Setting Up AIR-MS on Minerva OnDemand for Training Session II

September 19th 2025

Icahn School of Medicine at Mount Sinai

To prepare for AIRMS Training Session I, please follow the step-by-step setup guide. The guide will walk you through:

- Verifying prerequisites (accounts, software, connections)
- Accessing Minerva
- Launching our Training Jupyter Notebook
- Launching ollama on Minerva

The Jupyter Notebook will serve as a follow-along resource during the training session, so it is important to confirm access in advance. We recommend that you complete this setup a few days before the session to ensure everything is working smoothly and to allow time for troubleshooting if needed.

Pre-requisites

1. Possess a Mount Sinai School Network account
2. Request and obtain a Minerva account
3. Have an active Symantec VIP 2-factor authenticator (download here)
4. Request and obtain access to AIR-MS MSDW DEID Dataset
5. Be either onsite on Mount Sinai Network (NOT Guest Wi-Fi) or connected to VPN

For more information on how to fulfill these steps visit our getting started page

Step 1: Launch a Jupyter Lab via OnDemand

1. **Navigate to OnDemand via your browser**
2. **Sign in with your Mount Sinai credentials and password**
3. **Click on “Interactive Apps” and then select “Jupyter”**

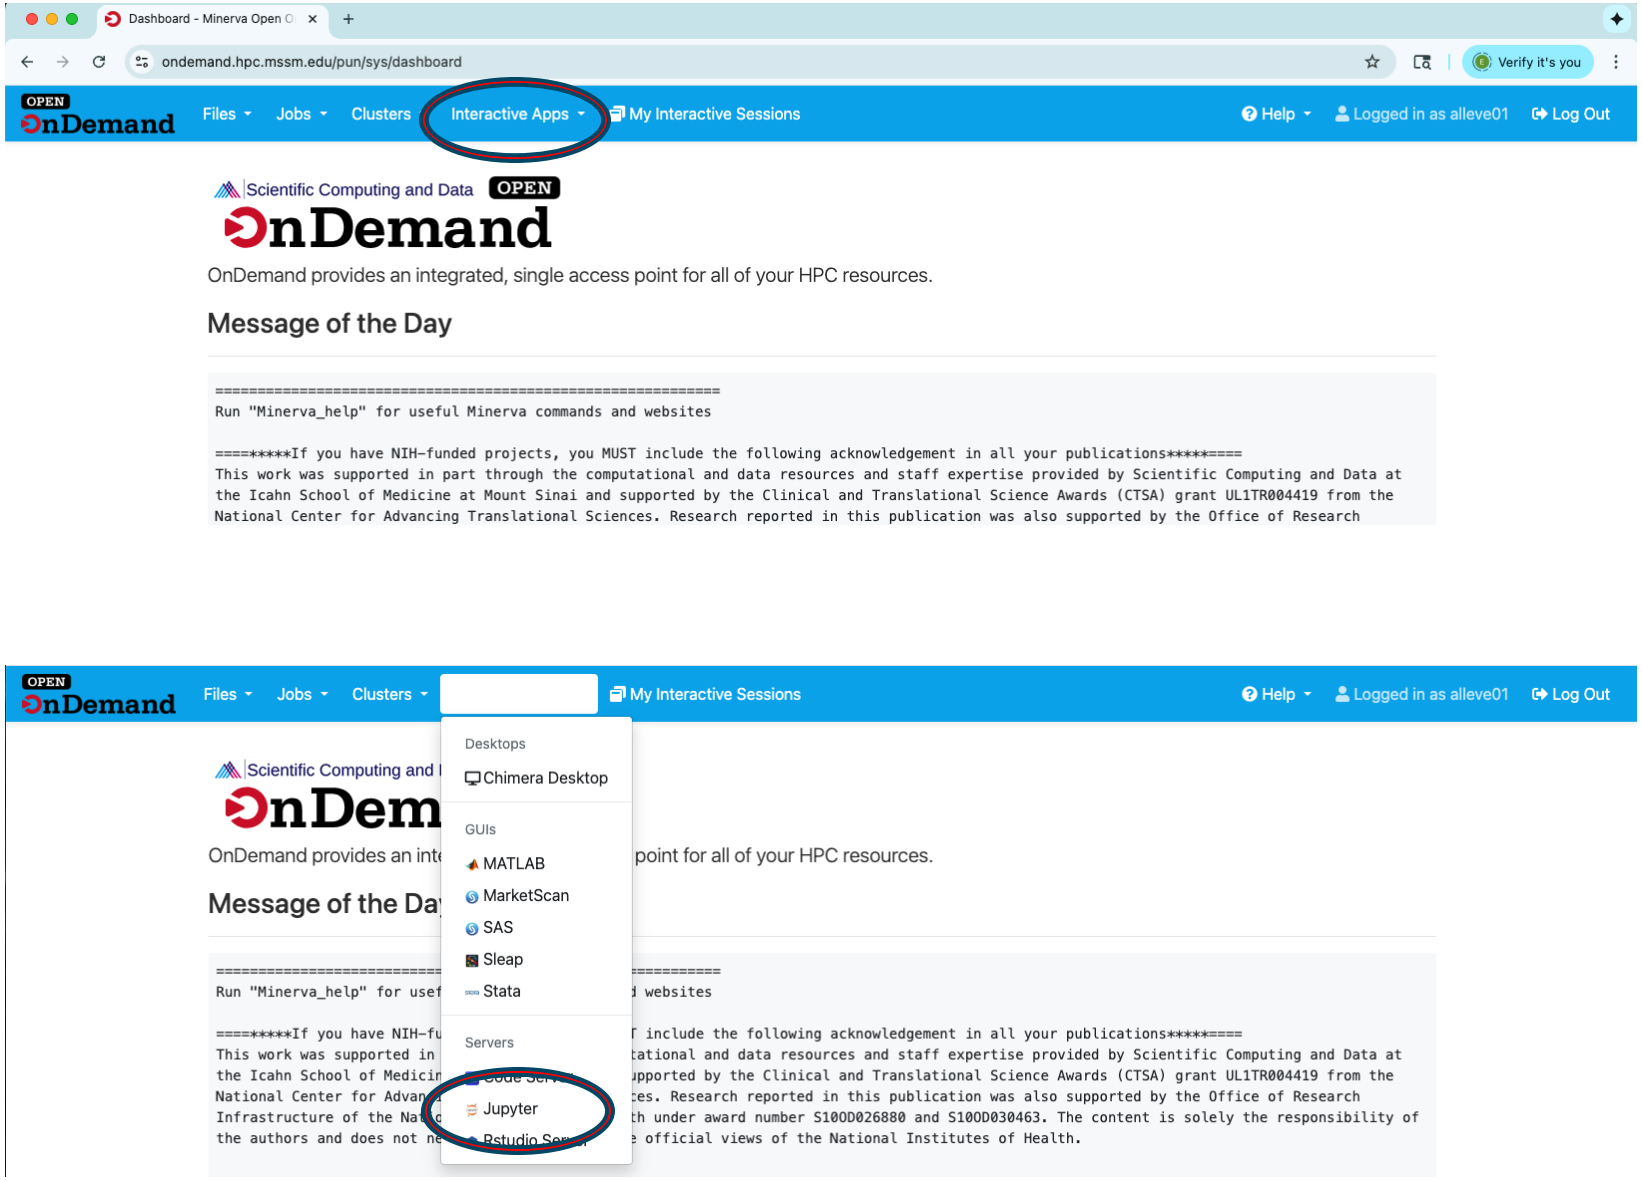
Dashboard - Minerva Open O
ondemand.hpc.mssm.edu/pun/sys/dashboard
Verify it's you
OPEN OnDemand
Files
Jobs
Clusters
Interactive Apps
My Interactive Sessions
Help
Logged in as alleve01
Log Out
Scientific Computing and Data
OPEN OnDemand
OnDemand provides an integrated, single access point for all of your HPC resources.
Message of the Day
Run "Minerva_help" for useful Minerva commands and websites
This work was supported in part through the computational and data resources and staff expertise provided by Scientific Computing and Data at the Icahn School of Medicine at Mount Sinai and supported by the Clinical and Translational Science Awards (CTSA) grant UL1TR004419 from the National Center for Advancing Translational Sciences. Research reported in this publication was also supported by the Office of Research
Desktops
Chimera Desktop
GUIs
MATLAB
MarketScan
SAS
Sleap
Stata
Servers
Jupyter
Rstudio Server

Step 1: Launch a Jupyter Lab via OnDemand

1. Navigate to OnDemand via your browser
2. Sign in with your Mount Sinai credentials and password
3. Click on “Interactive Apps” and then select “Jupyter”
4. **Now fill out the form with the following and press “Launch”:**
   - **Queue:** Express <= 12 hrs
   - **Project Account:** The name of your Minerva project allocation, usually acc_<project name>. If you recently requested an account and do not have a project it might be something like acc_<your PI username>
   - **Mode:** Jupyter Lab
   - **Working Directory:** /sc/arion/work/<INSERT YOUR USER NAME HERE>
   - **Number of cores:** 1
   - **Memory request (in GB):** 1
   - **Number of hours:** 1 (you might want to select 2 during the training session and 1 for testing)
   - **Python version:** Python 3
   - **Extra Modules:** leave blank
   - **Reservation ID**: leave blank

