

# MATLAB and NVIDIA GPUs with Open OnDemand

## Benefits and Best Practices

October 28, 2021

Alan Chalker – Director of Strategic Programs at the Ohio Supercomputer Center & PI for Open OnDemand

Andy Thé – AI / GPU Partner Manager at MathWorks

# Agenda

1. Overview: Open OnDemand
2. Why MATLAB with NVIDIA GPUs
3. Demo: Benefits of MATLAB + GPUs
4. Best Practices & OOD examples of MATLAB+GPU Configurations
5. Next Steps and Discussion

# Open OnDemand

<http://openondemand.org/>



## An intuitive, innovative, and interactive interface to remote computing resources

Open OnDemand helps computational researchers and students efficiently utilize remote computing resources by making them easy to access from any device. It helps computer center staff support a wide range of clients by simplifying the user interface and experience.

### Key Benefits & Impact

- Key benefit to you, the end user: You can use any web browser to access resources at a computing service provider.
- Key benefit to you, the computer center staff: A wide range of clients/needs can utilize your computing resources.
- Overall impact: Users are able to use remote computing resources faster and more efficiently.

### Key Things to Know

Open OnDemand is an NSF-funded open-source HPC portal in production use at over 200 research computing sites around the world, which provides those sites' clients:

- Plugin-free web experience
- Easy file management
- Command-line shell access
- Job management and monitoring across different batch servers and resource managers
- Graphical desktop environments and desktop applications

# Agenda

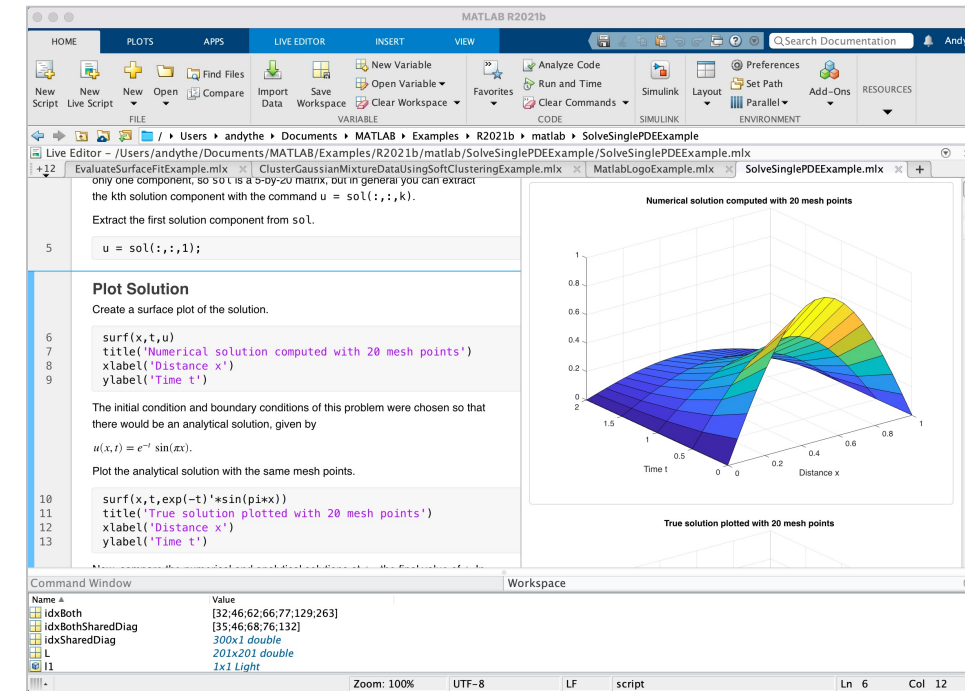
1. Overview: Open OnDemand
2. Why MATLAB with NVIDIA GPUs
3. Demo: Benefits of MATLAB + GPUs
4. Best Practices & OOD examples of MATLAB+GPU Configurations
5. Next Steps and Discussion



# Why MATLAB on Open OnDemand?

MATLAB

- Used at 6500+ colleges and universities
  - 100's -1000's of MATLAB users at each school
- MATLAB is used for science and engineering:
  - Neuroscience
  - Geoscience
  - Computer Science
  - AI / Robotics
  - Imaging / Computer Vision
  - Audio / Signal Processing
  - Wireless / Communications
  - Controls
- Appeal to new and current HPC users with MATLAB on Open OnDemand
  - MATLAB users can reap the benefits of OOD compute resources
  - MATLAB interface vs. Command Line Interface which can be challenging and intimidating



## 2 Types of MATLAB Users

### Casual User

- Happy w/ 5x-10x lift from GPUs
- Little to no GPU experience

#### Examples:

1. MATLAB Simulation: 1 hour on CPUs
  - Cut to 10 minutes with a GPU
2. Training a neural network: 1 day on CPUs
  - Cut to 3 hours with a GPU

