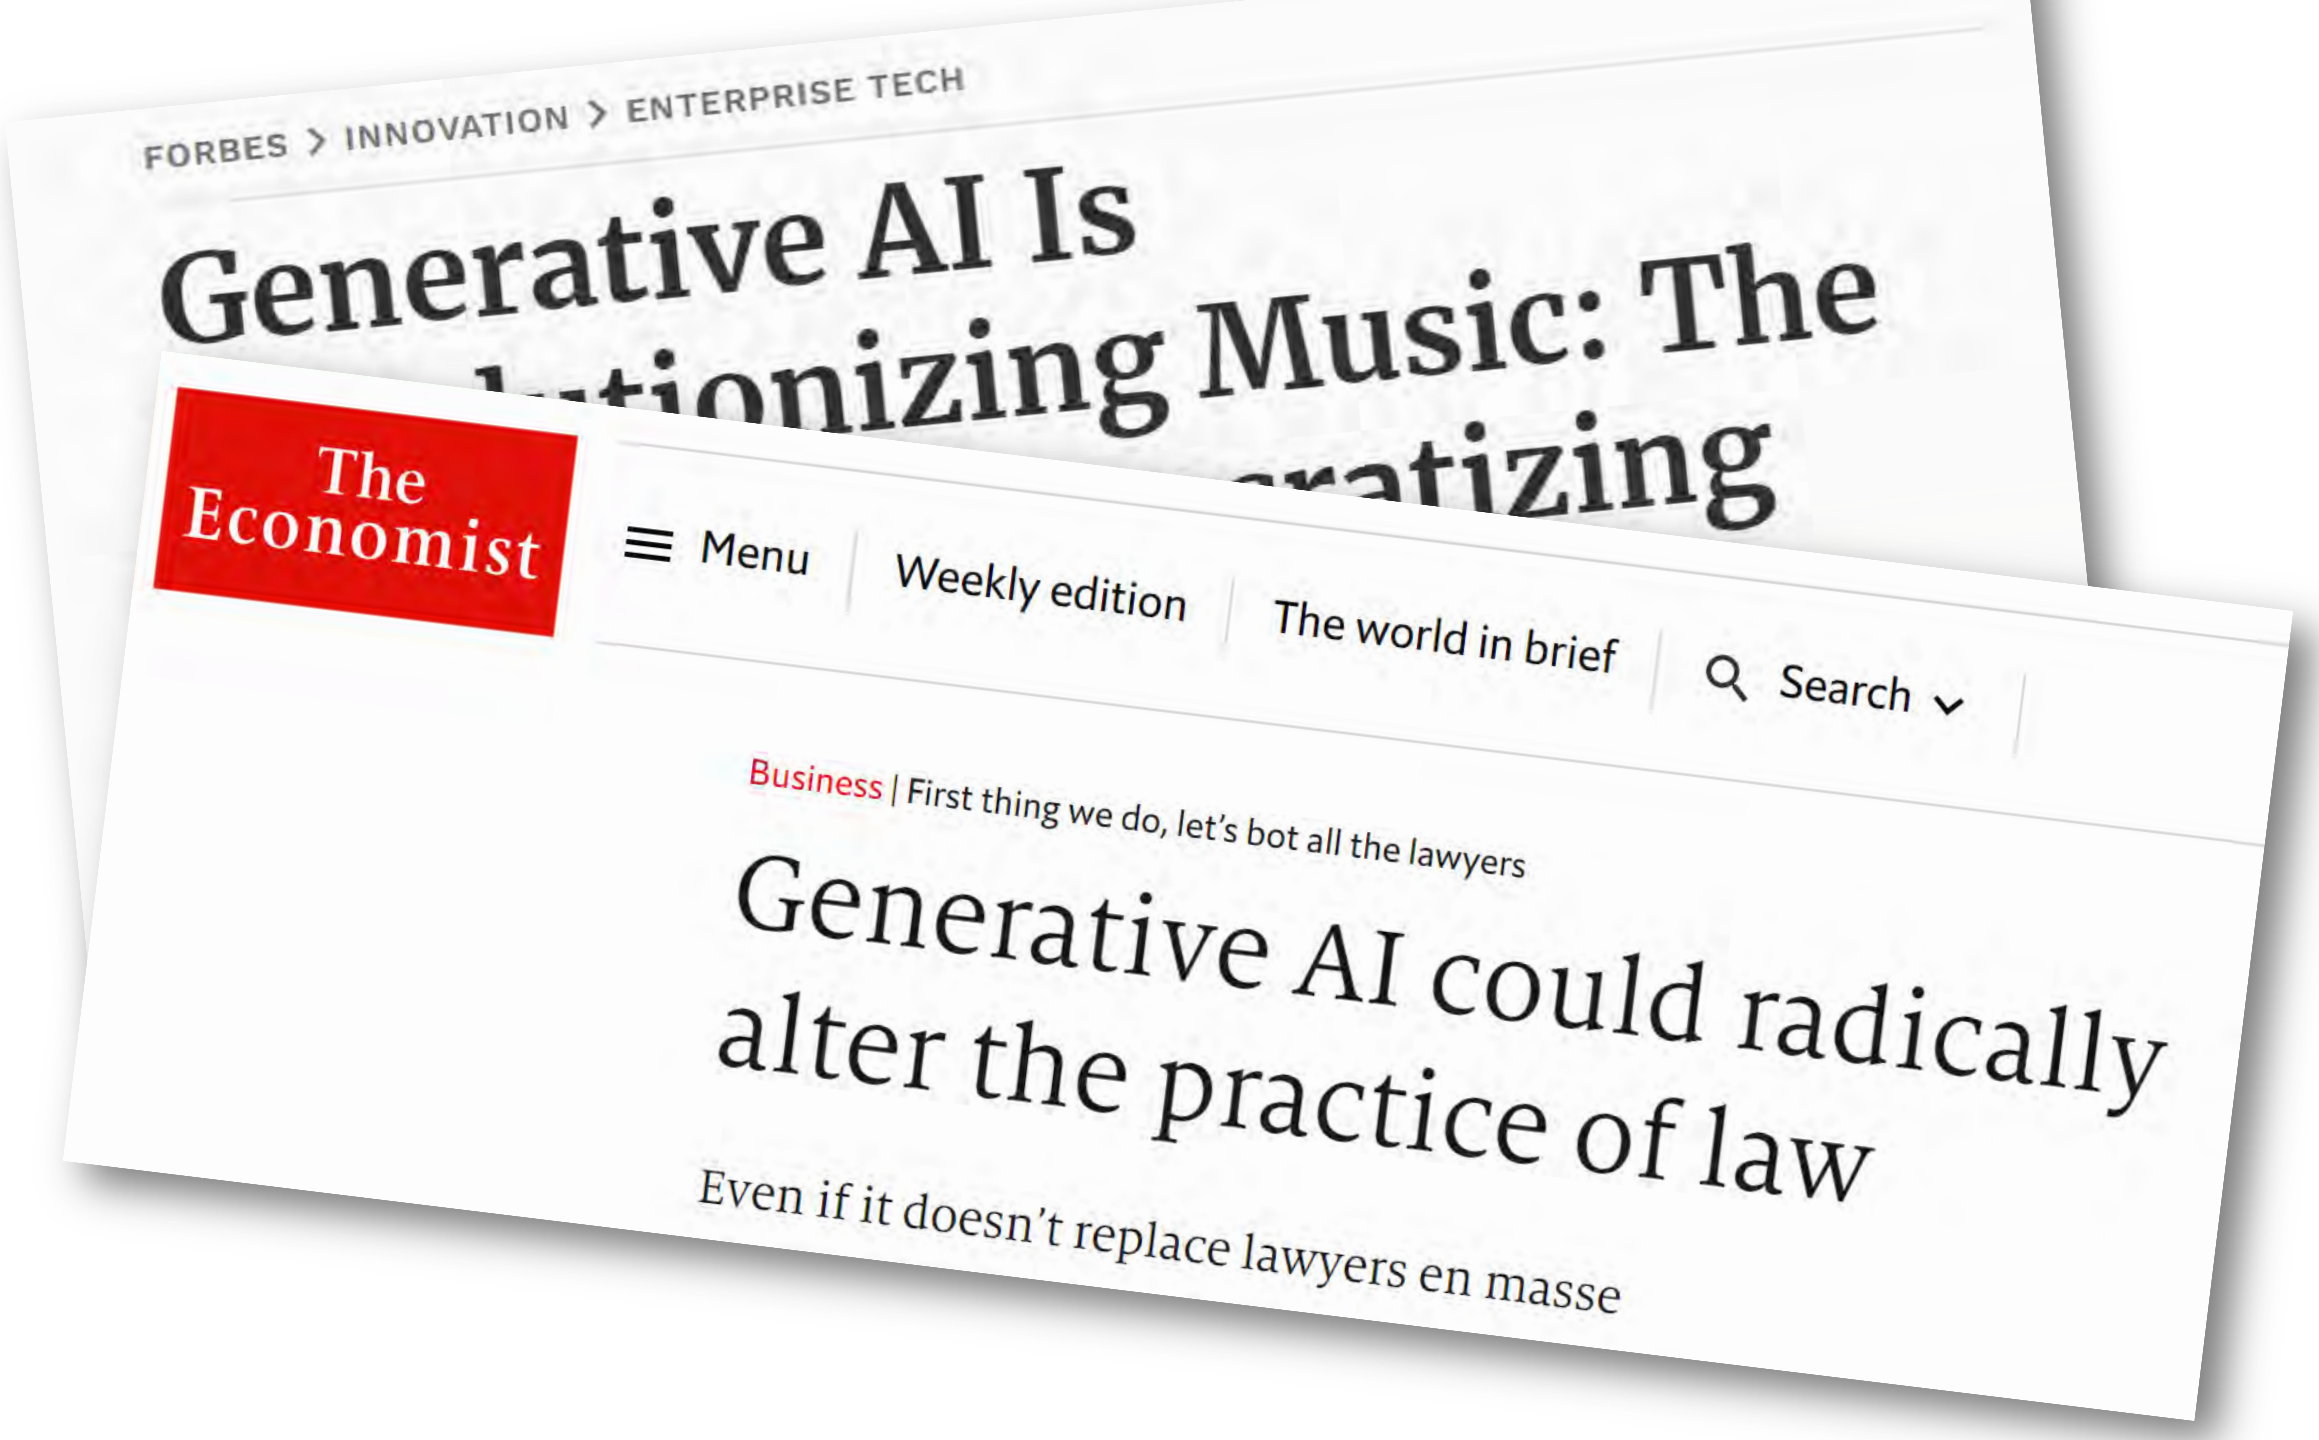

FORBES > INNOVATION > ENTERPRISE TECH

Generative AI Is
tionizing Music: The
ratizing

The Economist

Menu | Weekly edition | The world in brief | Search

Business | First thing we do, let's bot all the lawyers

Generative AI could radically alter the practice of law