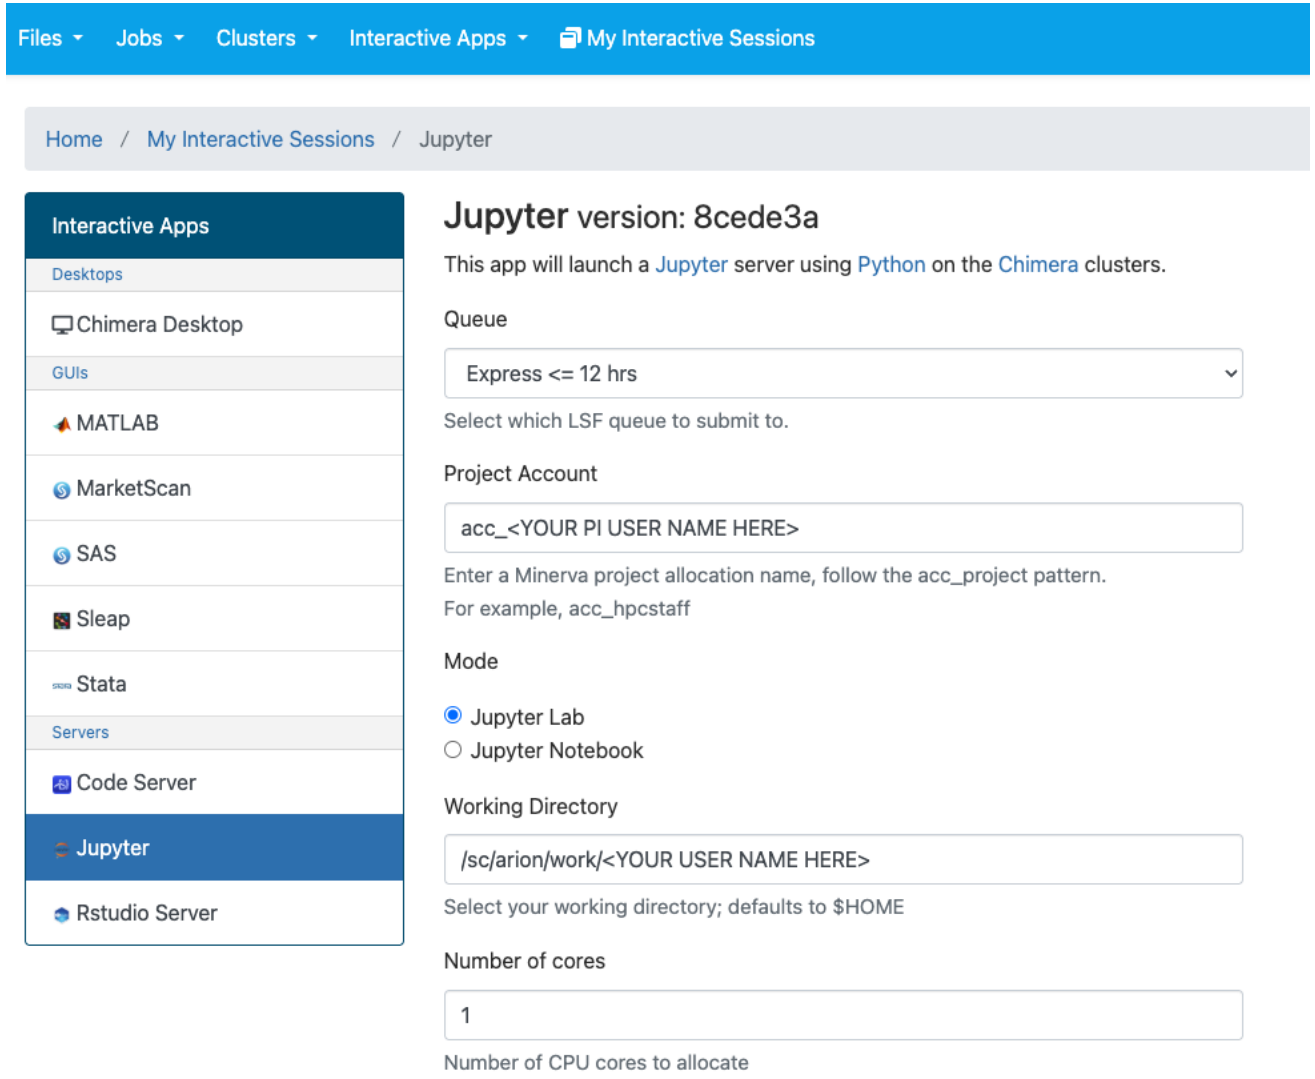
Files
Jobs
Clusters
Interactive Apps
My Interactive Sessions
Home / My Interactive Sessions / Jupyter
Interactive Apps
Desktops
Chimera Desktop
GUIs
MATLAB
MarketScan
SAS
Sleap
Stata
Servers
Code Server
Jupyter
Rstudio Server
Jupyter version: 8cede3a
This app will launch a Jupyter server using Python on the Chimera clusters.
Queue
Express <= 12 hrs
Select which LSF queue to submit to.
Project Account
acc_<YOUR PI USER NAME HERE>
Enter a Minerva project allocation name, follow the acc_project pattern.
For example, acc_hpcstaff
Mode
Jupyter Lab
Jupyter Notebook
Working Directory
/sc/arion/work/<YOUR USER NAME HERE>
Select your working directory; defaults to $HOME
Number of cores
1
Number of CPU cores to allocate

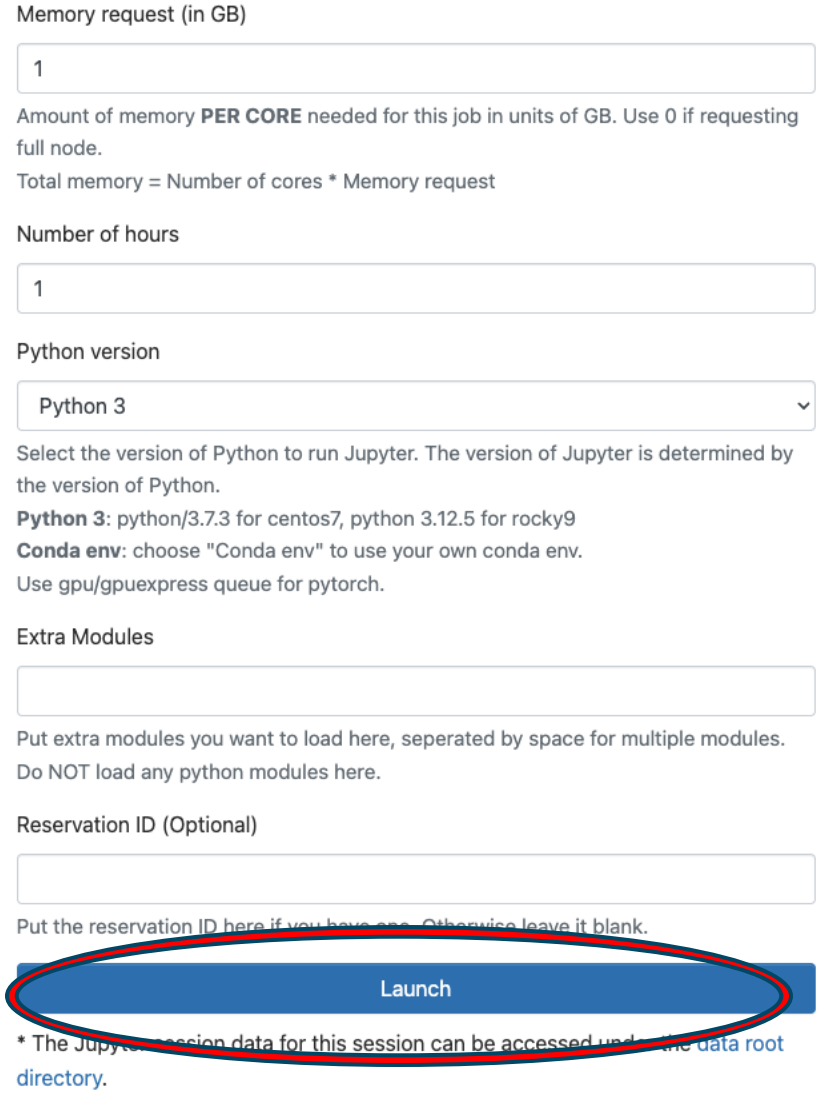
Memory request (in GB)
1
Amount of memory PER CORE needed for this job in units of GB. Use 0 if requesting full node.
Total memory = Number of cores * Memory request
Number of hours
1
Python version
Python 3
Select the version of Python to run Jupyter. The version of Jupyter is determined by the version of Python.
Python 3: python/3.7.3 for centos7, python 3.12.5 for rocky9
Conda env: choose "Conda env" to use your own conda env.
Use gpu/gpuexpress queue for pytorch.
Extra Modules
Put extra modules you want to load here, seperated by space for multiple modules.
Do NOT load any python modules here.
Reservation ID (Optional)
Put the reservation ID here if you have one. Otherwise leave it blank.
Launch
* The Jupyter session data for this session can be accessed under the data root directory.