### Expert

- Expects >10x lift from GPUs
- GPU experience


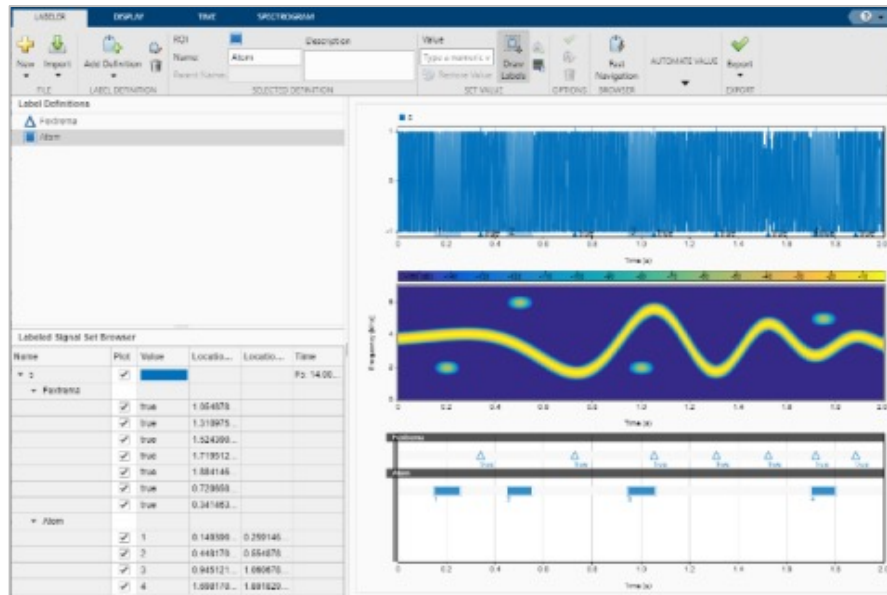
#### Examples:

1. MATLAB Simulation: 1 hours on CPUs
  - Cut to 2 minutes with a GPU
2. Training a neural network: 1 day on CPUs
  - Cut to 30 minutes with a GPU


# Why MATLAB with NVIDIA GPUs?




- Easy access to NVIDIA GPUs with 990+ GPU-enabled functions
- MATLAB users can work with NVIDIA GPUs without CUDA® programming.
- NVIDIA® GPUs accelerate many applications like AI / Deep Learning



# Accelerating MATLAB Applications with NVIDIA GPUs (vs CPUs)




**10x speedup**  
*Deep Learning Training*




**14x speedup**  
*Template Matching Routine*




**12x speedup**  
*Black-Scholes Modeling*



**44x speedup**  
*Simulating Celestial Objects in Motion*



**10x speedup**  
*K-means Clustering Algorithm*



**77x speedup**  
*Wave Equation Solving*

Run MATLAB functions on a GPU

# UT Austin Researchers Convert Brain Signals to Words and Phrases Using Wavelets and Deep Learning

## Challenge

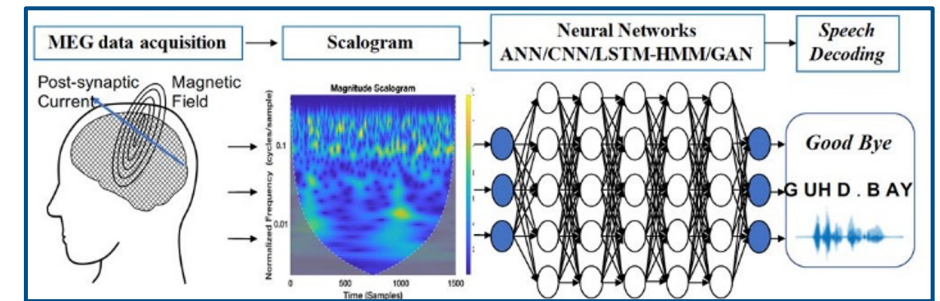
Create a speech-driven brain-computer interface to enable ALS patients to communicate by imagining the act of speaking specific phrases

## Solution

Use wavelet scalograms of MEG signals to train deep neural networks

## Results

- Classification accuracy of 96% achieved
- Wavelets and deep learning networks quickly combined
- Training times **accelerated by a factor of 10 with GPUs**



Classifying the brain signals corresponding to the imagined word “goodbye” using feature extraction and deep neural networks.