Even if it doesn't replace lawyers en masse

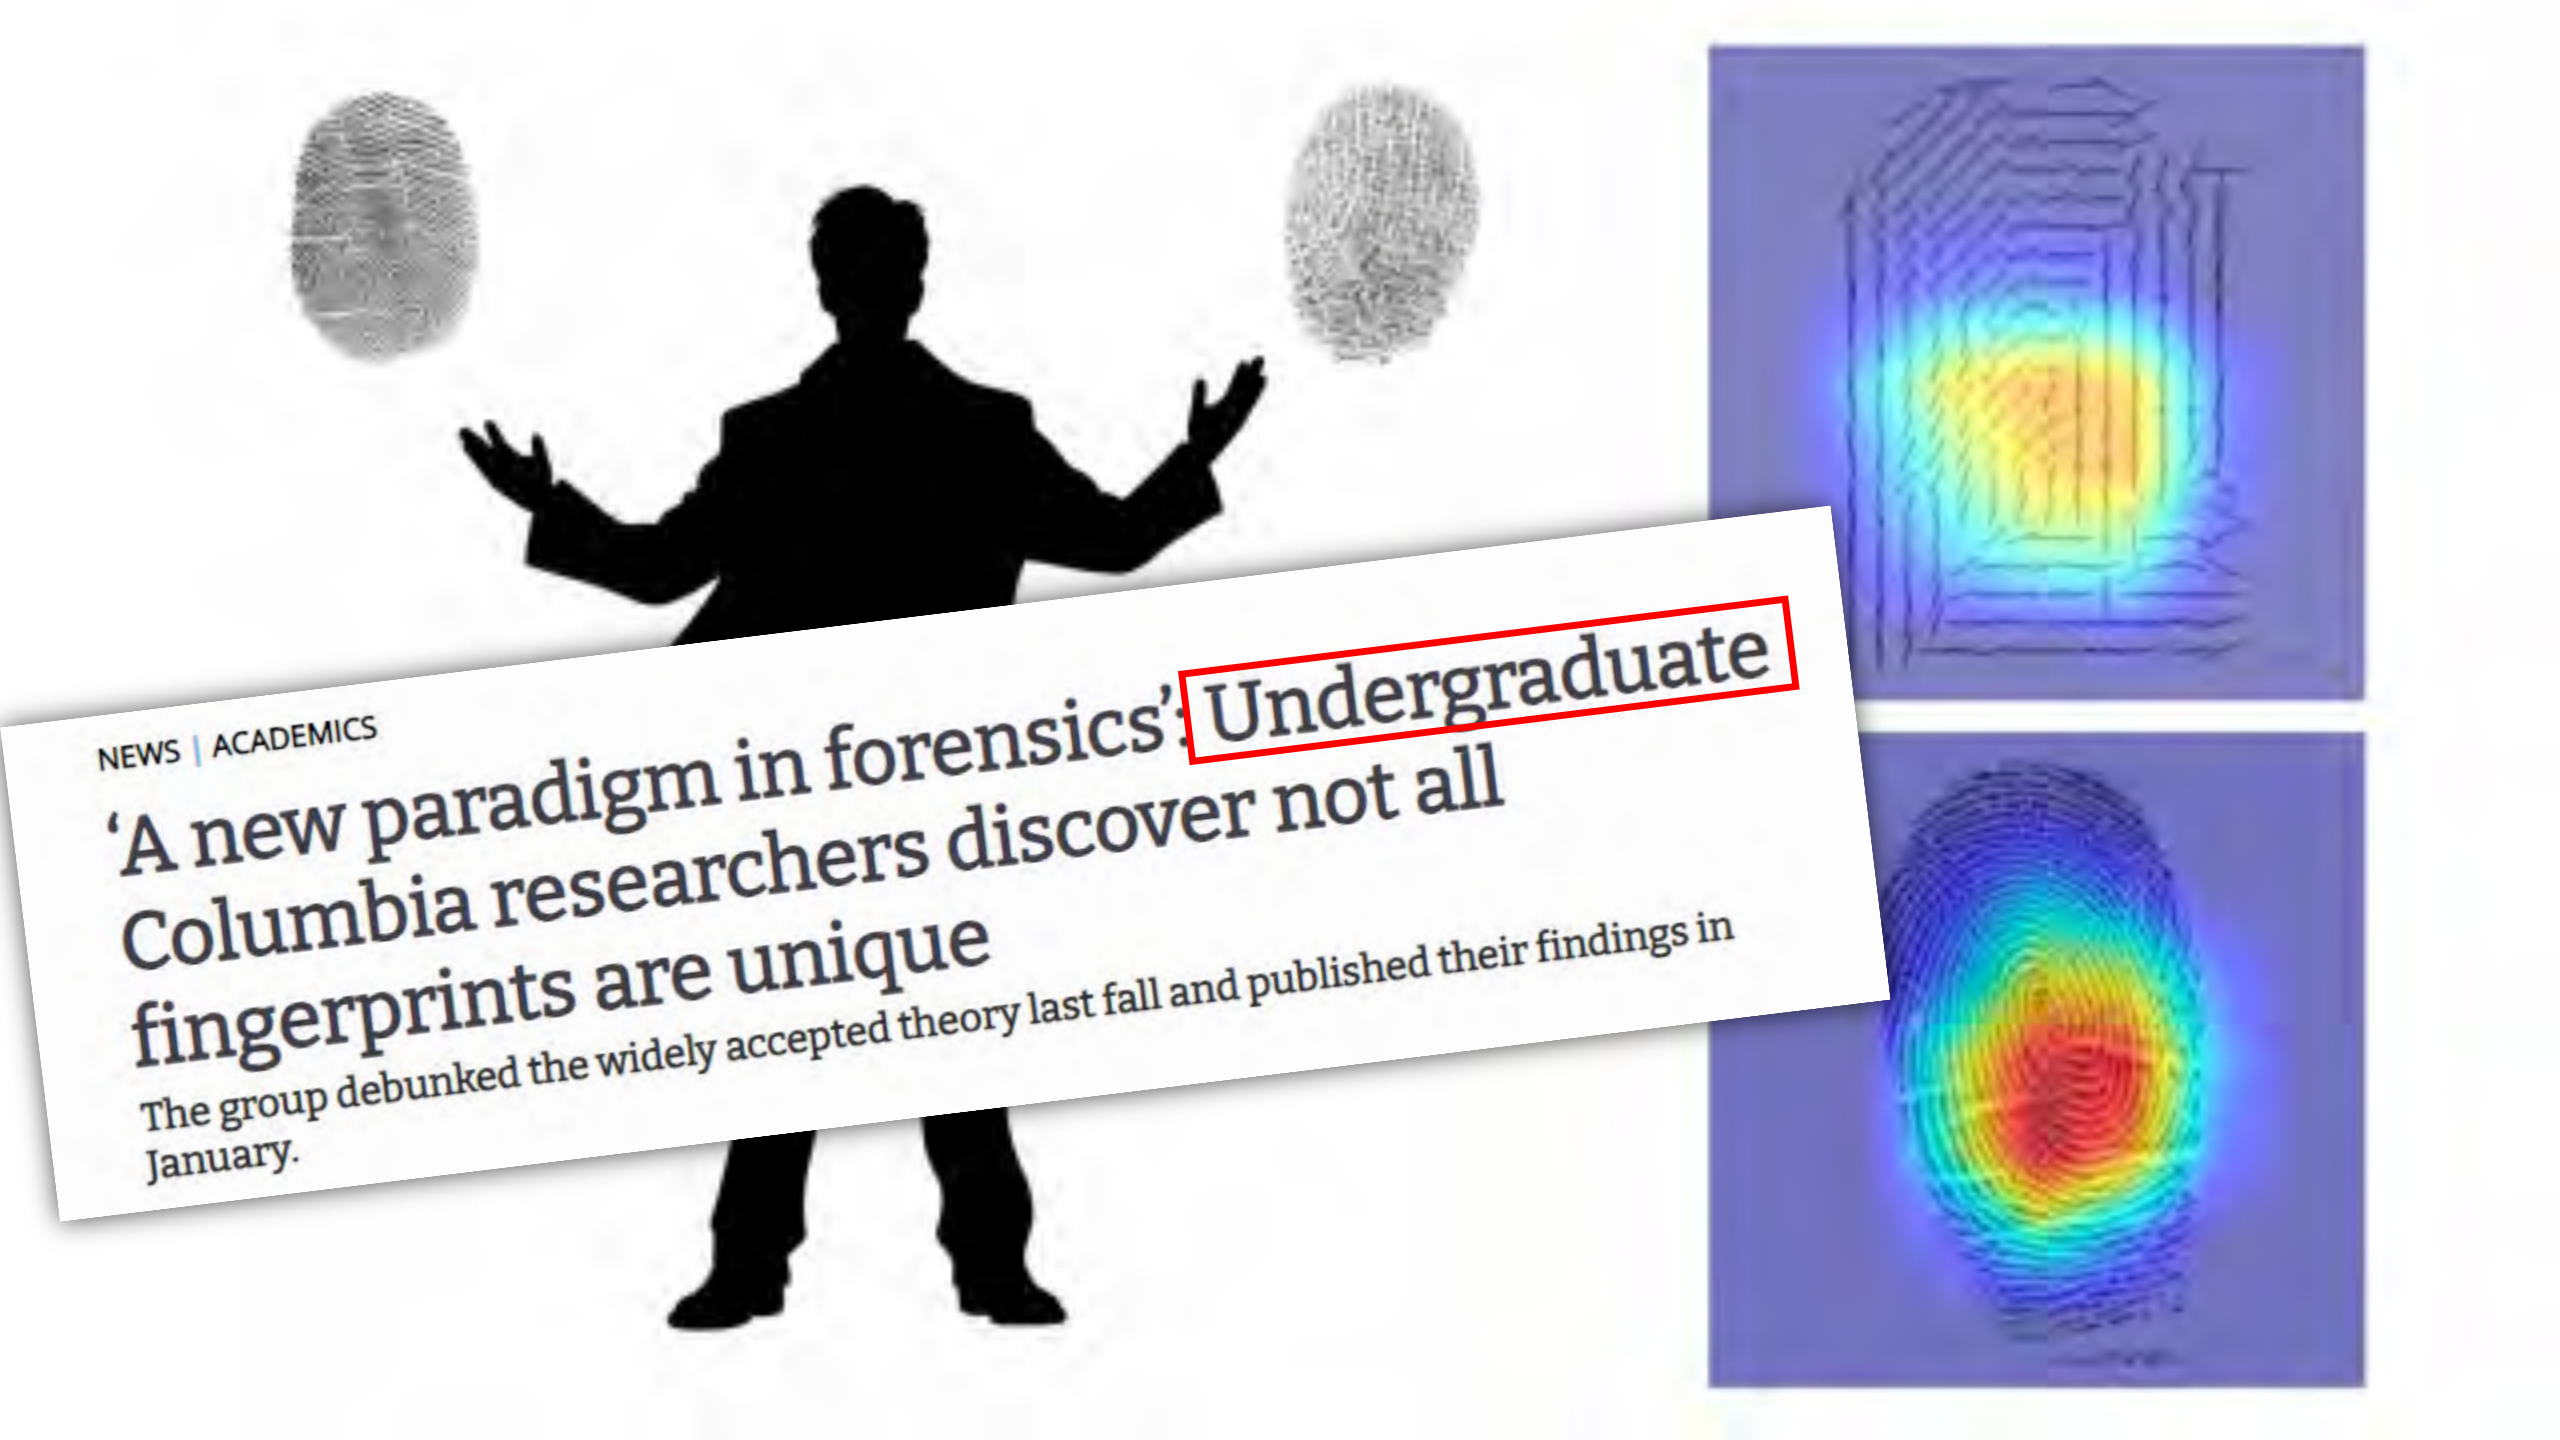

NEWS | ACADEMICS

‘A new paradigm in forensics’: Undergraduate Columbia researchers discover not all fingerprints are unique

The group debunked the widely accepted theory last fall and published their findings in January.

