Team Altair

Bernie Boscoe, Southern Oregon University

Team Mentors : Veronica Vergara & Mohamed Elbakary

Team Theme Song: New Order, Thieves like us remix (1987)
https://soundcloud.com/markaymufc/new-order-thieves-like-us-mk-instrumental-cover-kleptomaniac-mix

Goals:

To add a module to an undergraduate Intro to Data Science course that demonstrates how to use Jupyter Notebooks in the cloud, with a large dataset, and if I can, GPUs to train an ML model that would not be possible to do without a GPU-enabled device. Outcomes would be an understanding of accessing cloud interfaces, basic terminal commands, an overview of the Jupyter notebook as both a local and cloud tool, and if possible, how to test if GPUs are being seen. Update: possibly using JetStream2

What I need help with: what resources have Jupyter notebooks with GPU option? How can we all share a space, for example for 25 students? How do I handle accounts? How can we load/make available a dataset for them to access?

https://github.com/bboscoe/gateways23

Team Altair: Syllabus

In addition to pre-existing Course Goals and Outcomes:

Course Goals
– Students will explore an HPC cluster environment via a browser-based platform and terminal, and familiarize themselves with large compute infrastructures for scientific exploration

–Students will use Jupyter notebooks to load a pre-trained machine learning model and evaluate its performance with their own datasets and code.

–Students will develop an intuition for what resources are needed depending on the size and complexity of the data, and how much training is desired

Team Altair: Sample Exercises

1. Taking a Jupyter notebook they ran on Colab that shows GPU detection code.
2. Loading the Cassava data from storage and training their own model using the notebook provided by instructor
3. Downloading model to their Colab or local computer and write code to test their own images

If I can get all of these things to work without issue I will eat my hat

Team Altair: Gateways

ACCESS
TACC
SciServer (JHU)