Fundamentals of Deep Learning

Raymond Ptucha,
Rochester Institute of Technology
NVIDIA, Deep Learning Institute

Electronic Imaging 2020: SC20
January 28, 2019, 8:30am-12:45pm

IS&T International Symposium on
Electronic Imaging 2020
SCIENCE AND TECHNOLOGY

26–30 January 2020
Burlingame, California USA

1

Fair Use Agreement

This agreement covers the use of all slides in this document, please read carefully.

- You may freely use these slides, if:
  - You send me an email telling me the conference/venue/company name in advance, and which slides you wish to use.
  - You receive a positive confirmation email back from me.
  - My name (R. Ptucha) appears on each slide you use.

2

Agenda

- Part I- Intuition and Theory
  - 8:35-9:15pm: Introduction
  - 9:15-10:00pm: Convolutional Neural Networks
  - 10:00-10:40pm: Recurrent Neural Networks
- 10:40-11:00pm: Break
- Part II- Hands on
  - 11:00am-12:45pm: Hands-on exercises

4

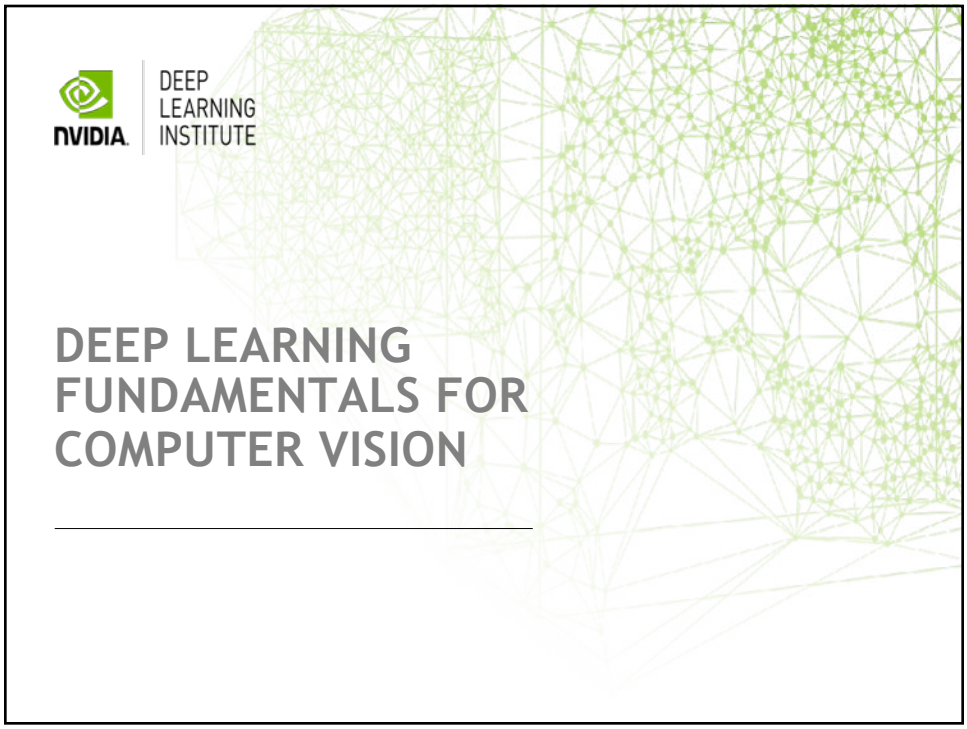

5

This Course: Deep Learning Fundamentals

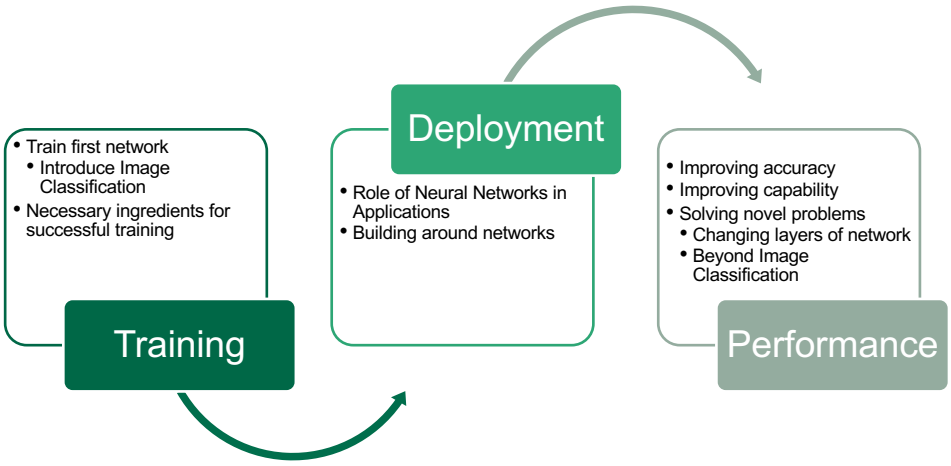

6

Work through Introduction Section

- Navigate to:
- courses.nvidia.com/dli-event
- Browser Recommendation: Chrome

- Event code: EIC_CV2.0_ AMBASSADOR_JA20
- Create an Account
- Work through the Introduction Section and 'Start' launching your first GPU task

7

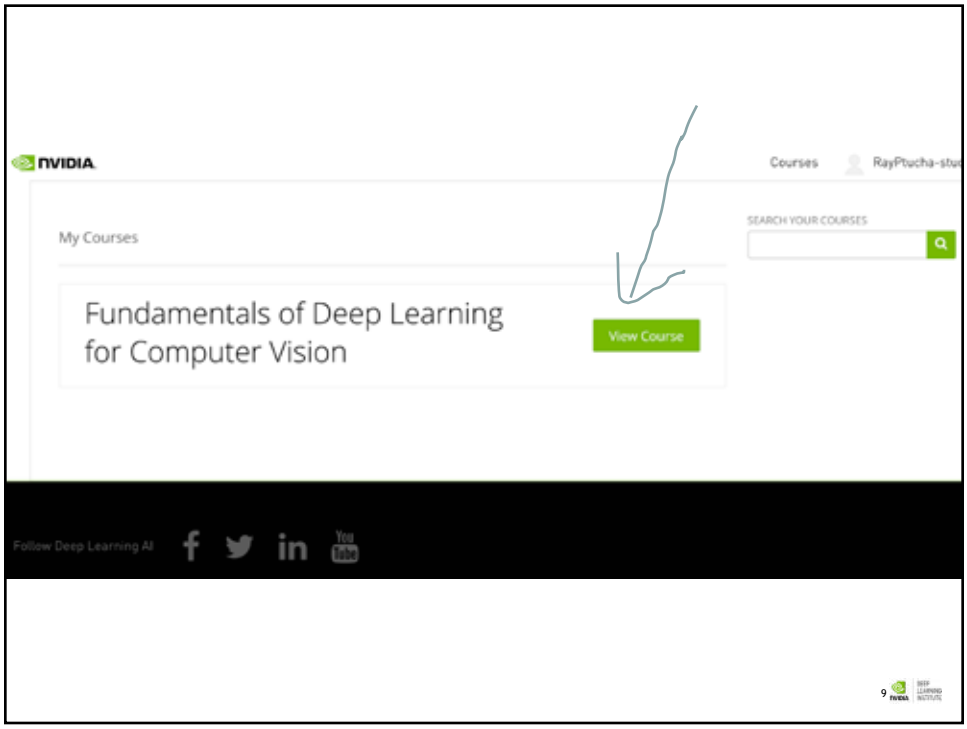
NVIDIA
Courses
My Courses
SEARCH YOUR COURSES
Fundamentals of Deep Learning for Computer Vision
View Course
Follow Deep Learning AI
9