Transformational AI

- Every discipline field is being transformed by AI
- Transformation is unlikely to stop or even slow down
- Universities must keep up to stay relevant
- Resources are scarce even for companies

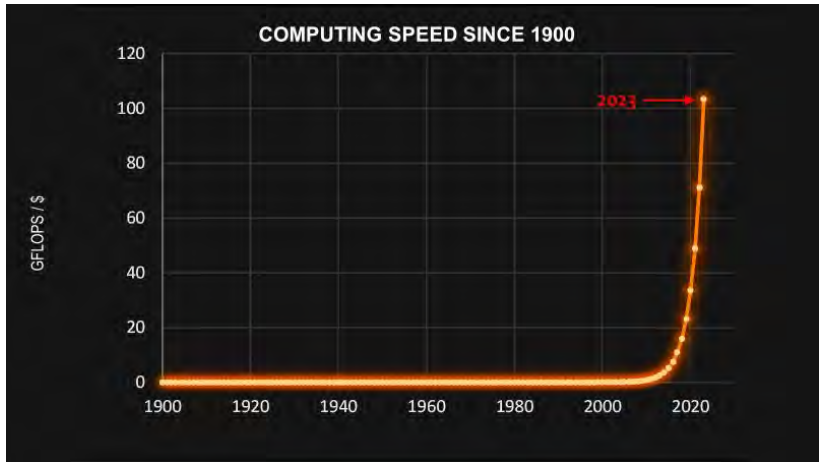

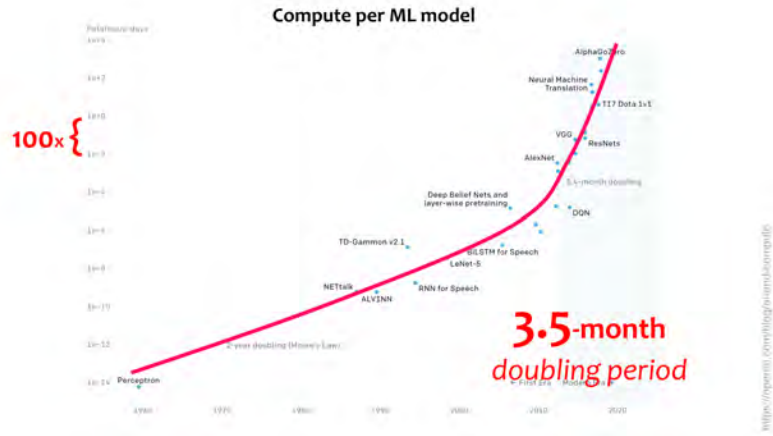

Year-to-date performance; Daily; Dec. 31, 2022, to May 31, 2023

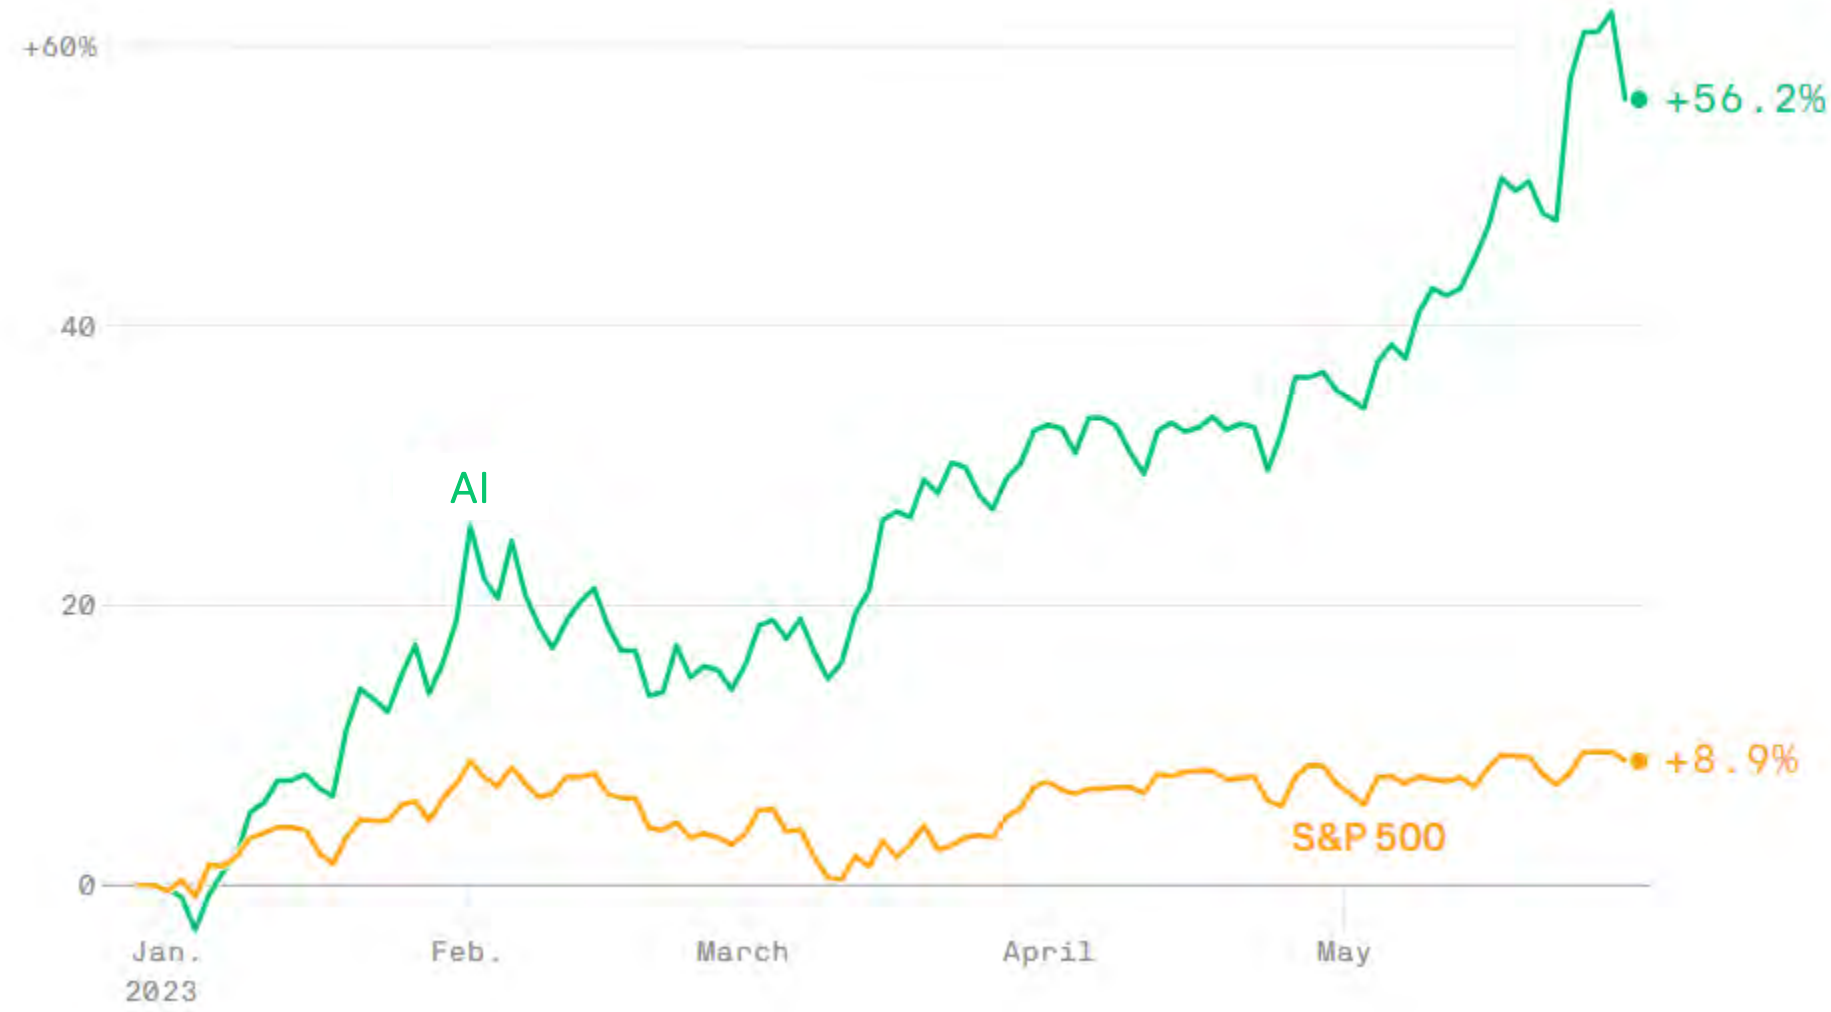

Data: FactSet; Chart: Axios Visuals

Survey

- Focus on future needs
  - *What research computing infrastructure would researchers require in the next decade to ~~keep up with~~ flourish, compete and lead your field?*

* Sample responses in appendix

From Mainframes to Datacenters to GenAI

- Evolution of shared needs
  - Mainframes / 1980s
  - Datacenters / 2000s
  - HPC/GenAI / 2020s
- Bifurcation in future needs
  - Advanced users need cost-effective HPC Resources (compute, storage)
  - Novice users need help in training and experimentation
- Many existing baseline infrastructure needs
  - Connectivity for every building, desktop for every researcher, etc
  - Not the main focus of this committee, but must be resolved