Step 1: Launch a Jupyter Lab via OnDemand

1. Navigate to OnDemand via your browser
2. Sign in with your Mount Sinai credentials and password
3. Click on “Interactive Apps” and then select “Jupyter”
4. Now fill out the form with the following and press “Launch”:
   - Queue: Express <= 12 hrs
   - Project Account: The name of your Minerva project allocation, usually acc_<project name>. If you recently requested an account and do not have a project it might be something like acc_<your PI username>
   - Mode: Jupyter Lab
   - Working Directory: /sc/arion/work/<INSERT YOUR USER NAME HERE>
   - Number of cores: 1
   - Memory request (in GB): 1
   - Number of hours: 1 (you might want to select 2 during the training session and 1 for testing)
   - Python version: Python 3
   - Extra Modules: leave blank
   - Reservation ID: leave blank
5. **Wait for the session to be in “Running” status and then press the “Connect to Jupyter” button**

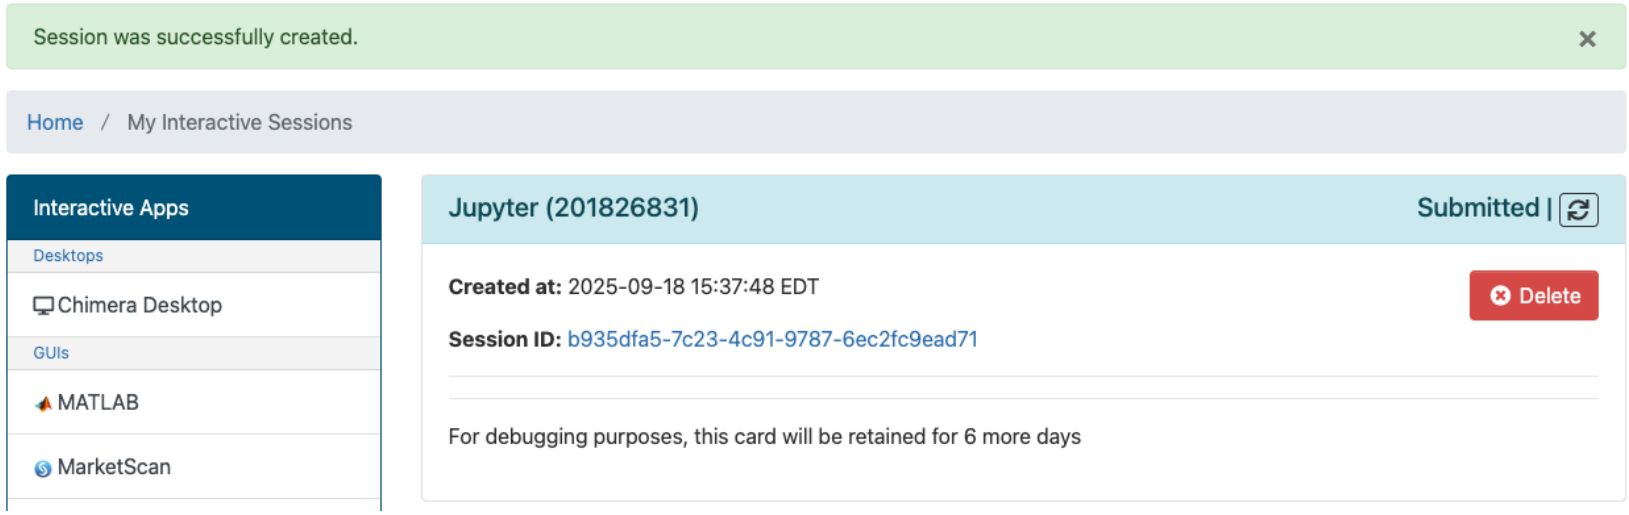

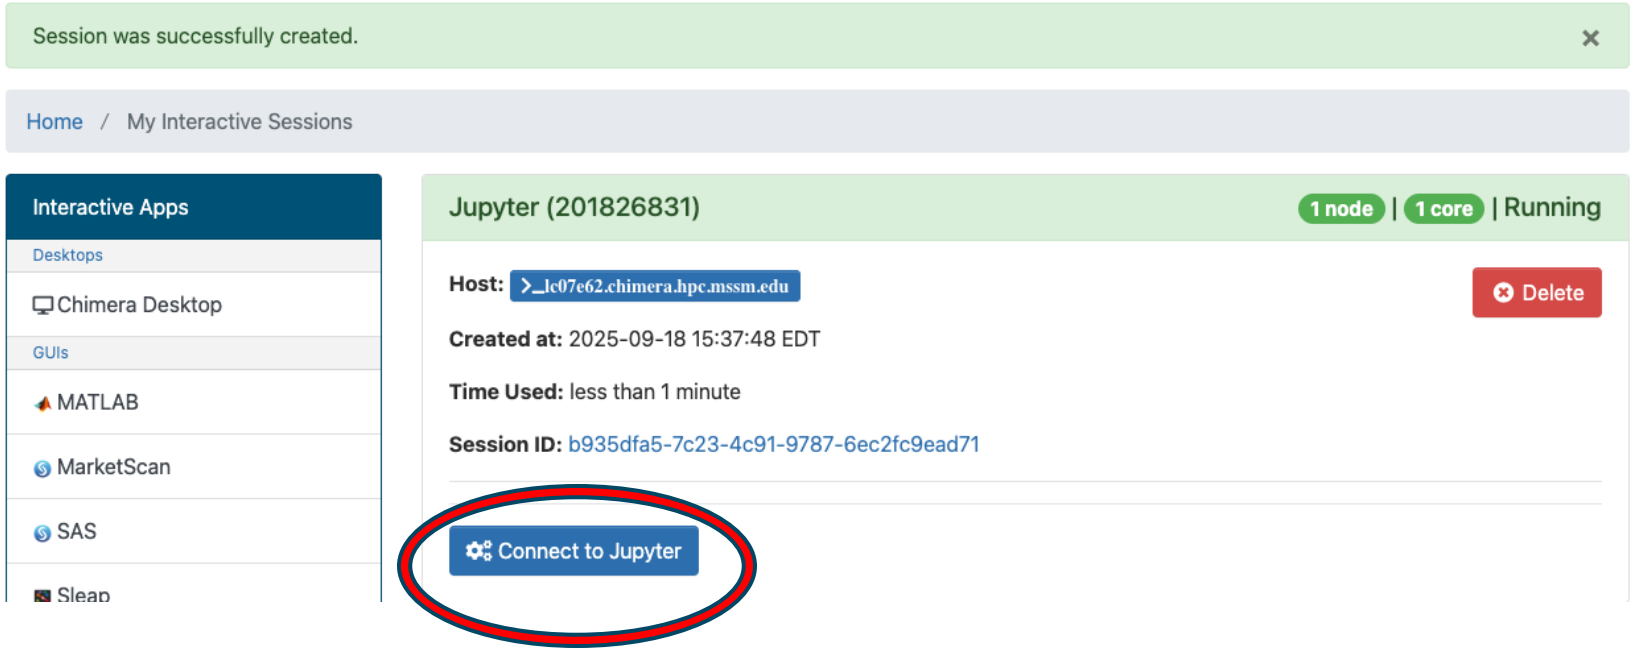

For help you can send an email to airms-support@mssm.edu

