UC Advanced Research Computing Center (ARC) and Core Facility

Tools and Services

Notebooks

The ARC web interface offers computational notebooks Jupyter Notebooks, RStudio and others on request.

Scientific Software

The ARC has a variety of software applications, compilers and support for software license management.

Containers

Docker and Apptainer (formerly Singularity) containers for reproducibility, transferability, sharing environments.

Remote Desktops

Seamlessly connect to the Linux infrastructure with a familiar environment regardless of your device or location.

Consultation

Onboarding, data workflow optimization and custom consulting available to help you scale up your research

Training & Education

Intro to HPC, SLURM, Workflow Optimization workshops
Knowledgebase articles
Slack Support Channel

Data Transfer

Share data with collaborators anywhere and seamlessly connect remote sites to the ARC's many research storage systems.

Grant Support

Language to include in proposals
Computing costs estimates/quotes
NSF Facilities, Equipment, and Other Resources doc

Computation/Analysis

GPUs, CPUs, High Memory CPUs
AI/ML, Modeling Simulation, Analysis/Visualization, genome sequence, statistics

Cloud Services

Research cloud services and virtual machines are provided and supported by the ARC team.

Storage

Lustre, Bulk and Tape Archive
Securely and easily store any research data by connecting to the system directly from your workstation or lab environments.

Secure Research

The ARC provides a secure research environment which is professionally maintained in the enterprise data center.

UC Advanced Research Computing Center (ARC) and Core Facility

Examples of Bioscience Research Acceleration at Scale

Machine Learning

Dept of Radiology postdoc was able to meet a paper submission deadline by using the ARC HPC cluster GPU nodes. The machine learning modeling and training was performed 100x faster on ARC GPU node than the local workstation with GPU cards (2 hrs vs. 200 hrs)

Parameter Sweeps

Hundreds of parameter sweeps were processed simultaneously instead of running one sweep at a time, freeing up the researcher's computer for other tasks and accelerating the time to results.

Secure Data Collection

CryoEM researcher is securely transferring data from their external collaborators instead of shipping hard drives back and forth, reducing their data acquisition time from weeks to minutes.

Statistical Analysis/ML

Pharmacy PhD student used Jupyter notebooks on the ARC HPC cluster to perform statistical analysis and machine learning to forecast how weather conditions impact the demand for migraine. The ARC computational scientists worked closely with the student to develop his workflow and optimize his code.

Deep Learning

Expanding articulatory information from ultrasound imaging of speech using MRI-based image simulations and audio measurements. Characterization and prediction of missing articulation information from ultrasound imaging via numerical acoustic wave simulations based on tissue maps segmented from MRI and deep learning using tongue contours and audio features.

Data Analysis

Metagenomic and metabolomic data analysis of "big data" sets by performing taxonomic and functional annotation

Deep Learning

A supervised deep learning approach is investigated to improve prediction of ablation zones and tissue temperature from 3D echo decorrelation images.

RNA Sequencing

To perform single cell RNA sequencing for our raw data and identify new cell populations in our area. We need computationally powerful CPU and large RAM for performing this analysis

Research Storage

Biomedical engineering researcher moved 44TB from their previous institution to ARC storage where it is securely stored, backed up and managed and easily accessible.

UC Advanced Research Computing Center (ARC) and Core Facility

HPC/AI Cluster Specifications

95 - Regular Memory Nodes (6,080 cores)

Regular Memory (RM) CPU nodes provide extremely powerful general-purpose computing, machine learning and data analytics, AI inferencing, and pre- and post-processing.

- Two AMD EPYC CPUS, each with:
- 32 cores
- 2.35-3.35GHz
- 128MB L3
- 256GB of RAM
- 8 memory channels
- SATA SSD (960GB)
- Mellanox ConnectX-6 HDR InfiniBand 100Gb/s Adapter

1 - Large Memory Node

Large Memory (LM) node provides 1TB of shared memory for, and other applications requiring a large amount of memory for which distributed-memory implementations are not available.

- Two AMD EPYC CPUS, each with:
- 32 cores
- 2.35-3.35GHz
- 128MB L3
- 8 memory channels
- 1024GB of RAM
- Mellanox ConnectX-6 HDR InfiniBand 100Gb/s Adapter

10 – GPU Nodes (20 GPU cards)

Two NVIDIA Tesla A100 40GB GPUs

- Two AMD EPYC CPUS, each with:
- 32 cores
- 2.35-3.35GHz
- 128MB L3
- 8 memory channels
- 1024GB of RAM
- SATA SSD (960GB)
- Mellanox ConnectX-6 HDR InfiniBand 100Gb/s Adapter

UC Advanced Research Computing Center (ARC) and Core Facility

Core Facility Services and Rates

Node Type	Processors per node	Cores per nodes	Memory per node	InfiniBand	Accelerator card	Speed	UC / Collaborators Rates per Core Hr	Industry Partner Rates
Regular Memory CPU	2 AMD EPYC 7452 CPUS	64	256GB RAM	100 GB/s	N/A	2.35-3.35 GHz	$0.003	Contact us
GPU	2 AMD EPYC 7452 CPUS	64	1024GB RAM	100 GB/s	2 NVIDIA Tesla A100-40GB GPU	2.35-3.35 GHz	$0.18	Contact us
Large Memory CPU	2 AMD EPYC 7452 CPUS	64	1024GB RAM	100 GB/s	N/A	2.35-3.35 GHz	$0.004	Contact us

Storage Type	Use case	Backed up (Y/N)	Cost per TB/Mo
Scratch Storage	High speed Lustre storage intended for data sets that are in active use by the HPC cluster resources	N	Free
Project Storage	High speed Lustre storage intended for data sets that are in active use by the HPC cluster resources	Y	$6.00
Bulk Storage	Available for storing large data sets that might be accessed/analyzed in the near future.	Y	$3.33
Archival Storage (Tape)	Intended for the long-term storage of data sets	Y	$1.67