Maslow's hierarchy of needs

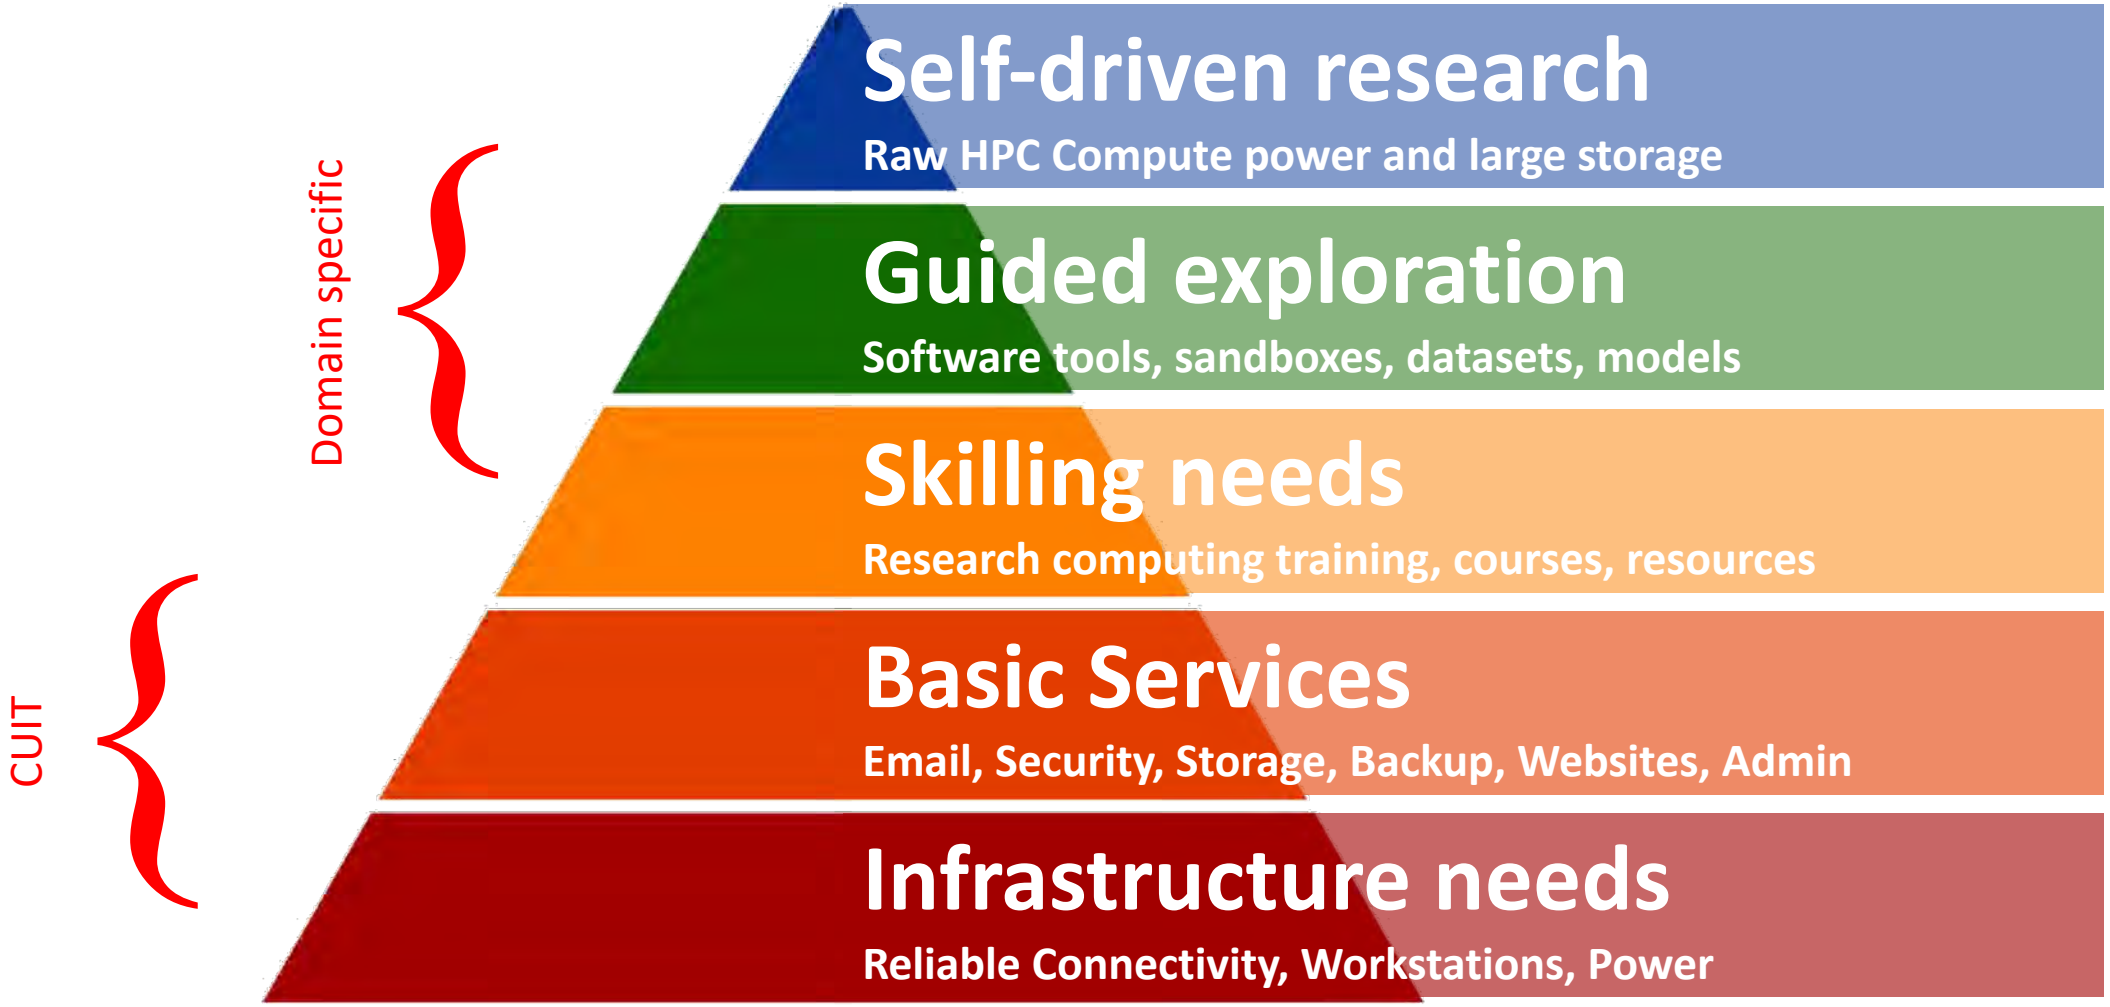

Research computing hierarchy of needs

Our main recommendation:

Discovery Accelerator

- Establish a new “lean” research computing center that provides subsidized access to high performance research <u>computing and education</u>
  - Support research infrastructure needs
  - Support research training needs

$100M Fundraising Target

UF

UNIVERSITY of FLO

SCIENCE LIFE H

UF announc artificial inte partnership

Stanford | High Performance Computing Center

Home People Events Teaching Student

Stanford University

High Performance Computing Cente

The intersection of research, te and industry.

AURORA RISING: A MASSIVE MACHINE FOR HPC AND AI

May 23, 2023 Timothy Prickett Morgan

As long as great science gets done on the final incarnation of the "Aurora" supercomputer at Argonne National Laboratory, based on Intel's CPUs and GPUs but not on its now defunct Omni-Path interconnect, people will eventually forget all of – well, most of – the grief that it took to get the massive machine to market.

1. Support HPC Infrastructure

- Low overhead
  - No new construction
  - Minimal additional staffing
  - Cost for end users < Edge computing
  - Renewable on 5-year cycle
  - Governing faculty body
- Tiered services
  - Parking spots for advanced users (AWS style)
  - Usable by all including undergrad researchers
  - Incude Empire/AI Cloud buy-in

$500K = 1 Full rack

2. Support Training / Education

- Grants for discipline-driven course and tool development
  - Intro/advanced short courses
  - Topic workshops / hackathons
  - Software/dataset purchases
  - Pretrained model acquisition
  - Sandboxes for experimentation
  - Subsidized consulting services
- Not good as centralized service
  - Evolving fast / Moving target
  - Difficult to hire and retain personnel
  - Many superb external resources
  - Doesn’t scale well / domain specific

Additional recommendations

- Basic pyramid needs must be met
  - Ensure every faculty has access to Wifi and workstation with minimal specs
- Empire-AI Cloud or subsidized commercial cloud
  - $2.5M/year, if proposal moves forward
- Use existing space if possible
  - Space requirement is minimal, e.g. Uris hall
  - Power/cooling is more of a problem
- Governing faculty body
  - Update minimum specs, new purchases. Fold in SRCPAC?
  - Administer grants / Identify new needs
- Initial + ongoing fundraising
  - As impact builds

Survey

What research computing infrastructure would researchers require in the next decade to ~~keep up with~~ flourish, compete and lead your field?

Broad Survey Results… very large range of needs, and technological savvy
Received roughly feedback from ~185 faculty, and ~50 Post-docs/ARS (in addition to feedback from members)

Many are complaints about *basic* infrastructure:
Shoddy network quality and access – poor wifi, low bandwidth, bad remote access.
Perceived lack of free or low cost access to any real compute resources
Limited mechanism to buy even basic desktop computers
Perceived lack of training – either non-existent, or not accessible, or poorly advertised
Slow IT responses
OneDrive, Zoom, free software (Adobe, etc.)

Many call out limited access to HPC (meaning CPU/GPU/Storage)
Some would like centralized specialty clusters in addition to “generic” HPC”
Want short queue times (few hour max) for big jobs
Experts that can facilitate running large jobs (run time engineers)
Very large RAM systems (1-2 TB), easier access to big data, gui beyond jupyter, etc.
Cost not competitive - meaning other places seemingly have more extensive free offerings

Desperate need access to GPU/CPU clusters for AI or ML, with concomitant storage, and gui access.
Need much bigger storage – needs to be accessible on campus, and to collabs – secure, but flexible sharing, backed up
Better data stewardship (can thing FAIR), and lifecycle management
Advanced training in specific areas – and basic training (esp. AI, cloud, some expert staff, approachable, and accessible)
Cloud integration – how and when to move, and manage cloud resources

Our report will clearly state that in “Compute is like electricity” …

1) Everyone needs reliable high bandwidth connectivity, regardless of dept. building, etc.
2) There needs to be a “minimal” level of compute that is accessible to everyone.
3) Sufficient training/education mechanisms must be available to promote efficient adoption and integration of compute
4) There will always be a need for local control. So support/resources need to span from Local to Central
5) Infrastructure investments are needed – where and how to house – balancing many constraints, but existing infra should be better utilized
6) Training is very domain specific, and mechanisms need to be in place for dept. run - central can’t do it all.

7) This is not “one and done” - there needs to be constant evaluation, steering, and stewardship of the university’s investments and direction – led by active faculty, with ex officio input from key stakeholders in CU administration.
8) Without advanced and significant local investment in compute/storage/expertise Columbia will not be a strong player in modern research. This means BIG fundraising and high-level central prioritization. Without real AI/GPU/CPU cluster in *our control*, we won’t compete – for faculty, for students, for research do

Modern Compute is really a discovery engine for every academic domain – without a good engine, Columbia will stall.

SRCPAC – Origin and Charter

Excerpt from the SRCPAC Charter, November 9, 2011:

"The Shared Research Computing Policy Advisory Committee (SRCPAC) will be a faculty-dominated group focused on a variety of policy issues related to shared research computing on the Morningside campus. As the use of computational tools spreads to more disciplines to create, collaborate, and disseminate knowledge, there is a commensurate rise in the costs of establishing and maintaining these resources. Shared resources have proven to leverage those available to individuals or small groups, but require careful consideration of the policies governing the shared resource and the basis of the operating model.

While final authority and responsibility for such policies customarily rests with the senior administrators of the University, it is vital that the research faculty examine and recommend the policies and practices they deem best suited to accomplishing the research objectives."

SRCPAC

Bi-weekly Meetings with SRCPAC Chair, CUIT, EVPR

Alex Urban, Chair of SRCPAC

Victoria Hamilton, AVP for Research Initiatives and Development

Axinia Radeva, Manager, CUIT Research Computing Services

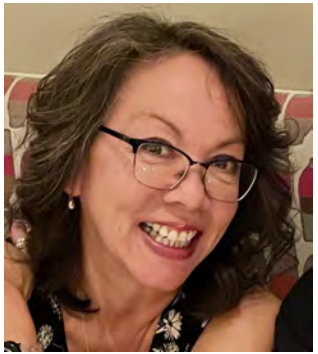

Halayn Hescock, Sr. Director CUIT Research Services

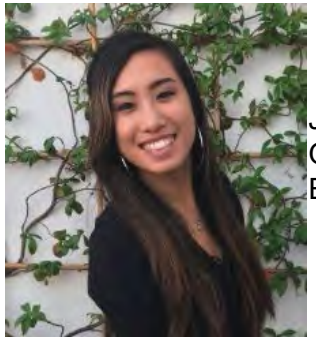

Joyee To, Operations Manager EVPR

Manager, High Performance Computing

High Performance Computing Updates

Alex Urban
Chair, Shared Research Computing Policy Advisory Committee (SRCPAC)

Current HPC Footprint

Terremoto Phase 2

- 18 Standard Nodes (192 GB)
- 4 High Memory Nodes (768 GB)
- 1 GPU 1x V100
- 3 GPU 2x V100

Ginsburg Phases 1, 2, and 3

Ginsburg has 286 nodes with a total of 9,152 cores (32 cores per node)

- 191 Standard Nodes (192 GB)
- 56 High Memory Nodes (768 GB)
- 18 GPU 2x RTX 8000 GPU modules
- 4 GPU 2x V100S GPU modules
- 9 GPU 2x A40 GPU modules
- 8 GPU 2x A100 GPU modules

Manitou - GPU Cluster

The cluster has 15 GPU nodes:

- 13 nodes with 1TB of memory 96 cores and 8
- A6000 GPUs with NVLink
- 2 nodes with 256G of memory 32 cores and 4 A6000 GPUs

Insomnia

Insomnia has 40 nodes with a total of 3,200 cores (80 cores per node)

- 24 Standard Nodes (192 GB)
- 10 High Memory Nodes (768 GB)
- 3 GPU 2 x A40
- 2 GPU 1 x H100 (backorder)
- 1 GPU 2 x H100 (backorder)

Free Tier

A portion of retired hardware, on a best-effort basis

Who is buying in?