9

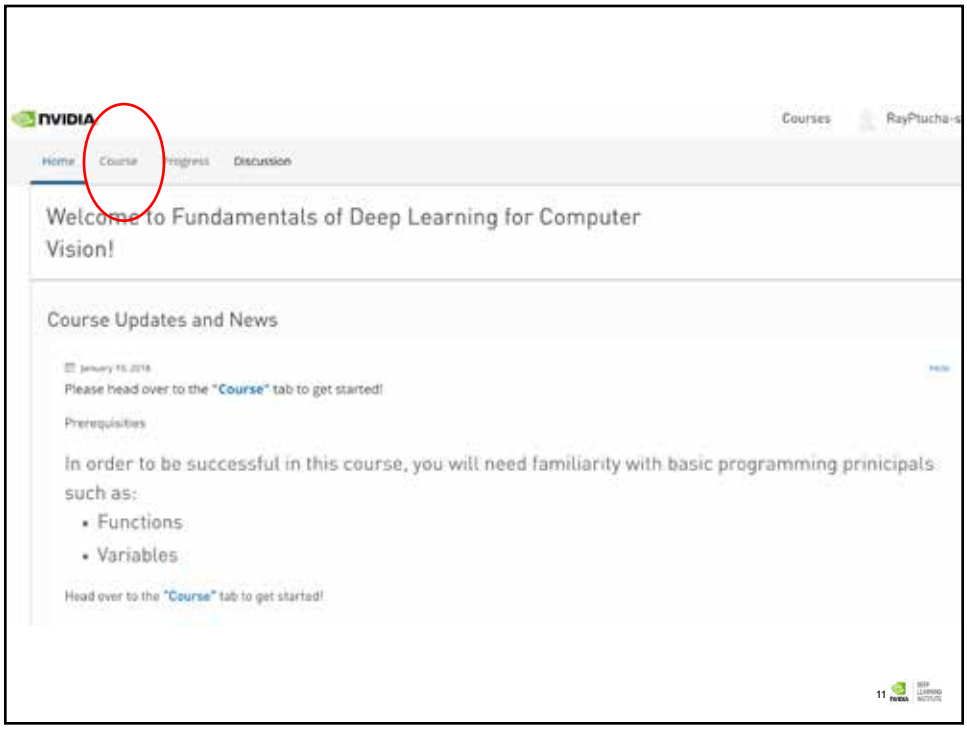
NVIDIA
Courses
Home
Course
Progress
Discussion
Welcome to Fundamentals of Deep Learning for Computer Vision!
Course Updates and News
Please head over to the "Course" tab to get started!
Prerequisites
In order to be successful in this course, you will need familiarity with basic programming prinicipals such as:
• Functions
• Variables
Head over to the "Course" tab to get started!
11