*“MATLAB is an industry-standard tool, and one that you can trust. It is easier to learn than other languages, and its toolboxes help you get started in new areas because you don’t have to start from scratch.”*

*- Dr. Jun Wang, UT Austin*



# NASA Langley Accelerates Acoustic Data Analysis with GPU Computing

## Challenge

Accelerate the analysis of sound recordings from wind tunnel tests of aircraft components

## Solution

- Use Parallel Computing Toolbox to process acoustic data
- Cut processing time by running computationally intensive operations on a GPU

## Results

- GPU computations completed **40 times faster**
- Algorithm **GPU-enabled in 30 minutes**
- Processing of test data accelerated

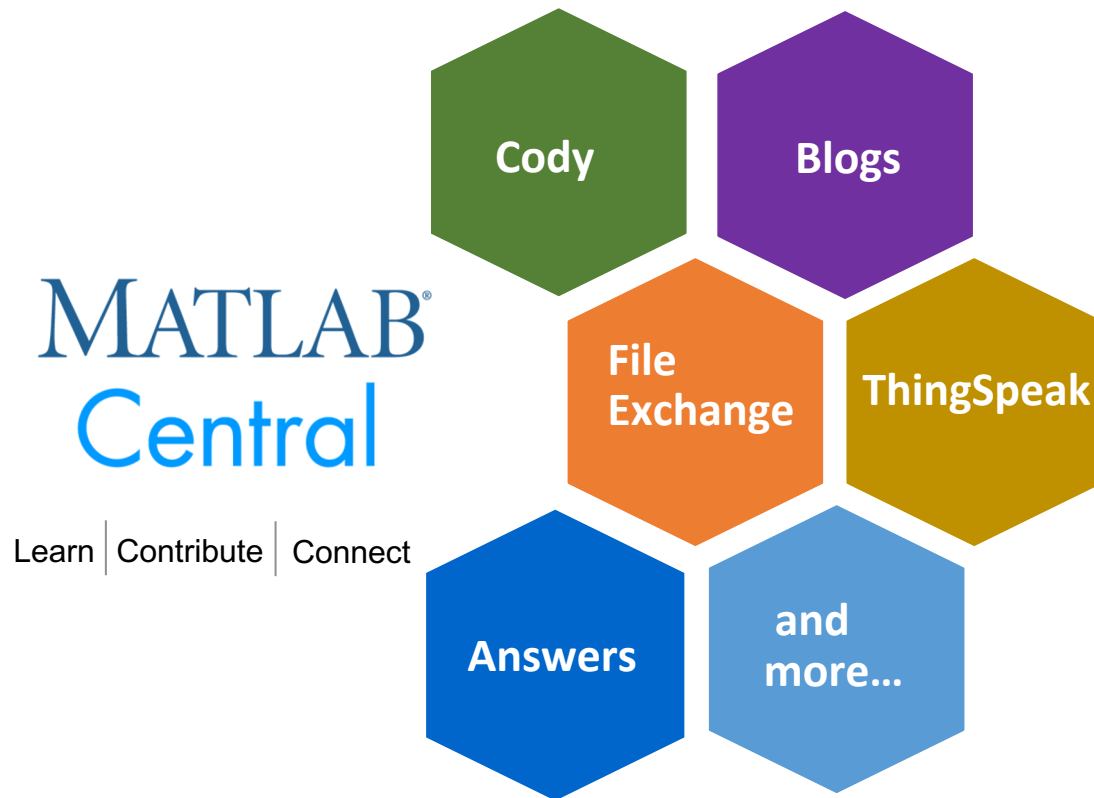


Wind tunnel test setup featuring the Hybrid Wing Body model (inverted), with 97-microphone phased array (top) and microphone tower (left).

*“Our legacy code took up to 40 minutes to analyze a single wind tunnel test. The addition of GPU computing with Parallel Computing Toolbox cut it to under a minute. It took 30 minutes to get our MATLAB algorithm working on the GPU—no low-level CUDA programming was needed.”*  
- Christopher Bahr, research aerospace engineer at NASA

# MATLAB Central Community

Every month, over **2 million** MATLAB & Simulink users visit MATLAB Central to download code, improve programming skills, and get questions answered.



**File Exchange**: Download code from a huge repository of free code including **tens of thousands** of open source community files

**Cody**: Sharpen programming skills while having fun

**Blogs**: Get the inside view from Engineers who build and support MATLAB & Simulink

**ThingSpeak**: Explore IoT Data

**MATLAB Answers**: Q&A forum; most questions get answered in only **60 minutes**

And more for you to explore...