Terremoto Phase 1
Chemical Engineering
Mechanical Engineering
Computer Science
APAM
Civil Engineering
Statistics
Astronomy

Terremoto Phase 2
Social Science Computing Consortium
Irving Institute for Cancer Dynamics
Statistics
Computer Science
Lamont-Doherty
Chemical Engineering
Zuckerman Institute
Department of Medicine
Chemistry

Ginsburg Phase 1
Ocean Climate Physics
Earth and Environmental Sciences
Mechanical Engineering
APAM
Biomedical Engineering
Chemical Engineering
Electrical Engineering
Astronomy
Biological Sciences
Chemistry
Psychology
Psychiatry
Neuroscience
Irving Institute for Cancer Dynamics
Computational Electrochemistry

# of users in the past year:
Terremoto: 1,129
Ginsburg: 1,070

Ginsburg Phase 2
Biological Sciences
Statistics
Astronomy
LDEO
Ecology, Evolution, and Environmental Biology
Biomedical Engineering
CCCE
Irving Institute for Cancer Dynamics
Physics
Astrophysics
Computer Science

Ginsburg Phase 3
Astrophysics
Earth and Environmental Engineering
Irving Institute for Cancer Dynamics
SSCC
APAM
Natural Sciences
SEAS Dean's Office
Zuckerman Institute
Chemical Engineering
Biostatistics
Environmental Health Sciences
HICCC

Insomnia
MSPH IT
Physics
Industrial Engineering and Operations Research
Irving Institute for Cancer Dynamics
Earth and Environmental Engineering
Statistics
Chemical Engineering
SIPA Center on Global Energy Policy
Biostatistics
Computer Science
Biomedical Engineering
APAM
Ecology, Evolution and Environmental Biology
Biological Sciences
Astrophysics

Manitou
Systems Biology
Computer Science

High Performance Computing Capacity

- High Performance Computing Capacity is based on FOUR major factors

Space ***Cooling*** ***Power*** ***Personnel***

- Space: We are currently occupying 13 of the 16 HD racks. Retiring hardware keeps racks rotating.
- Power:
  - Capacity: 16 HD racks fully loaded at 25kW = 400kW
  - We are currently using approximately 250kW
- Cooling: Expanding chilled water beyond the existing 16 racks will require capital investment.
- Personnel: The existing staff for administration of the current footprint is constrained.

Free Tier Resources

- The free tier is currently an unfunded POC to see the level of effort needed to support
- The free tier is intended for researchers who do not have resources on our other clusters.
- Support is limited to online documentation only.
- Compute nodes and storage are out of warranty, maintenance is on a best-effort basis.
- Users must be approved by a faculty member.

Habanero Retirement

- Habanero has been officially retired.
- Habanero was launched in early 2017, with a four-year life.
- Due to Covid, budget concerns, and remote work, we extended the life an extra year.
- Retired equipment was repurposed as a free-tier, which was widely utilized.

TERREMOTO

- Launched in December 2018
- Expanded in December 2019
- 5 year lifetime

- **Phase 1 retirement was retired in December 2023. Hardware was moved to free tier**
- **Phase 2 retirement - December 2024**

Terremoto - Cluster Usage in Core Hours

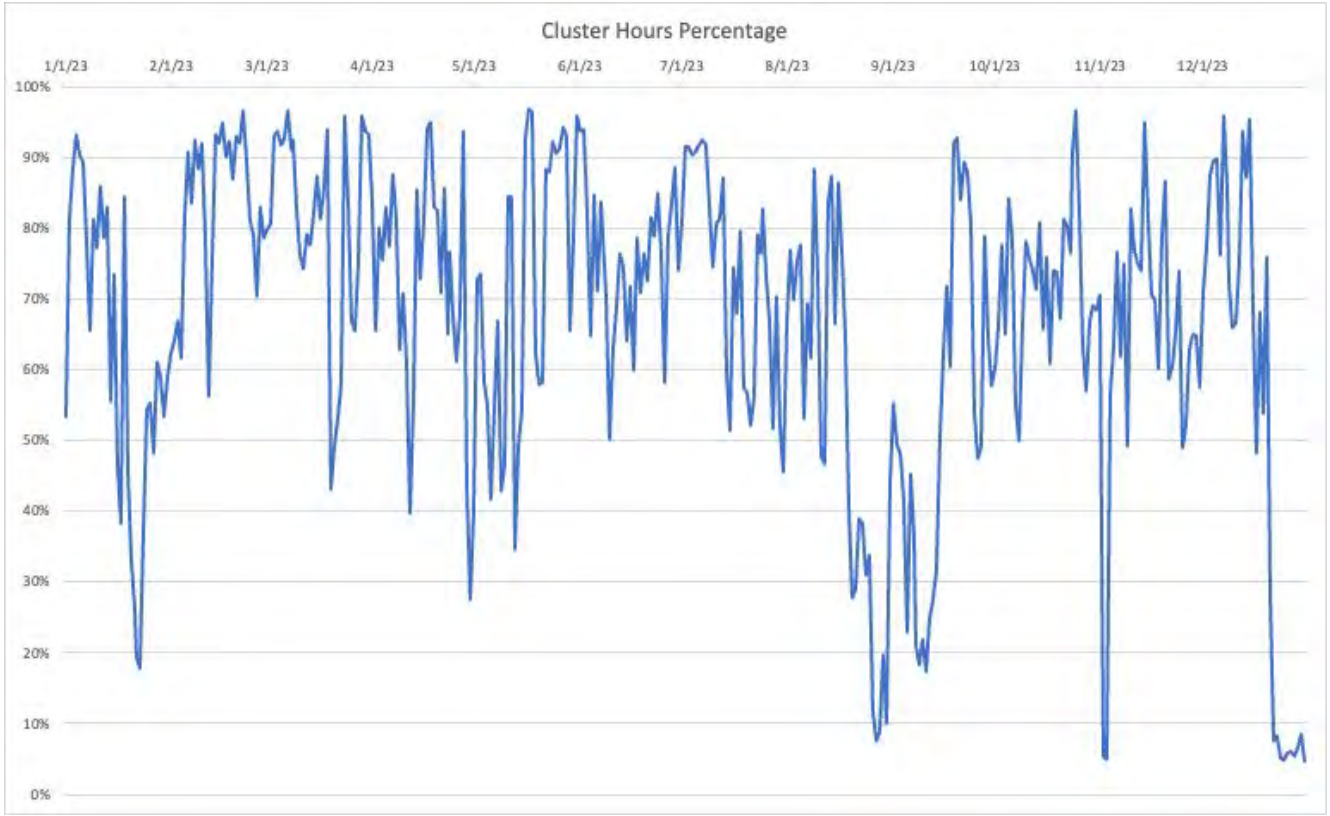

Ginsburg

- Ginsburg Expansion 2 (Phase 3) went live in December 2022
- Ginsburg now cluster total to **286 nodes, 9152 cores and 39 GPU hardware accelerated systems.**

Ginsburg Phase 1 retirement - 2025
Ginsburg Phase 2 retirement - 2026
Ginsburg Phase 3 retirement - 2027

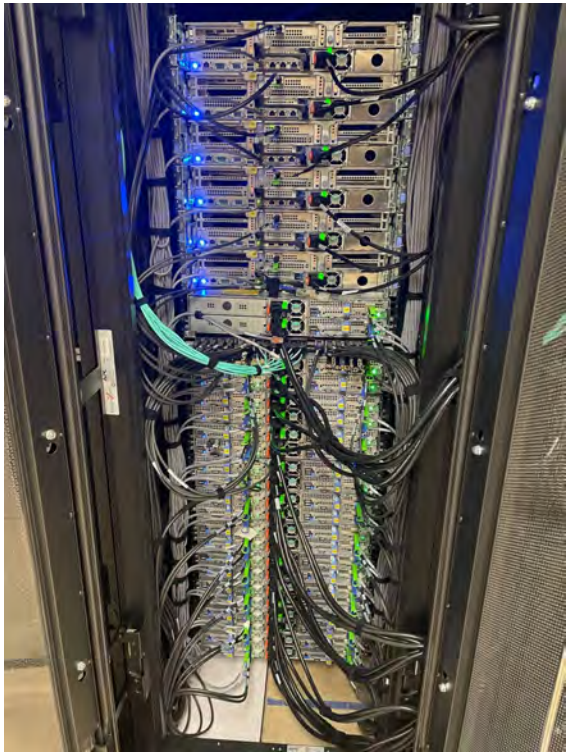

Ginsburg

Ginsburg

- Ginsburg has been experiencing issues with the performance of the DDN storage. We have had several downtimes to apply fixes.
- A configuration error was identified and corrected in January 2024. We have seen more than **10x I/O performance improvement** and we continue to monitor.

Ginsburg - Cluster Usage in Core Hours

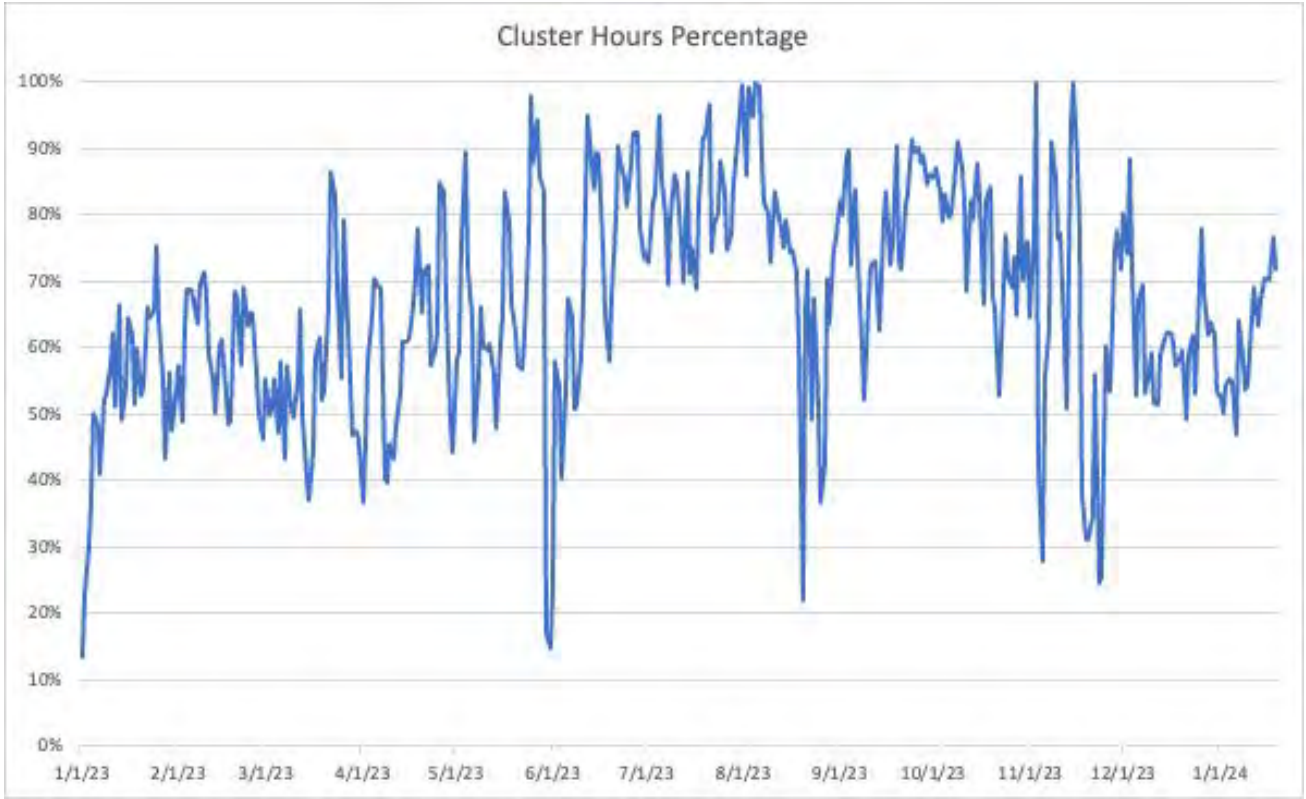

GPU Cluster - Manitou

- Manitou was delivered in late February 2023 and is currently live.
- Manitou Phase 2 expansion went live September 2023