11

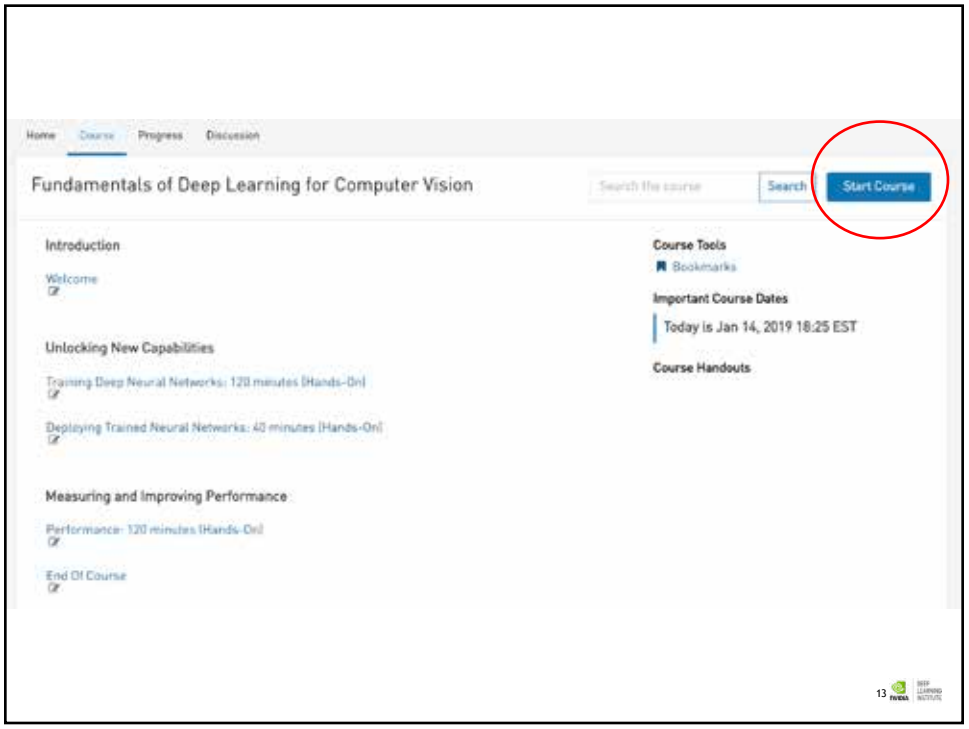
Home
Course
Progress
Discussion
Fundamentals of Deep Learning for Computer Vision
Search
Start Course
Introduction
Welcome
Unlocking New Capabilities
Training Deep Neural Networks: 120 minutes (Hands-On)
Deploying Trained Neural Networks: 40 minutes (Hands-On)
Measuring and Improving Performance
Performance: 120 minutes (Hands-On)
End Of Course
Course Tools
Bookmarks
Important Course Dates
Today is Jan 14, 2019 18:25 EST
Course Handouts
13

13

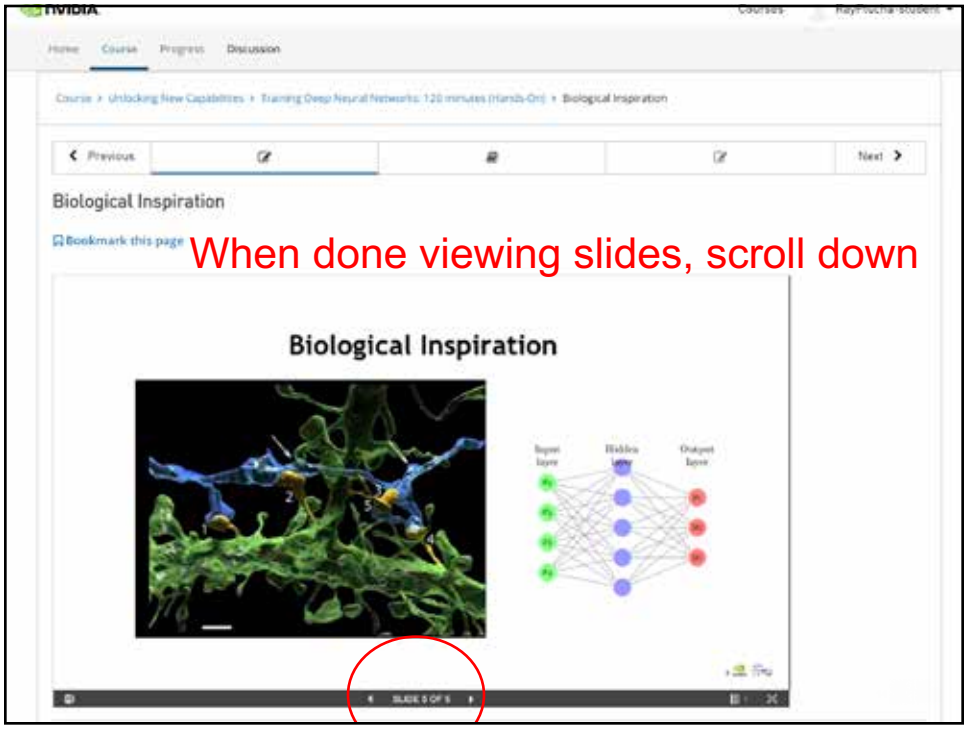
Home
Course
Progress
Discussion
Previous
Next
Biological Inspiration
Bookmark this page
When done viewing slides, scroll down
Biological Inspiration