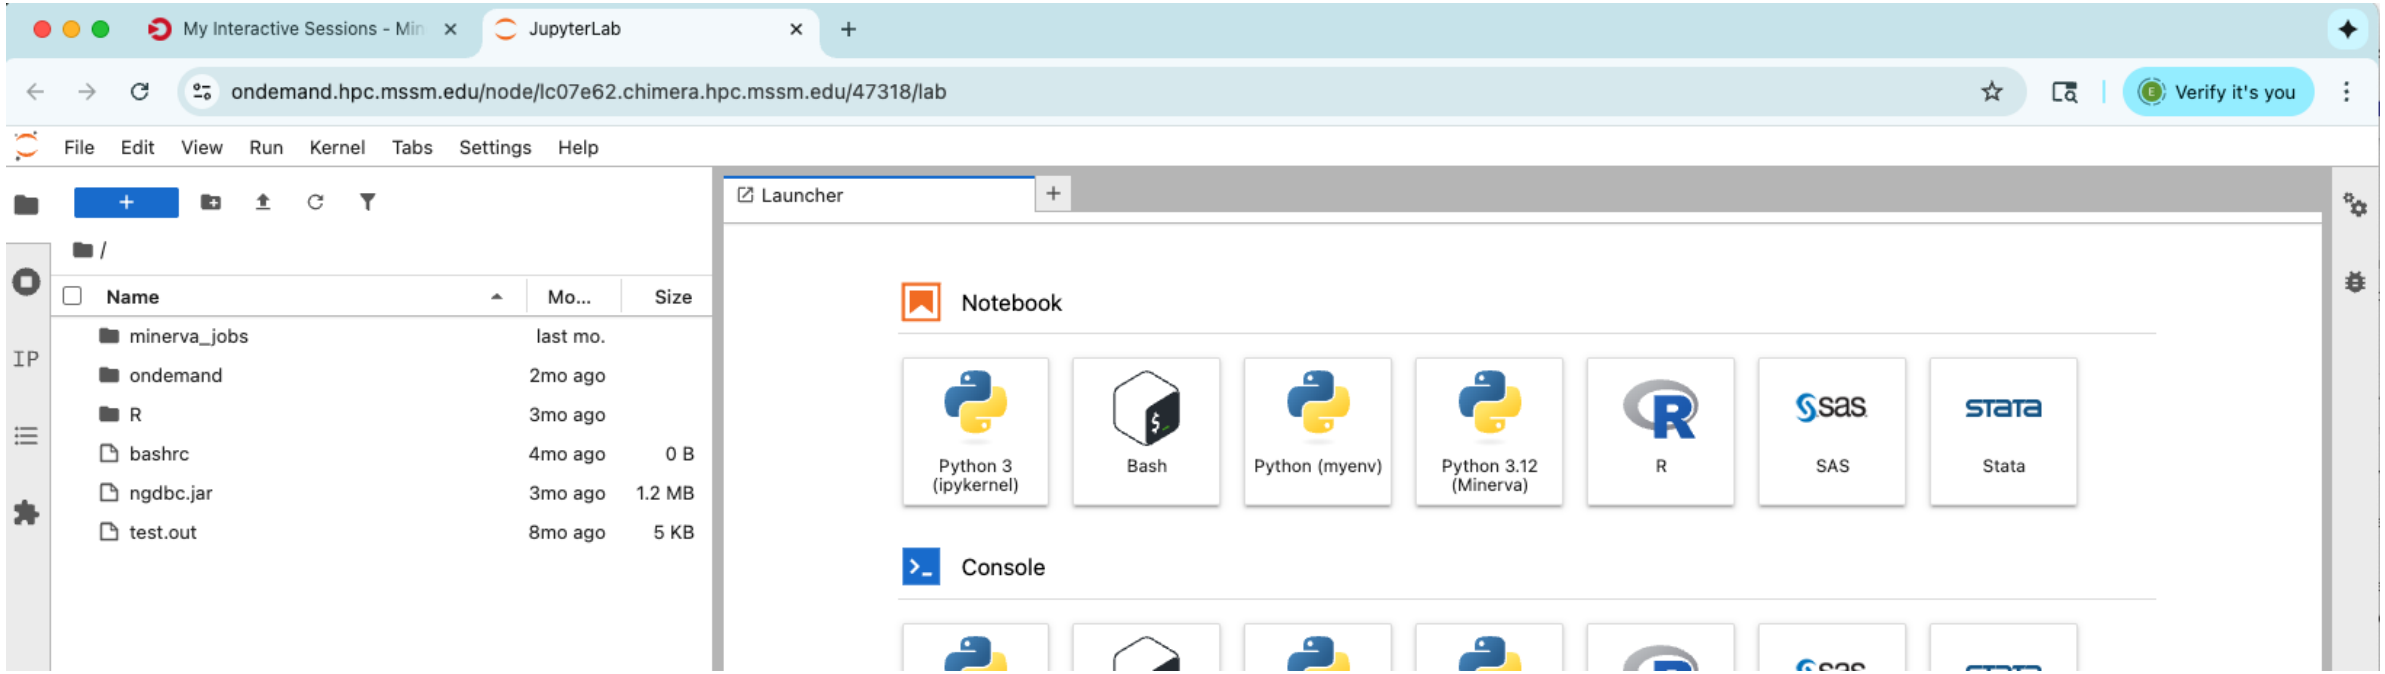
My Interactive Sessions - Mir
JupyterLab
ondemand.hpc.mssm.edu/node/lc07e62.chimera.hpc.mssm.edu/47318/lab
Verify it's you
File Edit View Run Kernel Tabs Settings Help
/
Name
Mo...
Size
minerva_jobs last mo.
ondemand 2mo ago
R 3mo ago
bashrc 4mo ago 0 B
ngdbc.jar 3mo ago 1.2 MB
test.out 8mo ago 5 KB
IP
Launcher
Notebook
Python 3 (ipykernel)
Bash
Python (myenv)
Python 3.12 (Minerva)
R
SAS
Stata
Console

Step 2: Run the Jupyter Notebook

1. **Download our example Jupyter Notebook at this link and drag & drop or upload (upper arrow icon button) it in your Jupyter Lab. Click on it once it appears on the left panel**

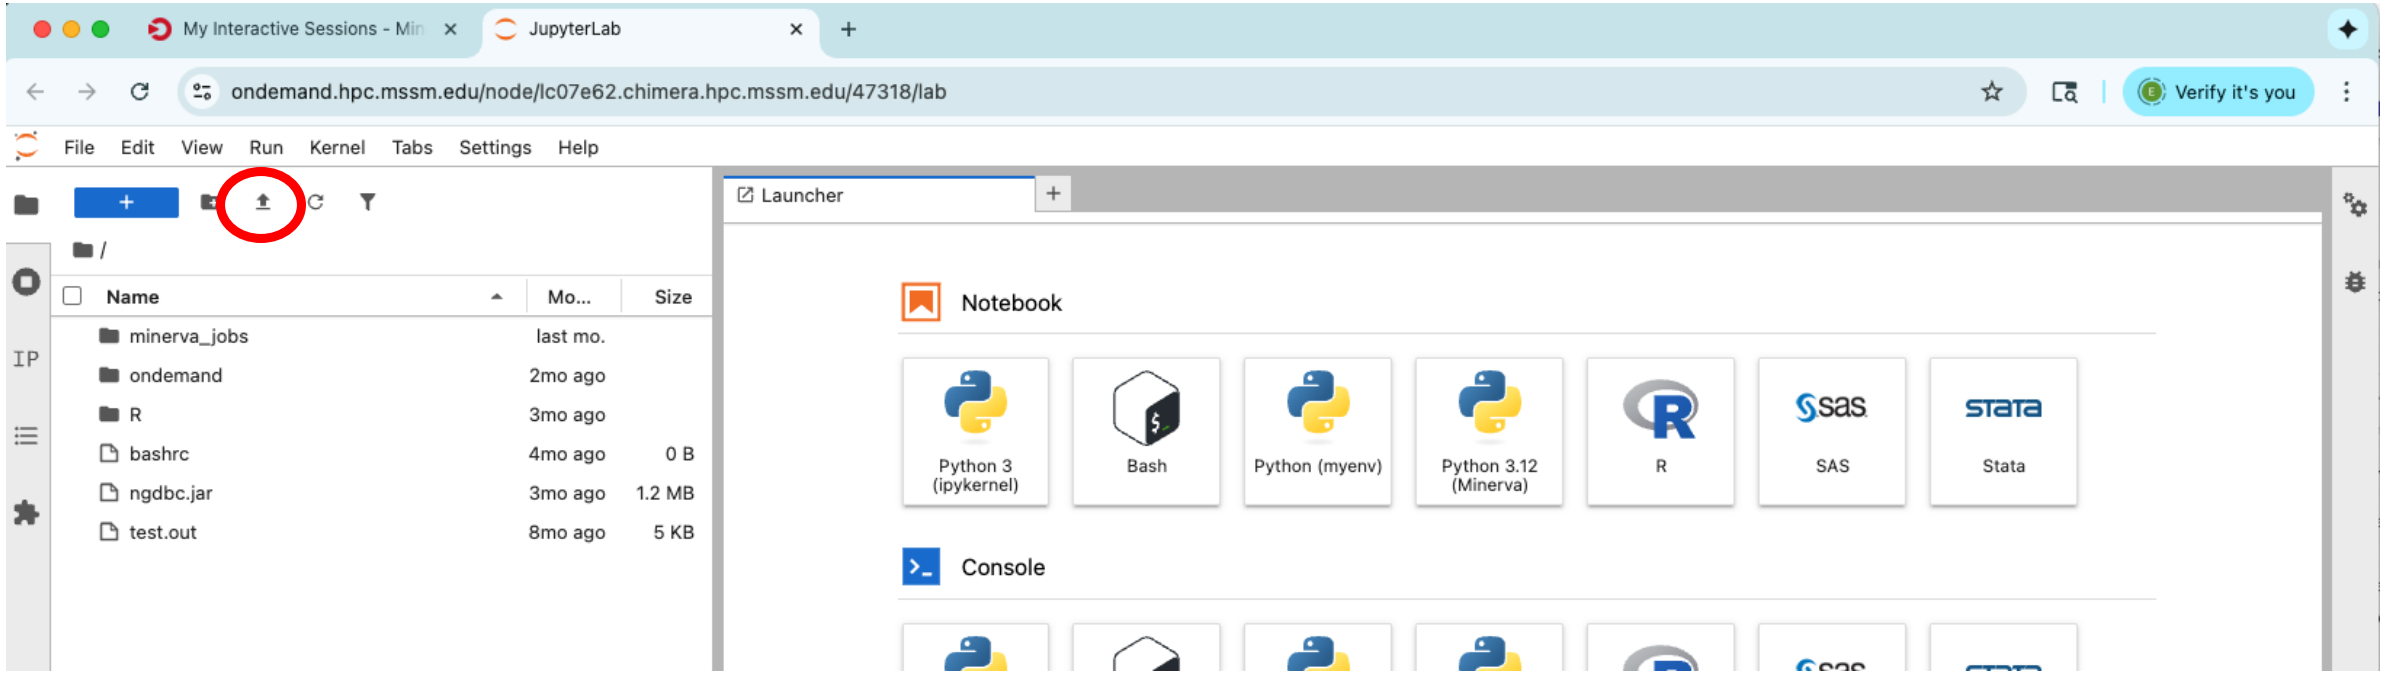
My Interactive Sessions - Mir
JupyterLab
ondemand.hpc.mssm.edu/node/lc07e62.chimera.hpc.mssm.edu/47318/lab
Verify it's you
File Edit View Run Kernel Tabs Settings Help
/
Name
Mo...
Size
minerva_jobs last mo.
ondemand 2mo ago
R 3mo ago
bashrc 4mo ago 0 B
ngdbc.jar 3mo ago 1.2 MB
test.out 8mo ago 5 KB
IP
Launcher
Notebook
Python 3 (ipykernel)
Bash
Python (myenv)
Python 3.12 (Minerva)
R
SAS
Stata
Console