Manitou Phase 1 retirement - May 2028
Manitou Phase 2 retirement - September 2028

- The cluster has 15 nodes:
  - 13 nodes with 1TB of memory 96 cores and 8 A6000 GPUs with NVLink
  - 2 nodes with 256G of memory 32 cores and 4 A6000 GPUs

New Cluster: Insomnia

Insomnia was racked and stacked last week and is currently being configured. It should be available for beta testing next week.

Insomnia has 40 nodes with a total of 3,200 cores (80 cores per node)

- 24 Standard Nodes (192 GB)
- 10 High Memory Nodes (768 GB)
- 3 GPU 2 x A40
- 2 GPU 1 x H100 (backorder)
- 1 GPU 2 x H100 (backorder)

The cluster that never sleeps!

Insomnia - new purchasing model

NEW PLANS FOR CLUSTER MANAGEMENT AND PURCHASING

- We plan on having one cluster going forward rather than a new cluster every couple of years.
  - One storage system
  - Better rack utilization
  - Central provisioning
- We are working with the vendor and Columbia Vendor Management on a purchase-on-demand model. We are working to get away from the annual purchase rounds.
- Offering quarter shares of standard nodes.
- Rental option still available.
- Communication to come out soon!

HPC Support Services

- **Email**
  - hpc-support@columbia.edu - all HPC questions
- **Office Hours (Online)**
  - Meet with HPC support staff via Zoom from 3pm – 5pm on 1st Monday of month: Registration required
- **Group Information Sessions**
  - HPC support staff meet with your group, upon request
- **Training Workshops every semester (Online)**
  - Introduction to Linux
  - Introduction to Scripting
  - Introduction to High Performance Computing

Research Computing Services Updates

Axinia Radeva, *Manager, Research Computing Services, CUIT*

Research Computing Services rcs.columbia.edu

Embedded Research Computing Support
We provide embedded research computing support to CPRC, SSW, Psych, and other affiliates on the Morningside and Medical Center campuses.

Secure Data Enclave (SDE)
A virtual platform used for working with **secure data sets**.

Electronic Research Notebooks with LabArchives
An online platform specialized in **organizing and storing research data**, as well as enabling information sharing and collaboration, all with automated backups and a comprehensive audit trail. Enterprise license is covered by CUIT and the Libraries.

Globus
Our enterprise Globus subscription helps you efficiently, securely, and reliably transfer data directly between systems, including between HPC clusters and Amazon S3, Google Drive, Box and more.

Cloud Research Computing Consulting
Our team can help you determine the best resources and configurations to support your needs and assist with onboarding.

Access National HPC Campus Contact
Columbia researchers can try out the **Columbia's Discover allocation** and receive guidance for applying for free Access national HPC resources.

SnapGene
A molecular biology software that allows users to plan, visualize, and document molecular biology procedures. CUIT offers the opportunity to purchase an annual SnapGene license at a **reduced price** through the University's bulk license.

Overleaf Professional
Online LaTeX and Rich Text collaborative writing and publishing tool that facilitates the writing, editing and publishing of scientific documents. Enterprise license is covered by CUIT and the Libraries.

Research Services: Enhanced Website and New Video Library

Find all your CUIT research resources in one place:
https://www.cuit.columbia.edu/research

- Research system login links
- Highlights of newly-offered tools
- Research Services portfolio overview
- FAQ
- Calendar of upcoming events and trainings

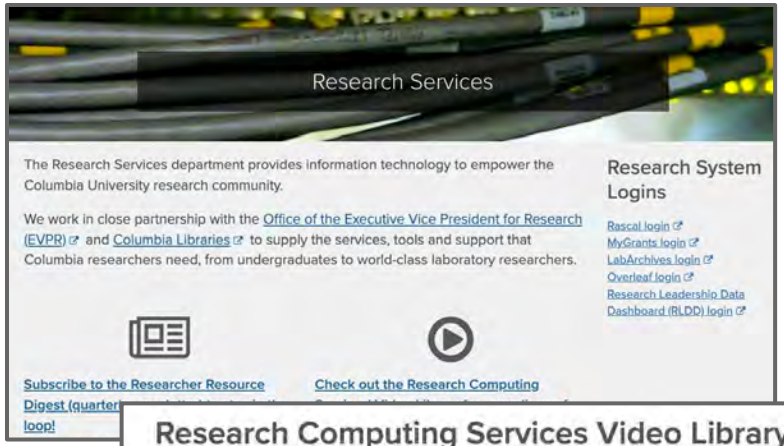

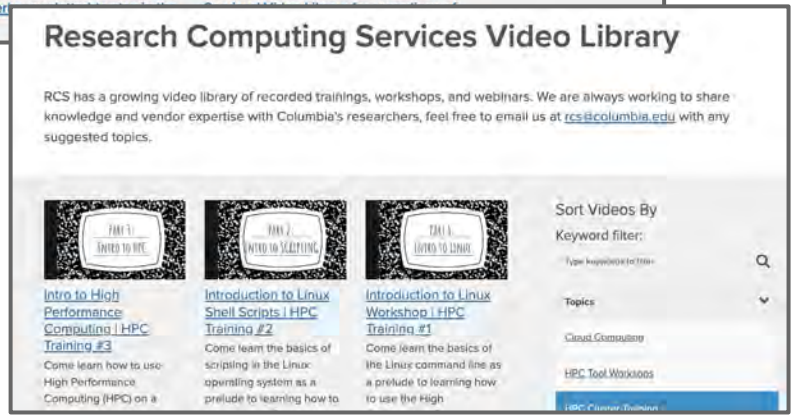

A growing repository of training recordings is now available:
https://www.cuit.columbia.edu/rcs-videos

- HPC cluster trainings
- Intel HPC tool workshops
- Cloud computing overviews
- SnapGene webinars

Overleaf

Overleaf Professional – Launched August 2023!

- Online LaTeX and rich text collaborative writing and publishing tool that facilitates writing, editing and publishing scientific documents
- CUIT and the Libraries partnered to provide an enterprise license.

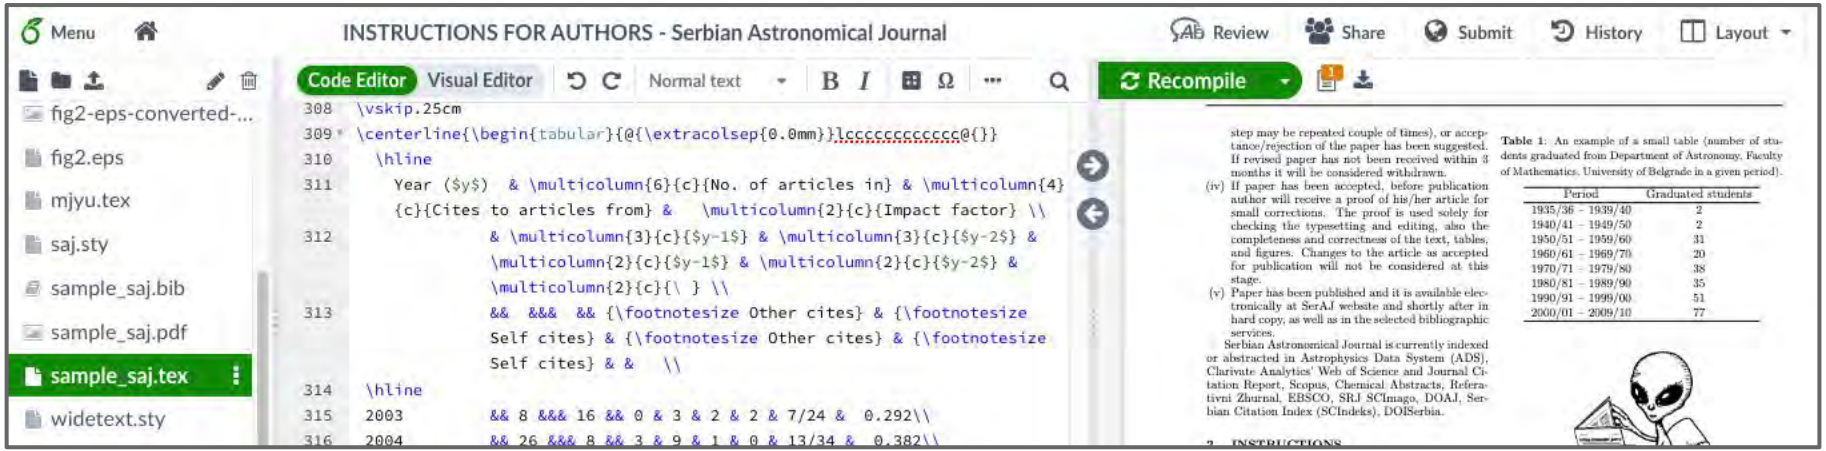

Overleaf

August 2023
- **1,657 Active Columbia users**

December 2023
- **4,558 Active Columbia** users

175% increase (2,901 users)

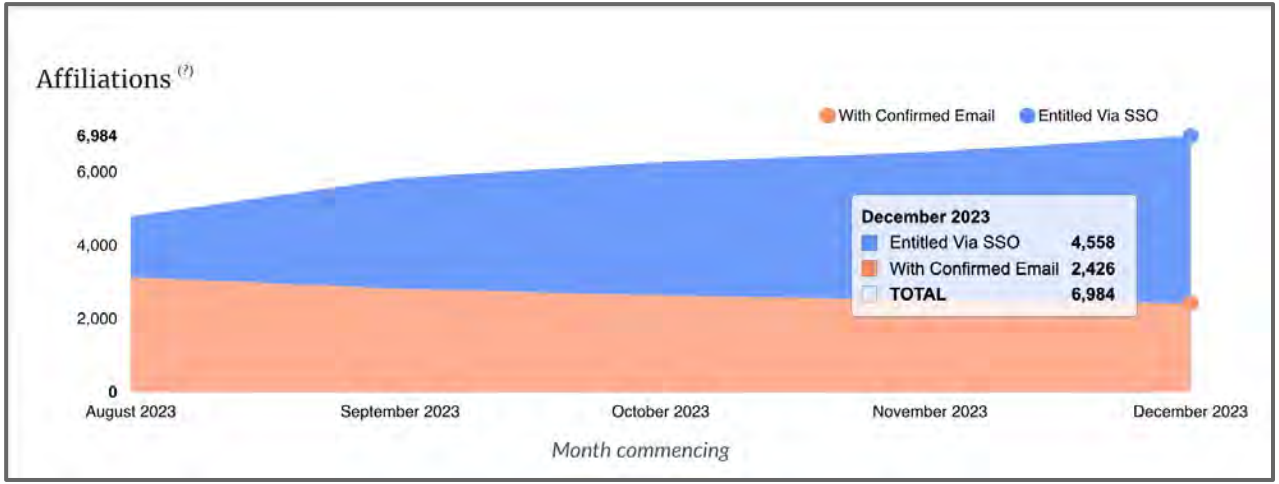

Overleaf

Number of users at other institutions that have collaborated with Columbia University in December

Secure Data Enclave (SDE)

Since 2018, SDE provides researchers with a virtual cold room to analyze and collaborate on projects with restricted data sets

- **SDE Linux environment was added in September, 2023**
- **Hardware Upgrade 2023**
  - More storage and compute have been added to address growing demand
  - New blades usable until 2030, allowing costs to be spread over 6 years
  - New blades can support GPU cards if there is demand and resources

Globus

Connectors are ready to use on SRCPAC HPC clusters

Columbia university-wide license for Globus Standard Subscription

Highly recommended for high speed file transfers to/from the HPC clusters!

