

Compute Ontario Summer School 2024

Pawel Pomorski
University of Waterloo
SHARCNET, Compute Ontario, DRAC

Summer school essentials

June 3 - June 21, 2024

3 weeks Monday to Friday, 9 am - 4:30 pm, morning and afternoon sessions

Offered online via Zoom and Moodle

Open registration, anyone can register

Website: <https://training.computeontario.ca/coss2024.php>

Registration opens 2 pm Monday May 13, 2024

History

Annual Summer Schools offered in various formats since 2002

https://helpwiki.sharcnet.ca/wiki/Summer_Schools

Initially in person, since 2020 online

Since 2020 schools offered by individual Ontario consortia merged into one online Compute Ontario Summer School

Last year's website:

<https://training.computeontario.ca/coss2023.php>

Accounts

Registered students will be issued guest accounts

If you have (or can get) an Alliance account, that is preferable, as it will allow you to continue your work on Alliance clusters after Summer School.

A few courses require an Alliance account.

Accounts are free to all Canadian university researchers. Students will need their supervisors to sponsor them.

<https://alliancecan.ca/en/services/advanced-research-computing/account-management/apply-account>

Registration

First come, first serve.

Default class size limit is 250, but some courses have lower limit if there are hardware limitations.

Courses offered in 2 streams, so make sure you don't register into two courses that overlap.

Very important: if you register and find you cannot attend, please unregister to free up space for other students.

Hardware

Summer School is built around hands on activities for students

Virtual CPU cluster provided via Magic Castle

Reservations on Alliance clusters for courses using GPUs

Instructors

Most of the instructors are from Ontario supercomputing consortia: CAC, SCINET and SHARCNET

Also have guest instructors from outside Ontario consortia (MathWorks, university faculty)

Over 40 instructors in total

Certificates

Issued for individual courses

Attendance at live session required to obtain certificate

Longer courses will have in class activities or assignments which must be completed satisfactorily to obtain the certificate.

Courses: Advanced Research Computing (ARC)

Introduction to Linux shell

Introduction to Advanced Research Computing

Introduction to Version control (Git)

Introduction to Python

High Performance Computing in Python

R

Introduction to C

Modern C++ Parallel Programming

Parallel Computing with MATLAB

Multicore parallel programming (OpenMP)

GPU programming: CUDA

SQL

Data Security

High-Performance I/O and Storage

Working with Jupyter on the clusters

Introduction to Scalable and Accelerated Data Analytics

Scientific visualization

Using Containers: Apptainer

Leveraging HPC for Computational Fluid Dynamics

Courses: Machine Learning (ML)

AI showcase

Text Mining

Machine Learning with MATLAB

Data Preparation

Machine Learning

Artificial Neural Networks (aka Deep Learning)

oneAPI library and programming model for image inferencing for both CPU and GPU

Courses: Research Data Management (RDM)

Research Data Management: Rationale for Reproducibility

From the I-Ching to ChatGPT: A Brief History of AI and Some Historical and Current Applications

Using generative AI tools for research data management

Introduction to Alliance RDM Services

Supporting research with Data Management Plans & the DMP Assistant!

Empowering Open Science: An Introduction to Depositing and Sharing Research Data and Code in Borealis

Academic Libraries and Machine Learning: Transforming the Library

Using Odesi for Survey and Public Opinion Research

Reproducible Research - Practices and Tools

Courses: Domain

Bioinformatics: Analysis of RNA-sequencing data

Bioinformatics: Long-read sequencing applications

Bioinformatics - Introduction and Metagenomics

Questions

Many instructors are present in this meeting, so feel free to ask about specific courses.