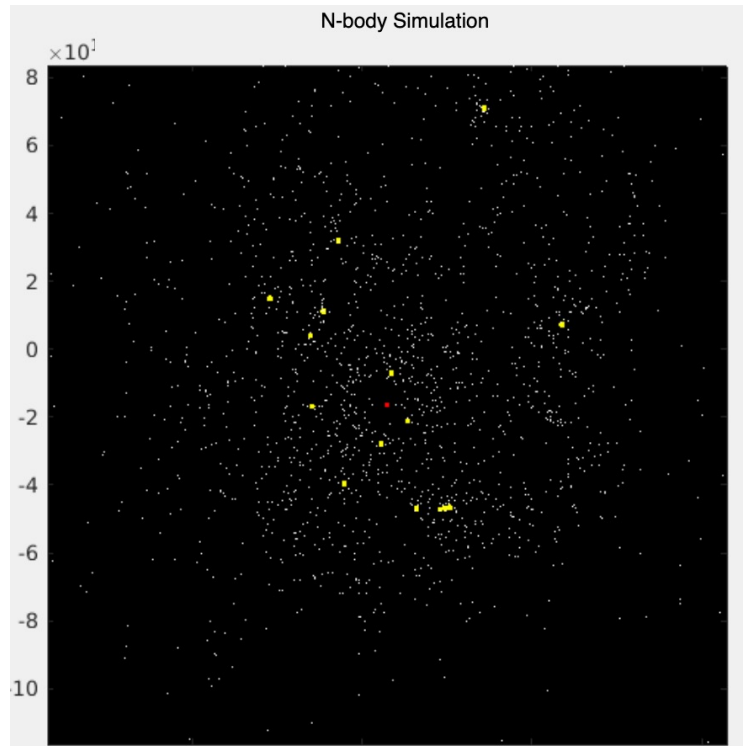


# Agenda

1. Overview: Open OnDemand
2. Why MATLAB with NVIDIA GPUs
3. Demo: Benefits of MATLAB + GPUs
4. Best Practices & OOD examples of MATLAB+GPU Configurations
5. Next Steps and Discussion

## 3 Demos: Performance Benefits and Ease of MATLAB + GPUs

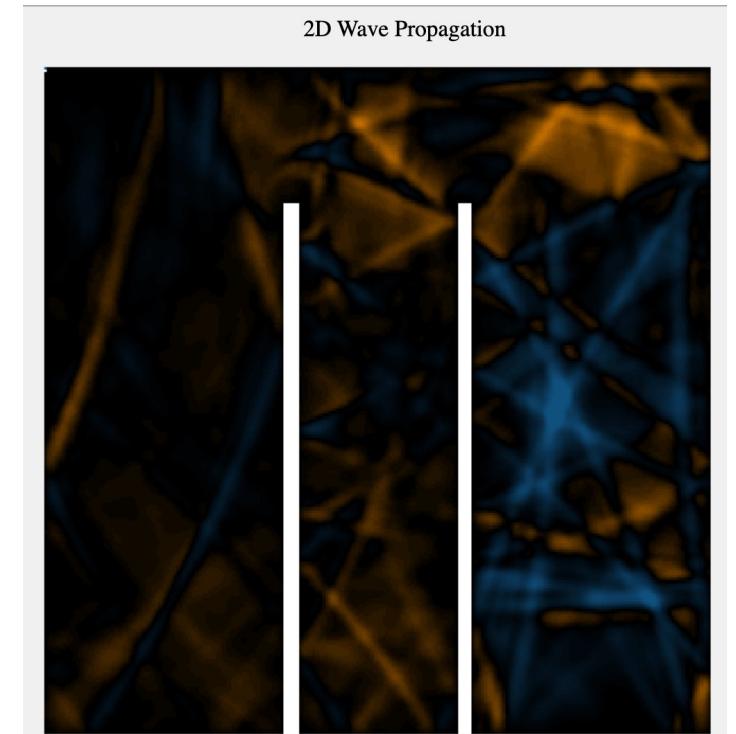
- Use a web browser to access MATLAB with NVIDIA GPUs
- Compare GPU vs CPU performance



N-Body



Particles



2D Wave Propagation

Home / My Interactive Sessions

IDE

Code Server

Eclipse

Nvidia-Nsight-Eclipse

PyCharm

Interactive Apps

AMDuProf

Ansys Workbench

DeepLabCut

Desktop

GAMS Studio

IGV

MATLAB (293017)

1 node | 12 cores | Running

Host: >\_tc-gpu001

Delete

Created at: 2021-10-27 13:11:06 EDT

Time Remaining: 3 hours and 55 minutes

Session ID: 131be365-5e26-40ad-8138-08a4cff10bb6

"Connect to Matlab"

MATLAB (292870)

Completed

Created at: 2021-10-27 12:02:48 EDT

Delete


Session ID: 86d9d359-8252-48f5-94e4-da399d77eb72

For debugging purposes, this card will be retained for 6 more days

MATLAB (292090)