- **FLEXIBLE:** Transfer datasets of any size to/from Amazon S3, Google Drive, Box, and more!
- **FAST:** Quicker than SCP, and won't affect other users by clogging the login nodes
- **FREE:** Globus is provided at no cost to you, and it's easy to get an account – simply email globus@columbia.edu with your UNI, and we'll send you an account invitation
- **RELIABLE:** Transfers automatically resume after temporary network disconnections
- **COLLABORATIVE:** Globus allows users to share data with colleagues at other institutions
- Send and receive data through Columbia HPC clusters from non-Columbia colleagues with Globus accounts

Globus Subscription Upgrade

- **Transition:** Globus Standard ➔ Globus High Assurance
  - Elevate data security with Globus' high assurance level.
  - HIPAA compliance for PHI included.
  - Streamlined, secure collaboration for subscribers.
- **CUIT and CUIMC are partnering** to provide Globus High Assurance subscription for all Columbia users (students, faculty, and researchers)

Launching February 2024!

XSEDE->ACCESS

XSEDE (Extreme Science and Engineering Discovery Environment) now known as **ACCESS (Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support)**, is an NSF-funded, nationwide collection of supercomputing systems available to researchers through merit-based allocations

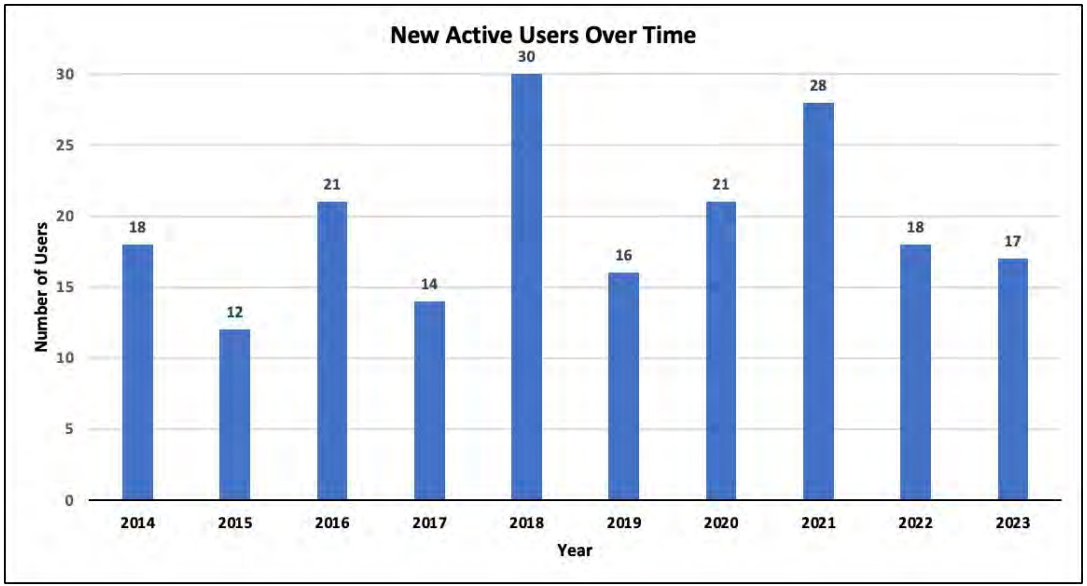

In September 2023, ACCESS approved a supplement of 750,000 ACCESS credits to our Columbia Discover Allocation, which is used for small scale testing and benchmarking for researchers

XSEDE->ACCESS

- **Additional metrics from 01/01/2023 to 01/16/2024:**
  - Total of 32 Allocations across Columbia University
  - Number of Active Users: 54
  - Number of Active PIs: 28
  - Total CPU Hours: 63,626,877.9
  - Total GPU Hours: 139,380.9903
  - Total Number of Jobs Executed: 94,837

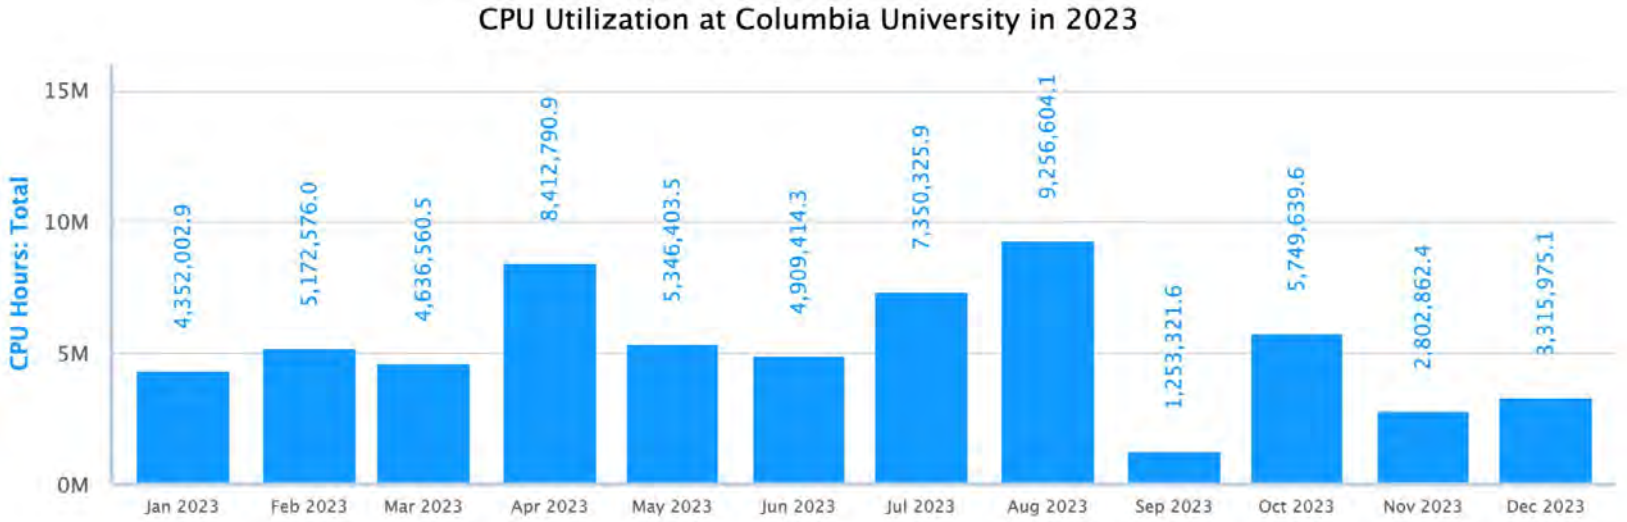

GraphPad Prism - discounted licensing

"Statistics for scientists–not statisticians"

- Guided statistical analysis pathways (no coding required)
- Exportable publication-quality graphs

Group subscription established after CUIT's embedded research engineers shared requests from their departmental users

- 52% discount from single academic license ($110 instead of $230)

Launched Feb 1 for 200+ users!

Research Computing Services

Research Computing Services (RCS) is available to discuss your research technology needs:

Email: rcs@columbia.edu

RCS: rcs.columbia.edu

Research Services: https://www.cuit.columbia.edu/research

SRCPAC Spring 2024 Update

February 6, 2024

Foundations Mission

Foundations for Research Computing provides an **informal introduction** for Columbia University graduate students and postdoctoral scholars to the fundamental skills for harnessing computation: core languages and libraries, software development tools, best practices, and computational problem-solving.

Purpose: to provide the investment in people and computational skills required to complement our investment in hardware, software and systems administration

Initial Design of Foundations

- **Novice Level**
  - Institutional Partnership with Software Carpentry
  - SC Bootcamps
- **Intermediate Level**
  - Intensives and Workshops
  - Python User Group/Python Club
  - Integration with Departmental Training (e.g. MechE)
  - Other modes (Distinguished Lecture series, CIG)
- **Advanced level**
  - Coordination with departmental curriculum

Initial Design of Foundations

- **Novice Level**
  - Institutional Partnership with Software Carpentry
  - SC Bootcamps
- **Intermediate Level**
  - Intensives and Workshops
  - Python User Group/Python Club
  - Integration with Departmental Training (e.g. MechE)
  - Other modes (Distinguished Lecture series, CIG)
- **Advanced level**
  - Coordination with departmental curriculum

New Hires

Anne Cong-Huyen, Ph.D.
Director of Digital Scholarship

Her portfolio includes Research Data Services, Academic Commons, Library Publishing Services, the Digital Humanities Center, and related services.

Anne was previously the Director of Digital Scholarship at the University of Michigan Library. She has a PhD from UC Santa Barbara.

Dan Woulfin, Ph.D., M.L.S
Computational Research Instruction Librarian

Dan oversees the Library's instructional program around computational literacy and practical skills. He works with partners in CUIT, EVPR, DSI, and others across campus.

Dan earned his Ph.D. from Stony Brook University - SUNY and his Masters of Library Science from Queens College - CUNY.