14

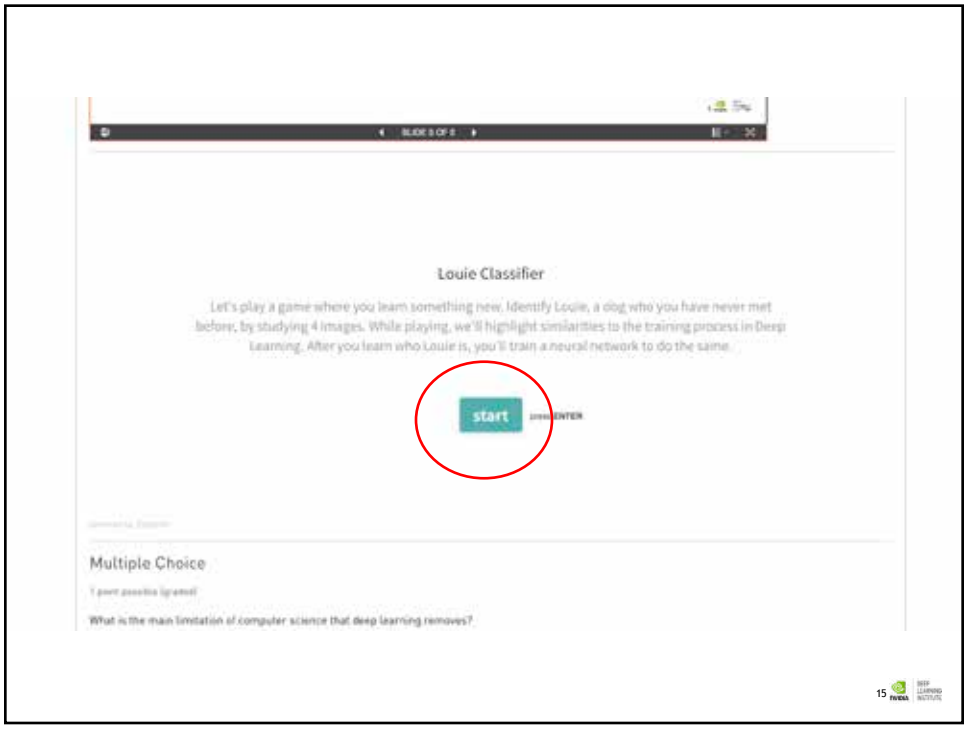
Louie Classifier
Let's play a game where you learn something new. Identify Louie, a dog who you have never met before, by studying 4 images. While playing, we'll highlight similarities to the training process in Deep Learning. After you learn who Louie is, you'll train a neural network to do the same.
start
Multiple Choice
What is the main limitation of computer science that deep learning removes?
15