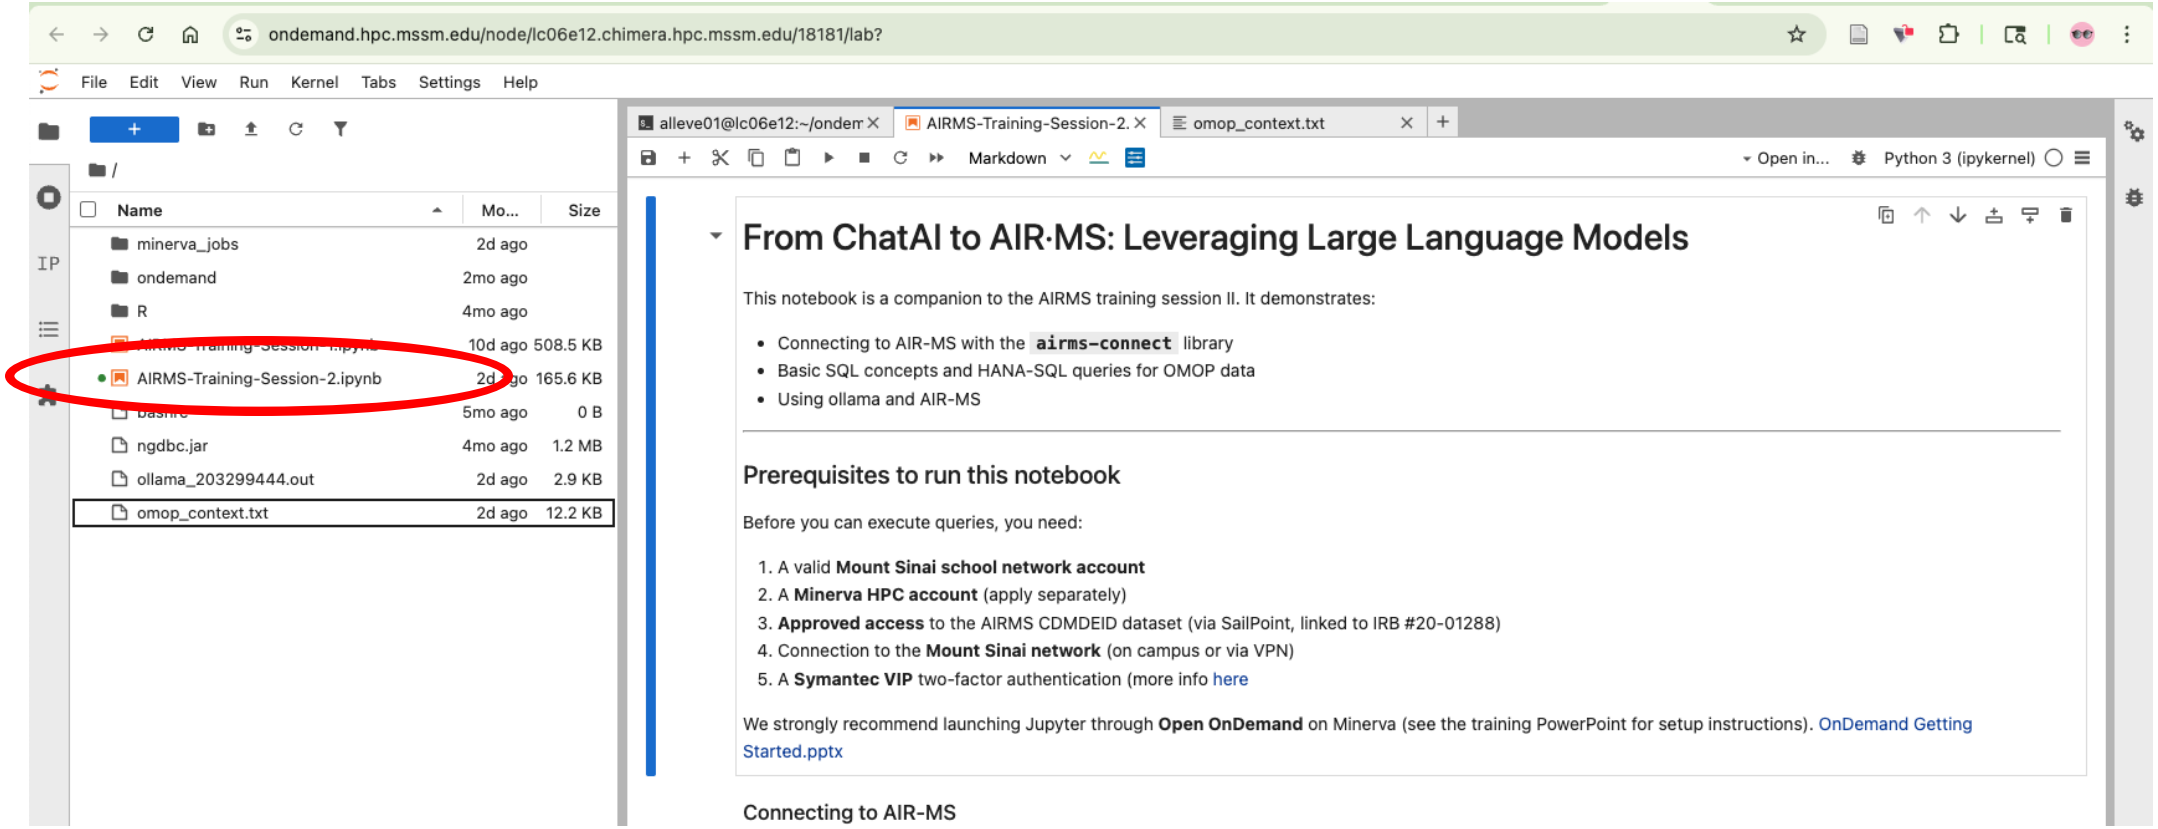
ondemand.hpc.mssm.edu/node/lc06e12.chimera.hpc.mssm.edu/18181/lab?
File Edit View Run Kernel Tabs Settings Help
alleve01@lc06e12:~/onden
AIRMS-Training-Session-2.
omop_context.txt
Markdown
Open in...
Python 3 (ipykernel)
Name
Mo...
Size
minerva_jobs 2d ago
ondemand 2mo ago
R 4mo ago
10d ago 508.5 KB
AIRMS-Training-Session-2.ipynb 165.6 KB
5mo ago 0 B
ngdbc.jar 4mo ago 1.2 MB
ollama_203299444.out 2d ago 2.9 KB
omop_context.txt 2d ago 12.2 KB
From ChatAI to AIR·MS: Leveraging Large Language Models
This notebook is a companion to the AIRMS training session II. It demonstrates:
Connecting to AIR-MS with the airms-connect library
Basic SQL concepts and HANA-SQL queries for OMOP data
Using ollama and AIR-MS
Prerequisites to run this notebook
Before you can execute queries, you need:
1. A valid Mount Sinai school network account
2. A Minerva HPC account (apply separately)
3. Approved access to the AIRMS CDMDEID dataset (via SailPoint, linked to IRB #20-01288)
4. Connection to the Mount Sinai network (on campus or via VPN)
5. A Symantec VIP two-factor authentication (more info here
We strongly recommend launching Jupyter through Open OnDemand on Minerva (see the training PowerPoint for setup instructions). OnDemand Getting Started.pptx
Connecting to AIR-MS