Completed

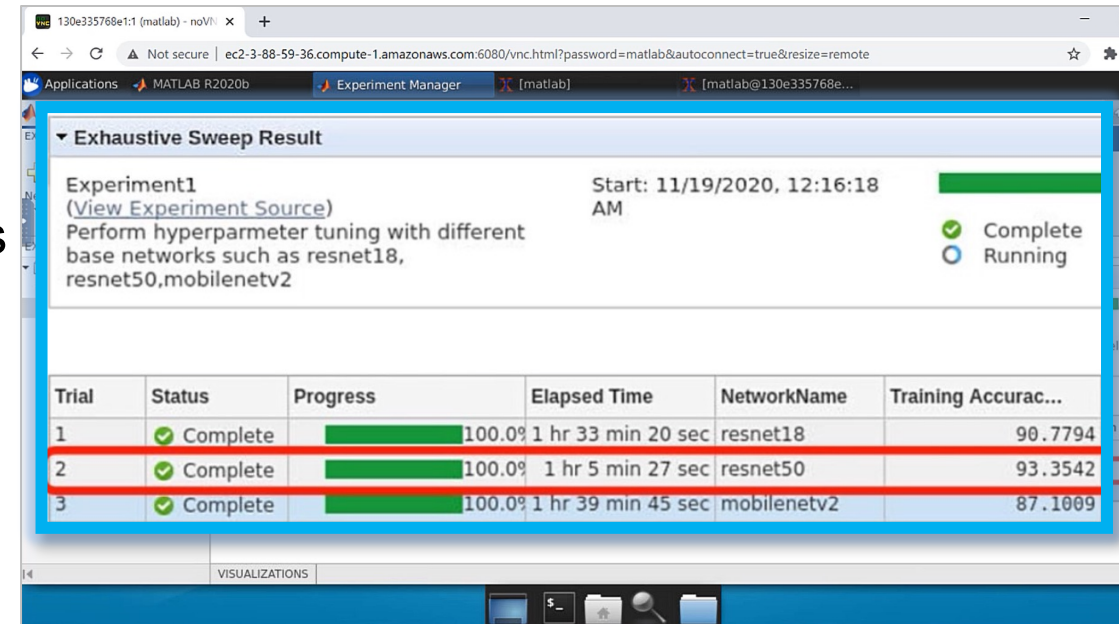
# What does HPC usage look like for Model Training neural networks?



# Use Multiple Nodes to Quickly Find the Optimal Network

## Experiment Manager App

- Manage experiments and reduce manual coding
- Sweep through a range of hyperparameter values
- Compare the results of using different data sets
- Test different neural network architectures



The screenshot shows the Experiment Manager App interface. A window titled 'Exhaustive Sweep Result' is open, displaying details for 'Experiment1'. The experiment description is 'Perform hyperparameter tuning with different base networks such as resnet18, resnet50, mobilenetv2'. The start time is '11/19/2020, 12:16:18 AM'. The status is 'Complete' with a green checkmark icon. Below this, a table lists the results of the sweep.

Trial	Status	Progress	Elapsed Time	NetworkName	Training Accurac...
1	Complete	100.0%	1 hr 33 min 20 sec	resnet18	90.7794
2	Complete	100.0%	1 hr 5 min 27 sec	resnet50	93.3542
3	Complete	100.0%	1 hr 39 min 45 sec	mobilenetv2	87.1009