Software Carpentries Workshops for Beginners

- 15 Workshops since Aug 2018 (2-3/year), initial demand ~800 applicants, ~120 attendees per bootcamp.
- 2020 shifts to online only for Covid
- Initial coordinator P. Smyth leaves mid 2021
- Interim leadership by the libraries
- Return to in-person, January 2023 (188 applicants, 28 attendees accepted)
- Dan and Anne hired (August/September 2023)
- By January 2024 demand returned (508 applicants, 58 attendees accepted)

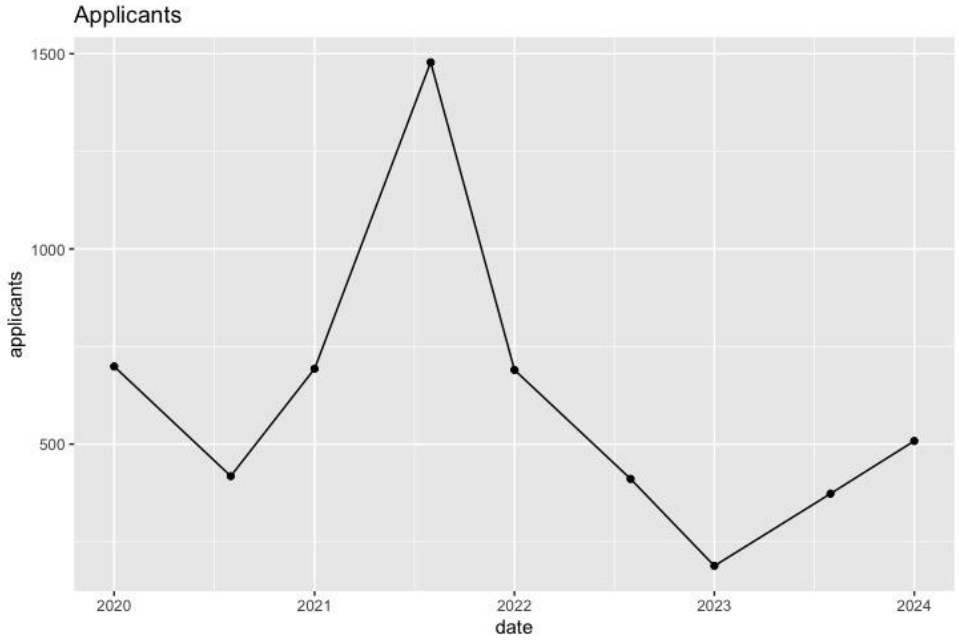

Foundations: Summary of Impact

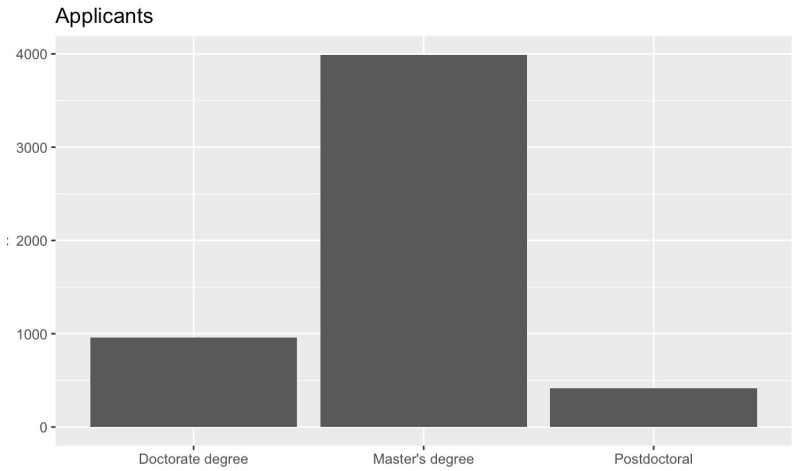

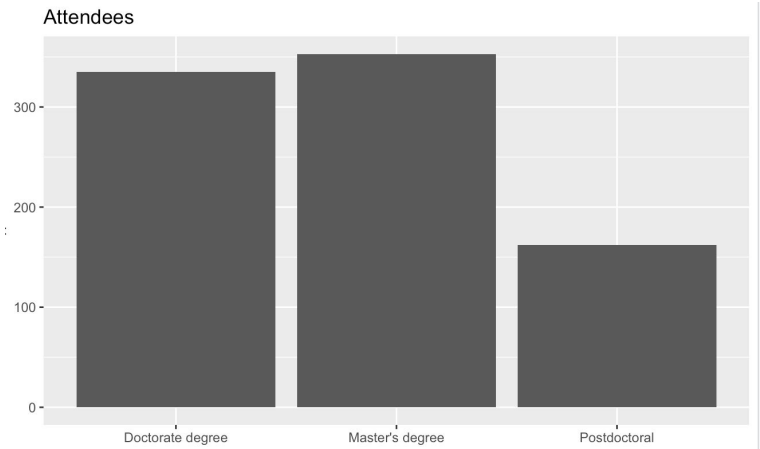

- **Who is the audience (applicants and attendees)?**
  - Since January 2020, we've had 5,458 applicants
    - Master's students - 3,991 applicants (73.4%) | Ph.D. Students (961 - 17.7%) | Postdocs (415 - 7.63%)
  - Attendees - 850 since January 2020
    - Master's students - 353 attendees (41.5%) | Ph.D. Students (335 - 39.4%) | Postdocs (162 - 19.1%)

Foundations: Summary of Impact (By School)

	GSAS	SEAS	SIPA	Business	Social Work	GSAPP	Law	Journalism	Arts	Nursing	Public Health	Physicians
Applicants	1130	1362	325	280	264	135	81	75	40	52	709	261
Attendees	256	171	19	11	39	13	17	8	5	17	94	102

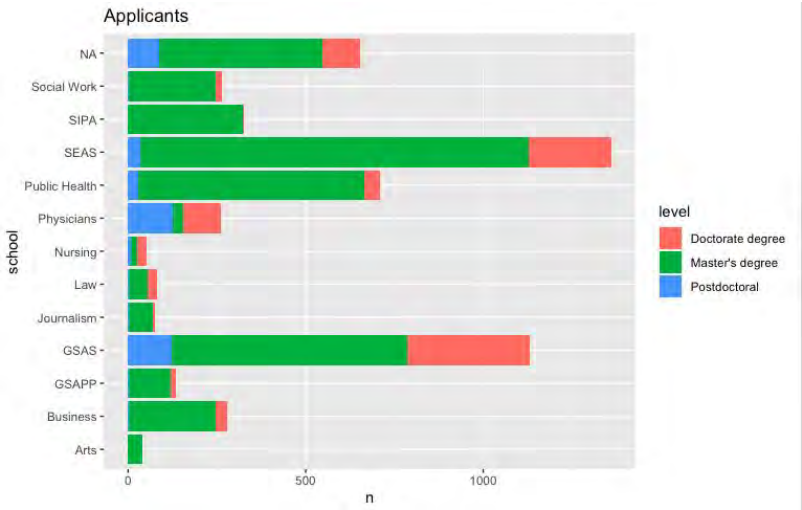

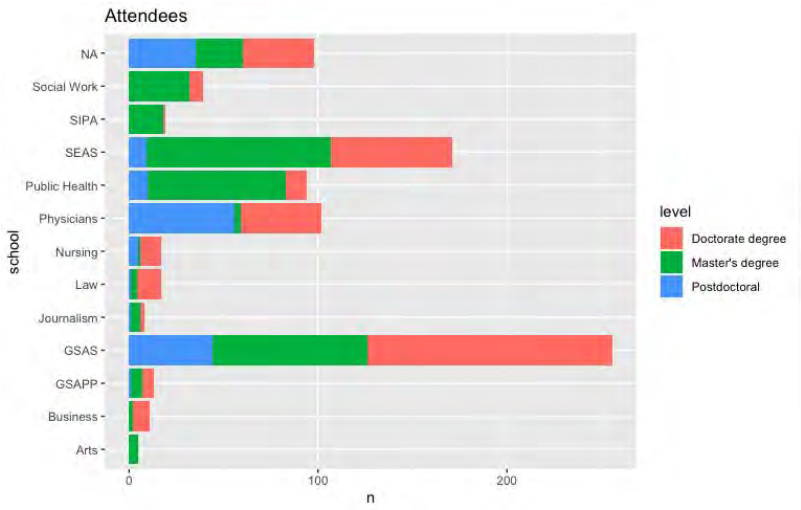

Foundations: Summary of Impact

- **How to scale to meet demand?**
  - The Carpentries model is very labor intensive and limits our capacity
    - Carpentries workshops requires a ratio of 1 instructor/helpers for every 8 learners.
    - The number of trained volunteers has decreased (55% attrition).
  - Capacity decreased post-Covid to two tracks, Python and R, with 60 total learners maximum
- Based on self-assessment of applicants, the average technical ability has risen, so most applicants have some programming experience. However, the average attendee has little to no experience.
- We're considering changes to our existing workshop model and developing additional programming to accommodate this need of intermediate programmers.

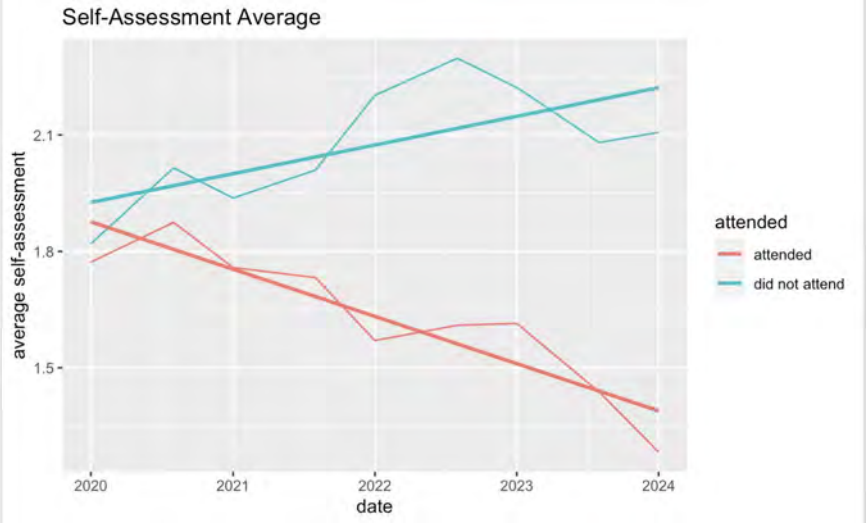

Foundations: Issues for Moving forward

Foundations was originally designed to introduce graduate students and postdocs to computational research skills. It's the first step of a lifelong journey. We remain committed to this mission.

After assessing the past program and trends, the following questions continue to come to mind:

- How do we maintain contact with faculty and their changing needs for their students?
- Are there alternatives to the labor-intensive model and curriculum we currently use?
- How can we better serve the increasing numbers of intermediate learners?
- How can we increase capacity and scale up?

Working Towards a New Framework

Addressing unmet needs by developing a more sustainable, expanded Foundations program that is data-informed, user-centered, and more modular.

Our current directions include the following steps:

- Re-engaging the Foundations faculty advisory group
- Revising the Foundations curriculum and instructor training to build local capacity
- Developing additional resources and opportunities for self-paced learning and/or more varied learning modalities
- Curating additional resources and linking learners to new and existing learning opportunities
- Leveraging staff and spaces to support learning communities, labs, cohorts, etc.
- Centering Computational Literacy as a structuring focus in the Library

Future Directions

SRCPAC – Where we are

- **Shared HPC**
  - Access policies work well for power users
  - Slow buy-in addressed with new continuous purchasing model
  - High cost addressed with quarter shares and rental option
- **Other research computing services**
  - More campus-wide licenses (e.g., Overleaf)
  - Growing number of training videos
  - Consultation options

SRCPAC – Where we could do better

- *Improve communication and marketing*
  - Apparently, many across campus are unaware of existing resources.
  - The access tiers appear not to be widely known.
- *Implement better entry-level training?*
  - Demand for entry-level bootcamps is ~5–10× the class size.
- *Address heterogeneous computing needs*
  - Long-running GUIs for data analysis (Mathematica, Matlab, Jupyter)
  - Software requirements (compilation, virtualization, campus licenses)
- *Figure out how to sustain the free tier*
  - Policy for phased out hardware, maintenance cost (personnel & repairs)
- *Increase faculty involvement*

Get involved

Nominate yourself or a colleague as SRCPAC Co-Chair

or

Offer to chair or join a temporary *ad-hoc* subcommittee on entry-level training.

Your ideas

Time for your
questions and suggestions