Step 2: Run the Jupyter Notebook

1. Download our example Jupyter Notebook at this link and drag & drop or upload (upper arrow icon button) it in your Jupyter Lab. Click on it once it appears on the left panel
2. **Run the cells in the "Connecting to AIR-MS" section to connect to AIR-MS**
3. **You will be asked to input**
   - **AIRMS_USER – which corresponds to your Minerva username**
   - **Password - which is your Mount Sinai password (without VIP token)**
4. You are now connected to AIR-MS and should be able to run all the following cells

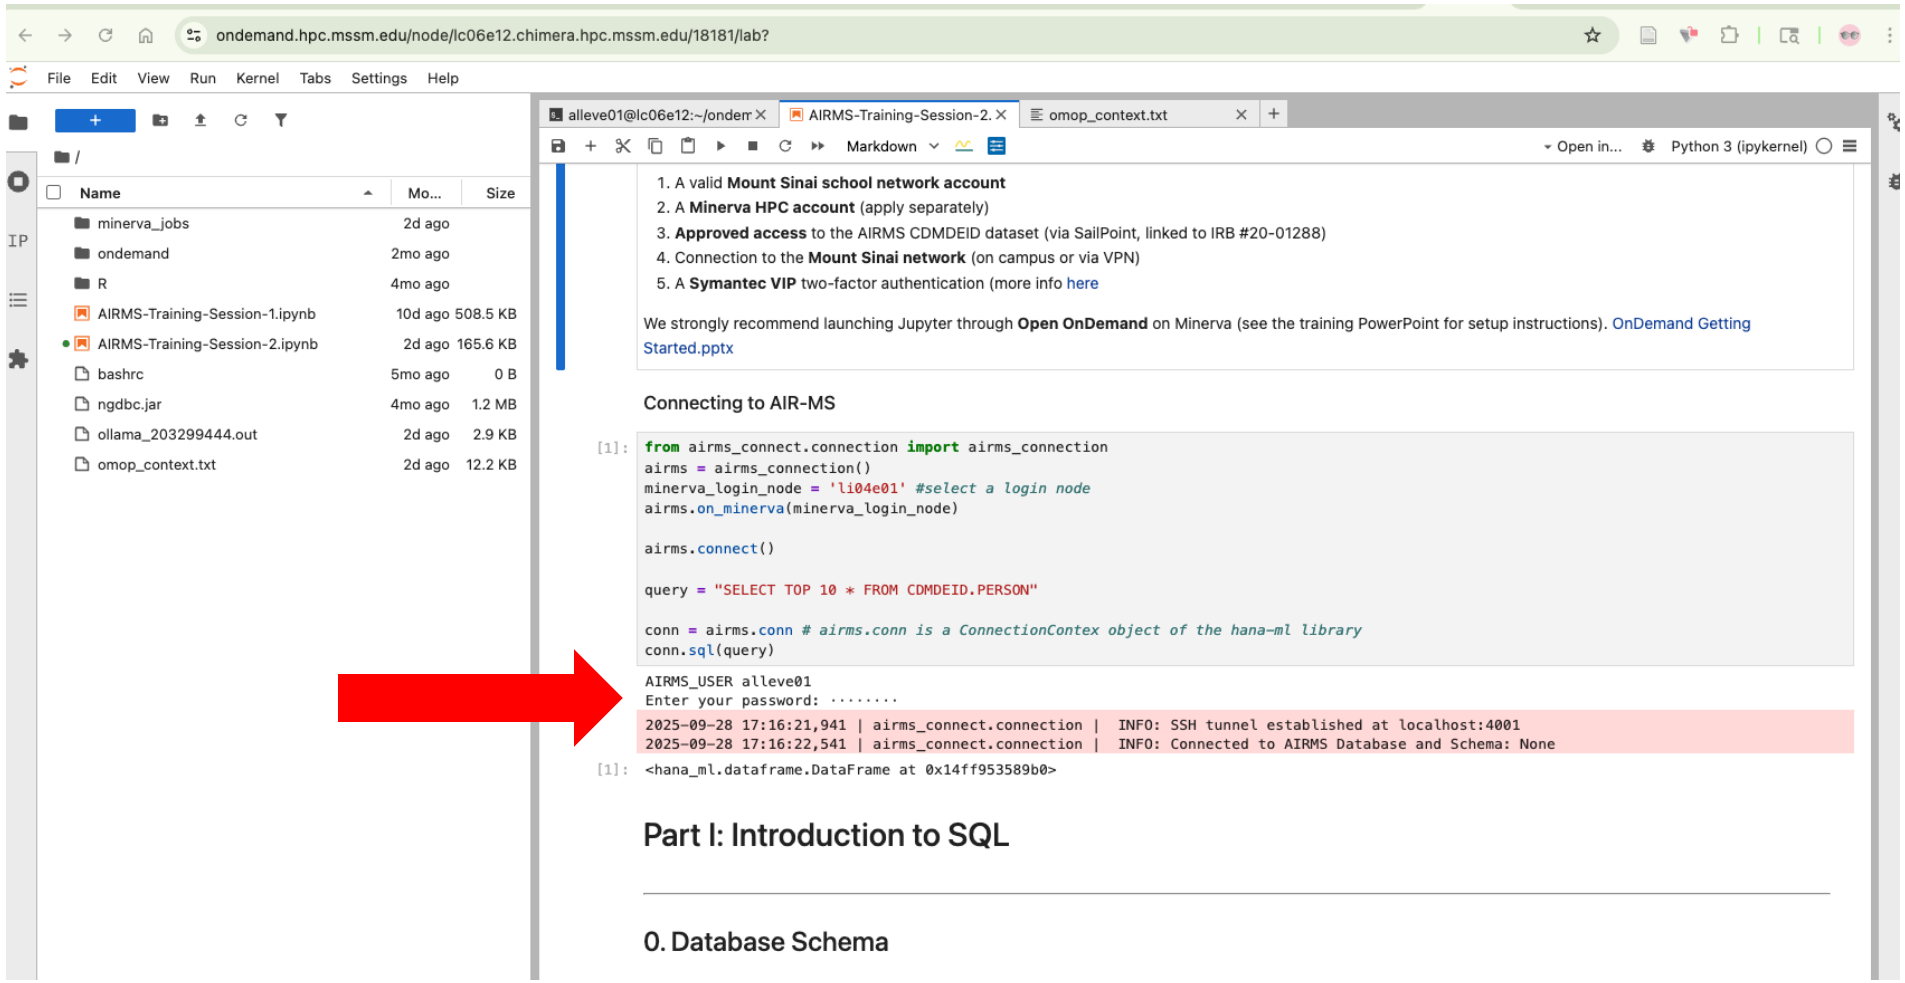
ondemand.hpc.mssm.edu/node/lc06e12.chimera.hpc.mssm.edu/18181/lab?
File Edit View Run Kernel Tabs Settings Help
alleve01@lc06e12:~/onder
AIRMS-Training-Session-2.
omop_context.txt
Markdown
Open in...
Python 3 (ipykernel)
Name
Mo...
Size
minerva_jobs 2d ago
ondemand 2mo ago
R 4mo ago
AIRMS-Training-Session-1.ipynb 10d ago 508.5 KB
AIRMS-Training-Session-2.ipynb 2d ago 165.6 KB
bashrc 5mo ago 0 B
ngdbc.jar 4mo ago 1.2 MB
ollama_203299444.out 2d ago 2.9 KB
omop_context.txt 2d ago 12.2 KB
1. A valid Mount Sinai school network account
2. A Minerva HPC account (apply separately)
3. Approved access to the AIRMS CDMDEID dataset (via SailPoint, linked to IRB #20-01288)
4. Connection to the Mount Sinai network (on campus or via VPN)
5. A Symantec VIP two-factor authentication (more info here
We strongly recommend launching Jupyter through Open OnDemand on Minerva (see the training PowerPoint for setup instructions). OnDemand Getting Started.pptx
Connecting to AIR-MS
[1]:
from airms_connect.connection import airms_connection
airms = airms_connection()
minerva_login_node = 'li04e01' #select a login node
airms.on_minerva(minerva_login_node)
airms.connect()
query = "SELECT TOP 10 * FROM CDMDEID.PERSON"
conn = airms.conn # airms.conn is a ConnectionContex object of the hana-ml library
conn.sql(query)
AIRMS_USER alleve01
Enter your password: ·····
2025-09-28 17:16:21,941 | airms_connect.connection | INFO: SSH tunnel established at localhost:4001
2025-09-28 17:16:22,541 | airms_connect.connection | INFO: Connected to AIRMS Database and Schema: None
[1]:
<hana_ml.dataframe.DataFrame at 0x14ff953589b0>
Part I: Introduction to SQL
0. Database Schema

Step 2: Run the Jupyter Notebook

1. Download our example Jupyter Notebook at this link and drag & drop or upload (upper arrow icon button) it in your Jupyter Lab. Click on it once it appears on the left panel
2. Run the cells in the “Connecting to AIR-MS” section to connect to AIR-MS
3. You will be asked to input
   - AIRMS_USER – which corresponds to your Minerva username
   - Password - which is your Mount Sinai password (without VIP token)
4. **You are now connected to AIR-MS and should be able to run all the following cells**

Setting Up ollama

In **Part II** of our training notebook, we will ask you to start an **ollama server** on Minerva to use open source LLMs.