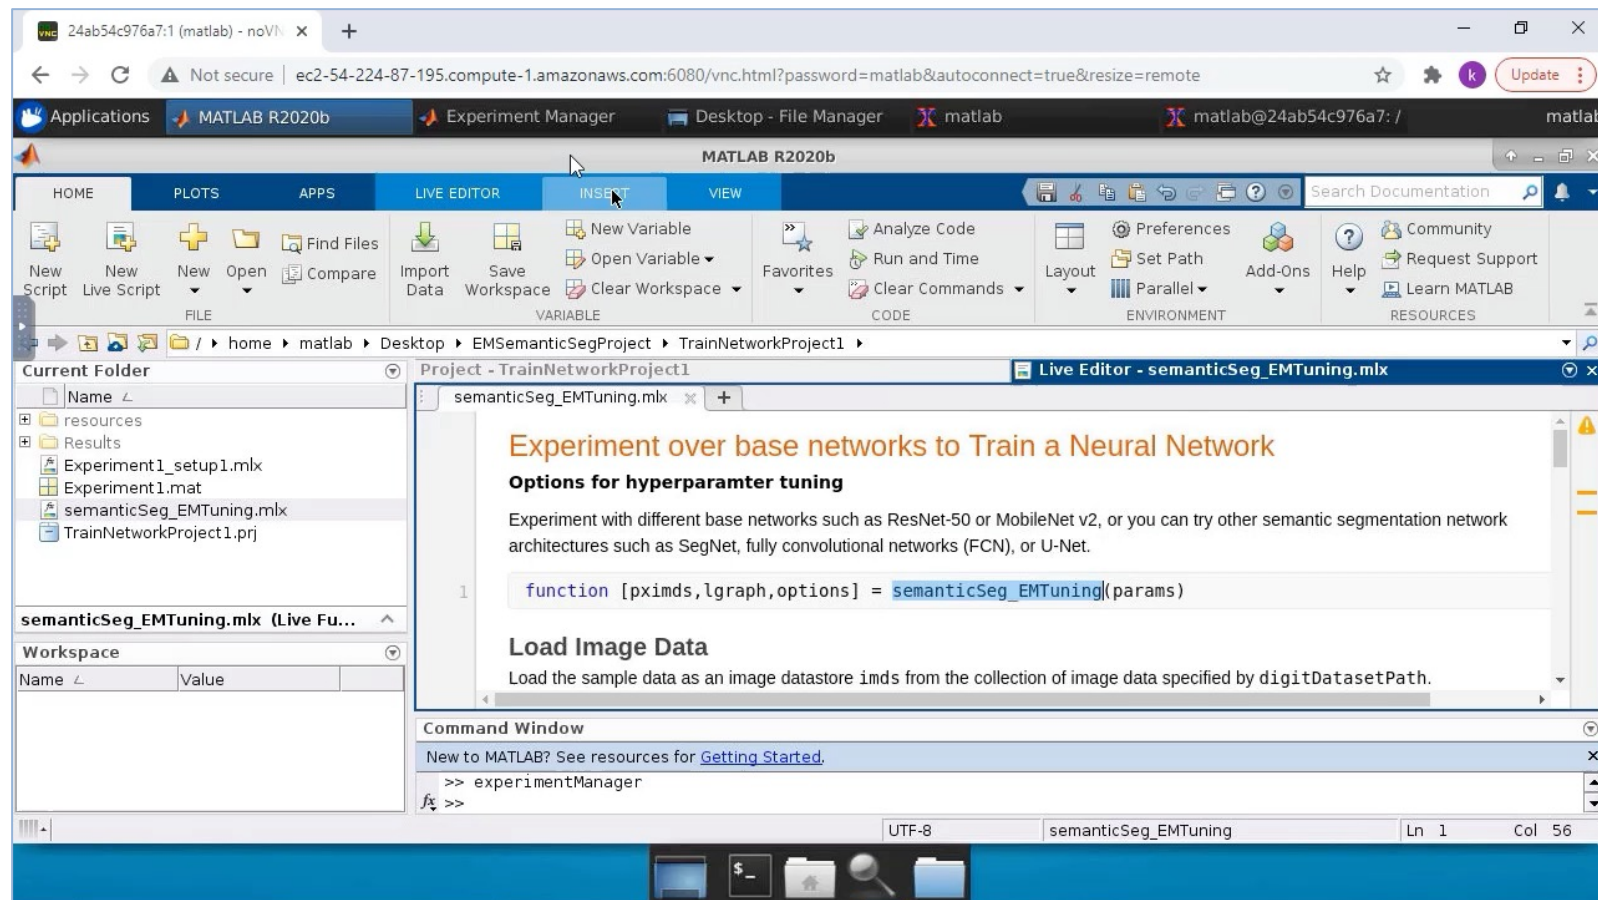
## MATLAB Parallel Server

- Enables multiple nodes to train networks in parallel -> greatly reduce testing time

# Use Multiple Nodes to Quickly Find the Optimal Network

## MATLAB Parallel Server

- Running many experiments to train networks and compare the results in parallel



# Agenda

1. Overview: Open OnDemand
2. Why MATLAB with NVIDIA GPUs
3. Demo: Benefits of MATLAB + GPUs
4. Best Practices & OOD examples of MATLAB+GPU Configurations
5. Next Steps and Discussion



# Example #1 - OOD Configuration for MATLAB & NVIDIA GPUs



- Select MATLAB version
- Select GPU enabled node type

Home / My Interactive Sessions / RHEL7 MATLAB

**Interactive Apps**

- Desktops
  - RHEL7 Interactive Desktop
  - RHEL7 QA Test (SOFTWARE TESTING)
- GUIs
  - RHEL7 ANSYS Workbench
  - RHEL7 AVIZO
  - RHEL7 Abaqus/CAE
  - RHEL7 COMSOL Multiphysics
  - RHEL7 MATLAB**
  - RHEL7 Tecplot
- Servers
  - BYOE Jupyter Server
  - RHEL7 Jupyter Server
  - RStudio Server

**RHEL7 MATLAB**

This app will launch a [MATLAB](#) GUI on [ACI-b](#). You will be able to interact with the MATLAB GUI through a VNC session.

**MATLAB version**

latest version

This defines the version of MATLAB you want to load.