To do this, follow these steps:

1) **Navigate to clusters>chimera Shell access**

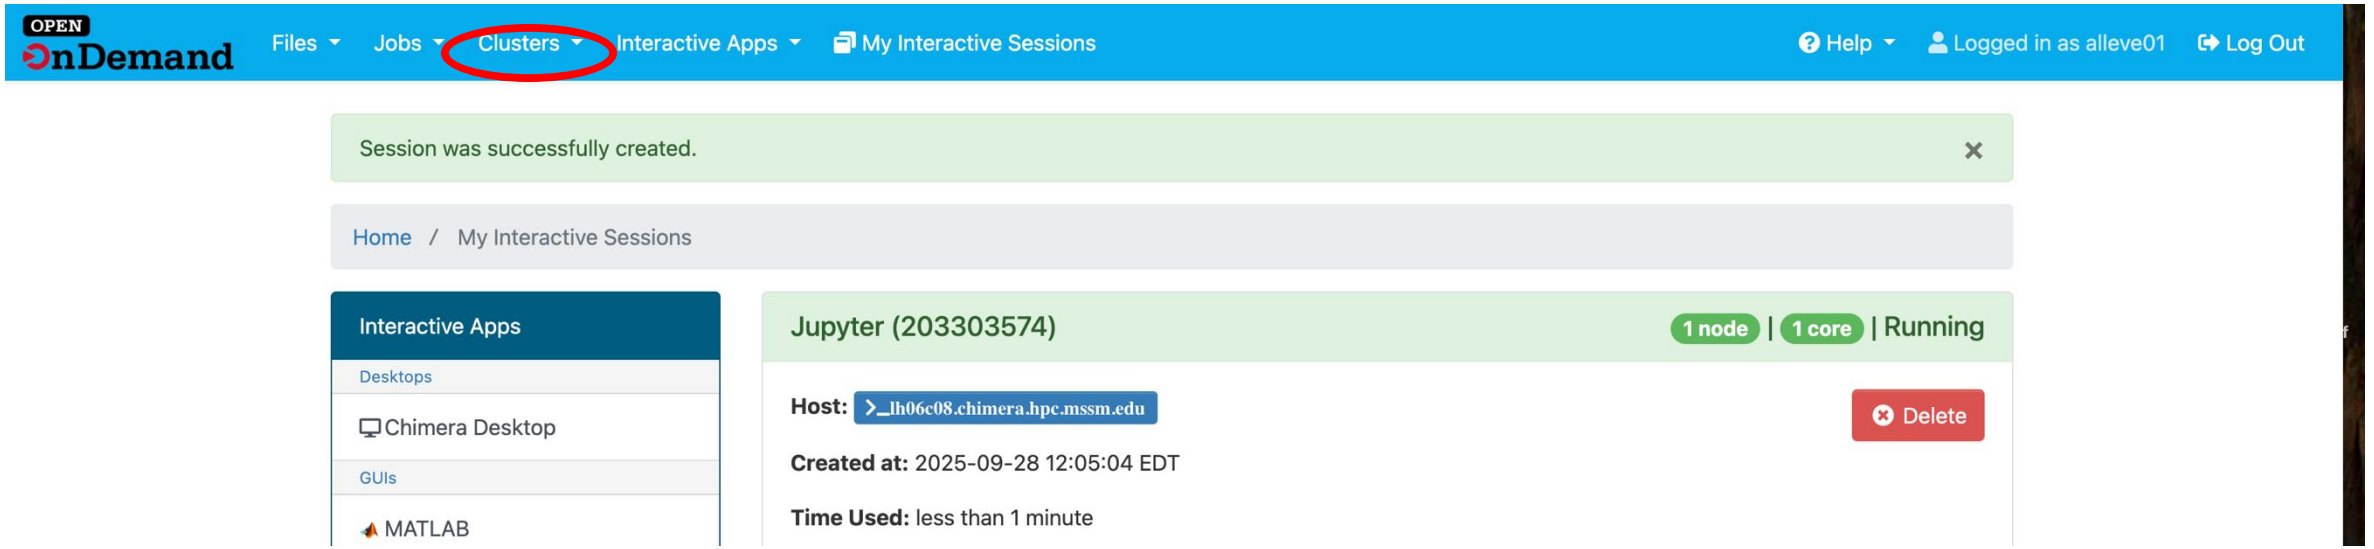
OPEN OnDemand
Files
Jobs
Clusters
Interactive Apps
My Interactive Sessions
Help
Logged in as alleve01
Log Out
Session was successfully created.
Home / My Interactive Sessions
Interactive Apps
Desktops
Chimera Desktop
GUIs
MATLAB
Jupyter (203303574)
1 node | 1 core | Running
Host: >_lh06c08.chimera.hpc.mssm.edu
Delete
Created at: 2025-09-28 12:05:04 EDT
Time Used: less than 1 minute

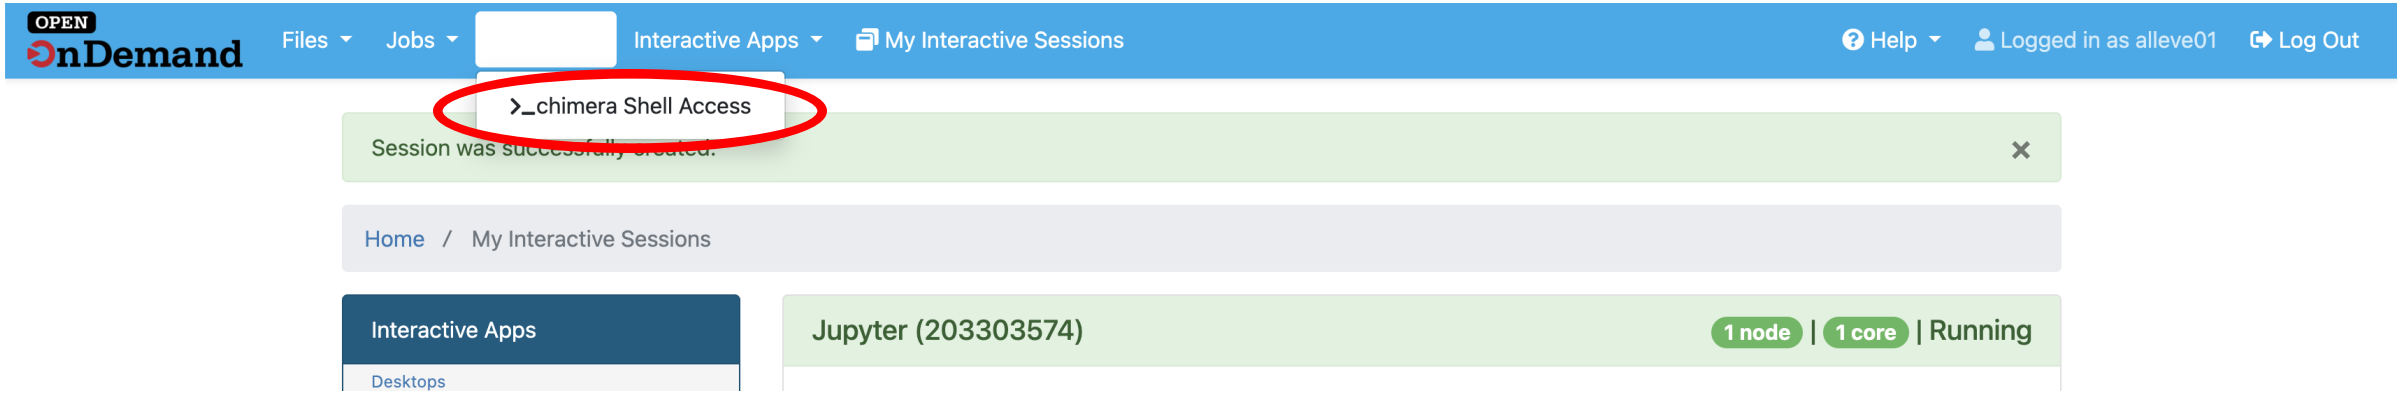
OPEN OnDemand
Files
Jobs
Interactive Apps
My Interactive Sessions
Help
Logged in as alleve01
Log Out
>_chimera Shell Access
Home / My Interactive Sessions
Interactive Apps
Desktops
Jupyter (203303574)
1 node | 1 core | Running

Setting Up ollama

In **Part II** of our training notebook, we will ask you to start an **ollama server** on Minerva to use open source LLMs.