**MATLAB version**

R2021a

**Number of Cores**

4

**Node type**

ACI-i

- **ACI-i** - Use an ACI-i node that has **GPU GL acceleration**.
- total RAM. Available for open account and allocation account submissions.
- **ACI-b Basic Core** - Use an ACI-b node without GPU GL acceleration, 20 available cores, and 128GB total RAM. Available for open account and allocation account submissions.
- **ACI-b Himem Core** - Use an ACI-b node without GPU GL acceleration, 40 available cores, and 1TB total RAM. Available for open account and allocation account submissions.

☐ I would like to receive an email when the session starts

**Launch**

\* All RHEL7 MATLAB session data is generated and stored under the user's home directory in the corresponding [data root directory](#).

powered by  
**OPEN OnDemand**

# Example #2 - OOD Configuration for MATLAB & NVIDIA GPUs



- Select MATLAB version
- Select GPU type and count

The screenshot shows the MATLAB OOD configuration interface. On the left, a sidebar lists various applications and clusters. The 'MATLAB' application is selected. On the right, the configuration details for MATLAB are shown. Two blue boxes with arrows highlight specific configuration steps:

- MATLAB version:** A box highlighting the 'MATLAB version' field, which is set to 'R2021a'. Below the field, it says 'This defines the version of MATLAB you want to load.'
- GPU type:count:** A box highlighting the 'GPU type:count' field, which is set to 'a100:1'. Below the field, a list of GPU types is shown: none, k80, 1080ti, 2080ti, titanv, 3090, p40, v100, and a100.

The configuration details for MATLAB include:

- Classes:** A list of classes including Atmospheric Sciences, ATMOS Synoptic, ATMOS5340, Bioinformatics, MIB2020, Chemical Engineering, CHEN Jupyter, and CHEN2450.
- Interactive Apps:** A list of interactive apps including Desktops, Basic Interactive Desktop, Expert Interactive Desktop, IDEs, ANSYS Electronics Desktop, ANSYS Workbench, Abaqus, COMSOL Multiphysics, IDL, Lumerical DEVICE Suite, MATLAB, Mathematica, RELION, STATA, Servers, Code Server, Jupyter Lab, Jupyter Notebook, R Shiny app, RStudio server, Visualization, and Coot.
- MATLAB:** This app will launch a MATLAB GUI on a HPC cluster or on a Frisco node. You will be able to interact with the MATLAB GUI through a VNC session. GPU specification is optional for the partitions that have them.
- Cluster:** notchpeak
- Number of hours:** 24
- Memory per job:** default
- GPU type:count:** a100:1

# Example #3 - OOD Configuration for MATLAB & NVIDIA GPUs



- Select GPU queue and count

Home / My Interactive Sessions / MATLAB

**IDE**

- Code Server
- Eclipse
- Nvidia-Nsight-Eclipse
- PyCharm

**Interactive Apps**

- AMDuProf
- Ansys Workbench
- DeepLabCut
- Desktop
- GAMS Studio
- IGV
- Jupyter Notebook -- test conda env/module
- Jupyter Notebook --Container
- LS-PrePost
- MATLAB**
- MATLAB via VNC
- PSPP
- ParaView
- Rstudio
- Rstudio -- Rstudio 1.4.1717, R>4.1
- VisIT
- qGIS
- Cascades-DragonsTooth-Huckleberry

**MATLAB**

This app will launch an interactive Matlab session via HTML on one node.

Cluster  
tinkercliffs

Account  
openondemand2

- The allocation you would like to use for SLURM.

Reservation

**Partition**  
a100\_normal\_q

- To request a GPU enabled queue, preface it with v100\_. Example:  
v100\_normal\_q

Number of hours (min=1, max=48)  
1

Number of nodes (min=1, max=2)  
1

Number of cores per node (min=1, max=128)  
12

**Number of GPUs per node (min=0, max=2)**  
1

Required modules

- This should be spaced list of -compatible- modules as you would see in 'module spider' for the node type targeted.
- Click the blue hash and look at output.log and script.sh to troubleshoot module errors.


☒ I would like to receive an email when the session starts

Launch

\* The MATLAB session data for this session can be accessed under the [data root directory](#).

# Example #4 - OOD with XDMoD for MATLAB+GPU Usage

My Interactive Sessions All Apps

 **Ohio Supercomputer Center**  
An OH·TECH Consortium Member

OnDemand provides an integrated, single access point for all of your HPC resources.

**Message of the Day**

2021-07-07 - Please review your usage of /fs/ess

Recent improvements to a routine OSC accounting process resulted in larger adjustments to fill usage than anticipated for some accounts. The updated metrics are more accurate but some accounts have less free space than expected. Please review your usage closely and, if necessary, remove files or request a quota expansion via OSCHelp to avoid exhausting your available space.

2020-03-16 - OSC Classroom Support

CLASSROOM RESOURCES FOR DISTANCE LEARNING

If your class has lost or limited access to computer labs, the Ohio Supercomputer Center might be able to help by providing no-cost access to cloud computing resources. Classes and workloads of any size can gain access. OSC's web-browser interface to its substantial Linux computer systems provides novice users with virtual desktops preloaded with applications, such as MATLAB, RStudio, or Jupyter Notebook.

As an example, an OSU undergrad statistics class recently used iPads to remotely access RStudio on OSC systems. We can provide online demonstrations or evaluations and potentially add additional software packages.

Please contact [OSCHelp@osc.edu](mailto:OSCHelp@osc.edu) to talk to OSC about distance-learning support options available to you.

**Jobs Efficiency Report -** 2021-09-27 to 2021-10-27 [Open XDMoD](#)

17.8% efficient 82.2% inefficient

37 inefficient jobs / 45 total jobs

22.4% efficient 77.6% inefficient

155.2 inefficient core hours / 200.1 total core hours

**Recently Completed Jobs -** 2021-09-27 to 2021-10-27 [Open XDMoD](#)

ID	Name	Date	CPU
5819917	ondemand/sys /dashboard /dev/codeserver	10/3	01.6
5822861	ondemand/dev /dashboard /sys/bc_osc_jupyter	10/3	00.0
15303938	ondemand/dev /dashboard	10/3	00.0

GPU Utilization

# Agenda

1. Overview: Open OnDemand
2. Why MATLAB with NVIDIA GPUs
3. Demo: Benefits of MATLAB + GPUs
4. Best Practices & OOD examples of MATLAB+GPU Configurations
5. Next Steps and Discussion

## Next Steps

1. Enable and improve GPU access with MATLAB on your cluster.
2. Ensure the interface clearly shows GPU access and MATLAB capabilities.
3. Set up the MATLAB App on your site.
  - GITHUB link to Virginia Tech's OOD MATLAB App
  - [https://github.com/AdvancedResearchComputing/OnDemandApps/tree/main/bc\\_vt\\_matlab\\_html](https://github.com/AdvancedResearchComputing/OnDemandApps/tree/main/bc_vt_matlab_html)

Contact us for help with setting up MATLAB on your cluster.

- [nvidia@mathworks.com](mailto:nvidia@mathworks.com)
- **Or** put your contact info in the chat

# Technical Questions and Content Requests

- For technical questions and support go to the [Open OnDemand Discourse](https://discourse.osc.edu/c/open-ondemand/5)
  - <https://discourse.osc.edu/c/open-ondemand/5>
- What would you like to see next?
  - Contact us: [nvidia@mathworks.com](mailto:nvidia@mathworks.com)

The screenshot shows the Open OnDemand Discourse forum interface. At the top, the header includes the Ohio Supercomputer Center logo and navigation links. Below the header, there are filters for 'Open OnDemand', 'all', 'all tags', and 'all', along with a 'Latest' button and counts for 'New (1)', 'Unread (1)', and 'Top'. The main content area is divided into three columns: 'Announcements', 'Get Help', and 'Feature Requests and Roadmap Discussion'. The 'Announcements' column lists 'Open OnDemand sessions at Gateways'21 and SC'21 Conferences', 'Open OnDemand Monthly Open Office Hours', and 'Open OnDemand tips and tricks calls'. The 'Get Help' column lists 'New OOD install: shows me the Apache welcome screen. delete that conf, shows...', 'Timed out waiting for RStudio Server to open port', and 'Unable to launch fully containerized Rstudio'. The 'Feature Requests and Roadmap Discussion' column lists 'Trello board roadmap', 'What If we combined Job Templates and Interactive App Forms?', and 'Customize form field layout'. Below these columns, there is a 'Collaborator Discussions' section. The bottom section shows a list of topics with columns for 'Topic', 'Replies', 'Views', and 'Activity'. The topics listed are 'About the Open OnDemand category', 'What If we combined Job Templates and Interactive App Forms?', 'New OOD install: shows me the Apache welcome screen. delete that conf, shows the autoindex. delete that, just get "/.', and 'Timed out waiting for RStudio Server to open port'.

Topic	Replies	Views	Activity
About the Open OnDemand category	0	602	Sep '19
What If we combined Job Templates and Interactive App Forms?	13	144	1h
New OOD install: shows me the Apache welcome screen. delete that conf, shows the autoindex. delete that, just get "/.	9	45	2h
Timed out waiting for RStudio Server to open port	2	39	6h



# Thank you!

## Q&A / Discussion

Contact us: [nvidia@mathworks.com](mailto:nvidia@mathworks.com)