To do this, follow these steps:

1) Navigate to clusters>chimera Shell access
2) **Enter your password followed by the Symantec VIP token and press enter to access Minerva**

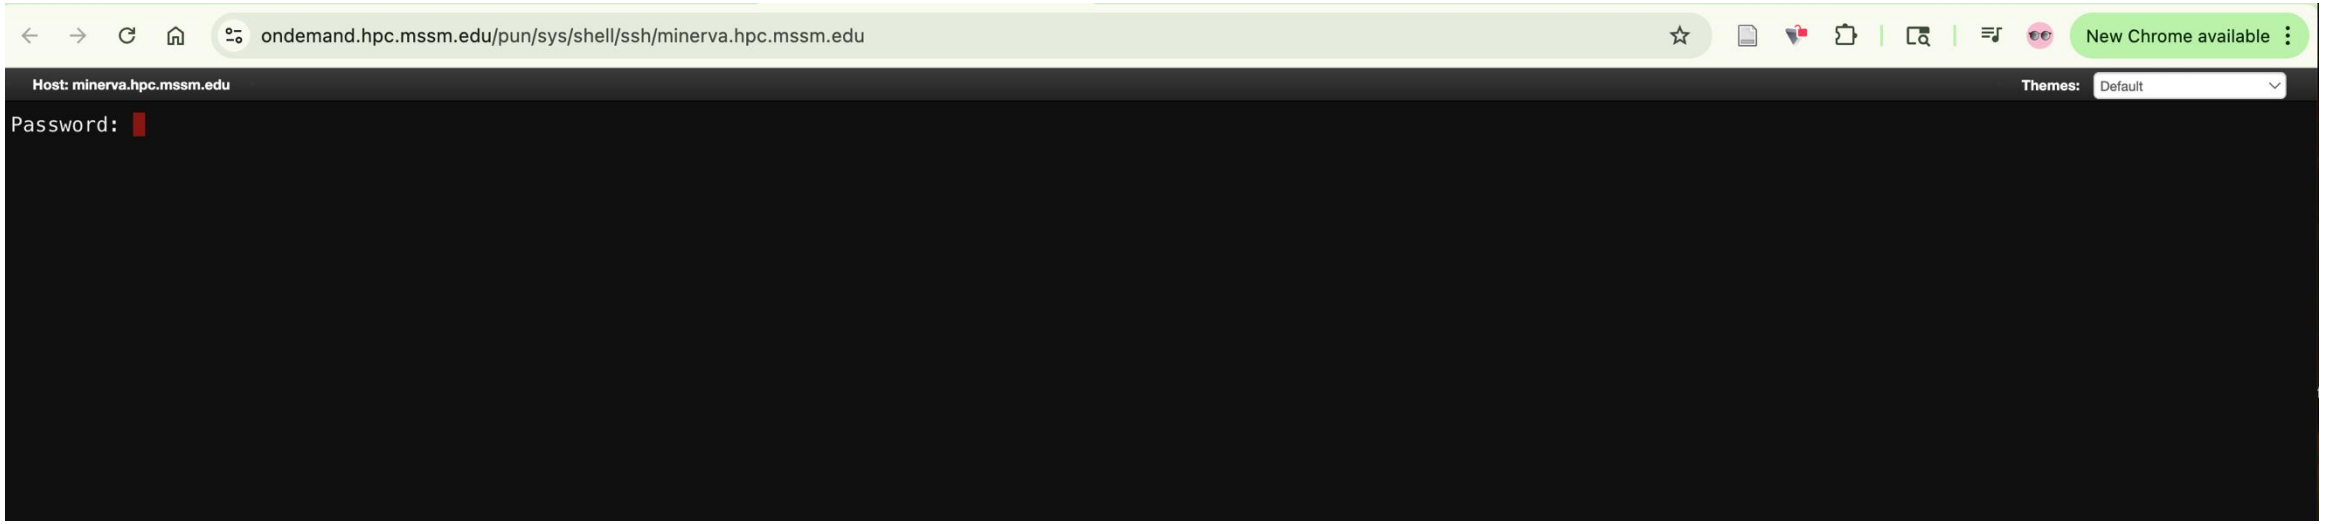


```
=========================Minerva Fall 2025 Training Sessions=========================
There will be 9 Training Sessions (in person and zoom).
Every Tuesday and Friday starting Sep 16 and ending Oct. 14
More info at https://labs.icahn.mssm.edu/minervalab/resources/the-minerva-user-group-and-training-classes/
=====================================================================================
    === Send tickets to hpchelp@hpc.mssm.edu ===

!!!!!WE DO NOT BACKUP USER FILES
          !!!!PLEASE ARCHIVE/BACKUP YOUR IMPORTANT FILES
============================================================

    Free GPUs
-----------------
queue/model         h100nvl         h10080g           a100           v100           l40s
       gpu          56/188           8/8            11/40          7/40          20/32
 gpuexpress          0/64             8/8            11/40          7/40          20/32
    private          35/40            0/0             4/4           0/0            0/0
```

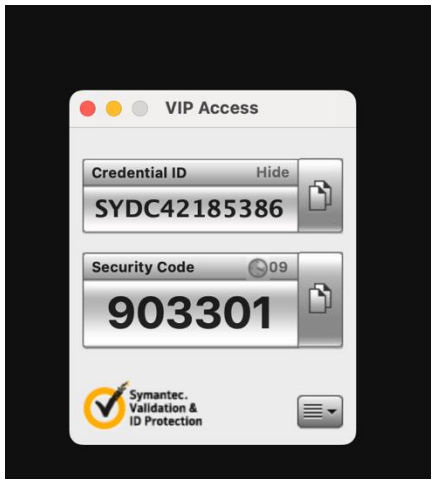

Setting Up ollama

In **Part II** of our training notebook, we will ask you to start an **ollama server** on Minerva to use open source LLMs.

To do this, follow these steps:

1) Navigate to clusters>chimera Shell access
2) Enter your password followed by the Symantec VIP token and press enter to access Minerva
3) **Run minerva-ollama-web.sh –o /sc/arion/work/<your user name> from the terminal and wait until the ollama server is up and running**
4) **Copy the line with ollama_client = Client(host=‘ […] into our training notebook**

```
(base) [alleve01@li04e01 ~]$ minerva-ollama-web.sh -o /sc/arion/work/alleve01/
[INFO] Image not specified, check if previously used
[INFO] Found previously used image /sc/arion/work/alleve01//minerva_jobs/ollama_jobs/ollama_v0.3.10.sif. Using it.
[INFO] Project is not specified, or is acc_null, using 1st avail project.
[INFO] Project to use is acc_pasml
[INFO] Parameters used are:
[INFO] -n       4
[INFO] -M       3000
[INFO] -W       6:00
[INFO] -P       acc_pasml
[INFO] -J       ollama
[INFO] -q       gpu
[INFO] -R       v100
[INFO] -g       1
[INFO] -o       /sc/arion/work/alleve01//minerva_jobs/ollama_jobs
[INFO] -i       /sc/arion/work/alleve01//minerva_jobs/ollama_jobs/ollama_v0.3.10.sif
[INFO] Submitting Ollama job...
[INFO] Job ID: 203303626
Ollama is started on compute node lg03a09, port 9999
[Authentication Info]
User ID: alleve01
Token: 4b3419f634a99264b04bae5be950c8e0
Access URL: http://10.95.46.94:53626

*** You can access Ollama from Python as shown below *****
from ollama import Client
ollama_client = Client(host='http://10.95.46.94:53626', headers={"Authorization": "Bearer alleve01:4b3419f634a99264b04bae5be950c8e0"})
```

```python
# Setting Up OLLAMA on Minerva
from ollama import Client
### PASTE THE COPIED LINES FROM THE TERMINAL HERE
client = Client(host='http://10.95.46.94:54325', headers={"Authorization": "Bearer alleve01:7e9806cffbf739b7a3686e3bd404b639"})
```

Python

Thank you.

For help, please reach out to airms-support@mssm.edu

For additional resources visit:

- Our Researcher Tutorials (Accessible on Mount Sinai Network or VPN)
- Our Website

Icahn School of Medicine at Mount Sinai