15

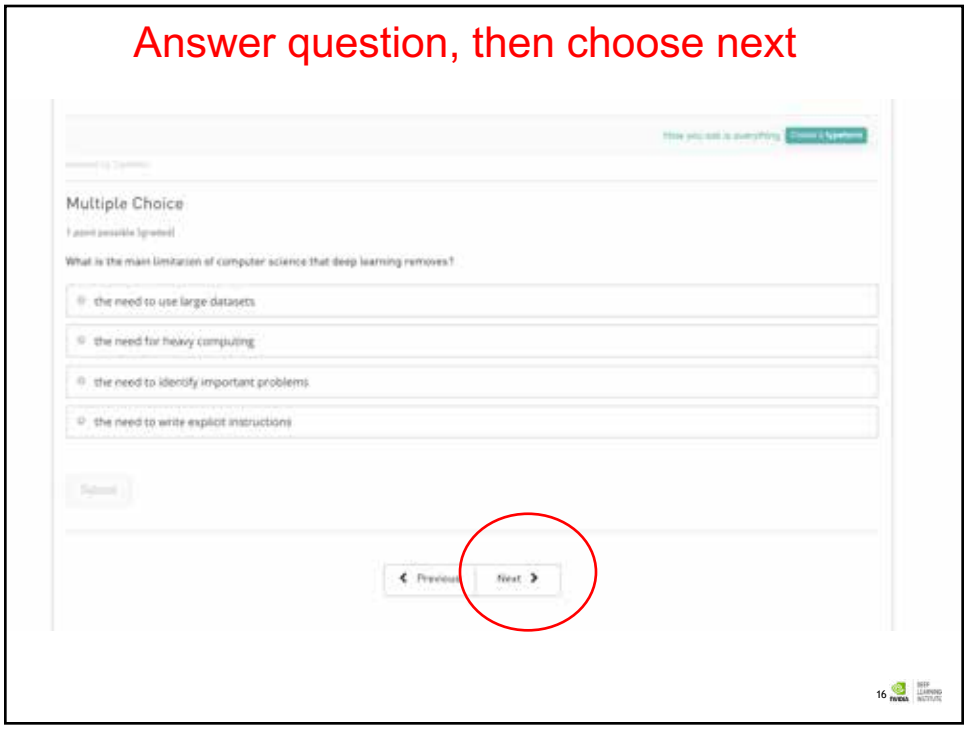
Answer question, then choose next
Multiple Choice
What is the main limitation of computer science that deep learning removes?
the need to use large datasets
the need for heavy computing
the need to identify important problems
the need to write explicit instructions
Previous
Next
16

16

When done viewing slides, scroll down

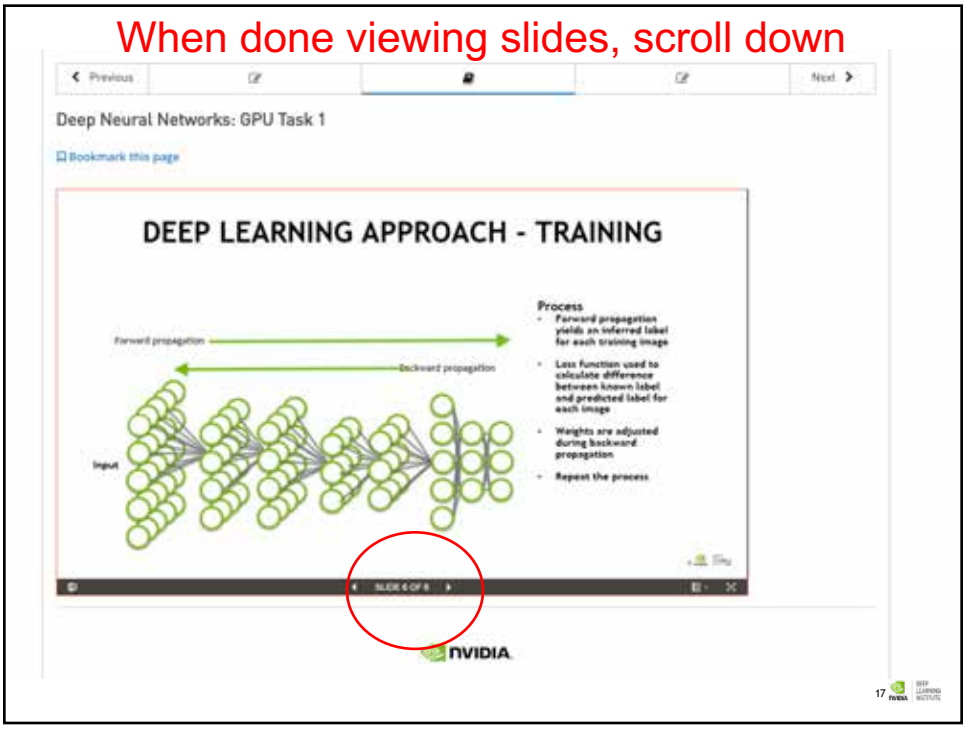

17

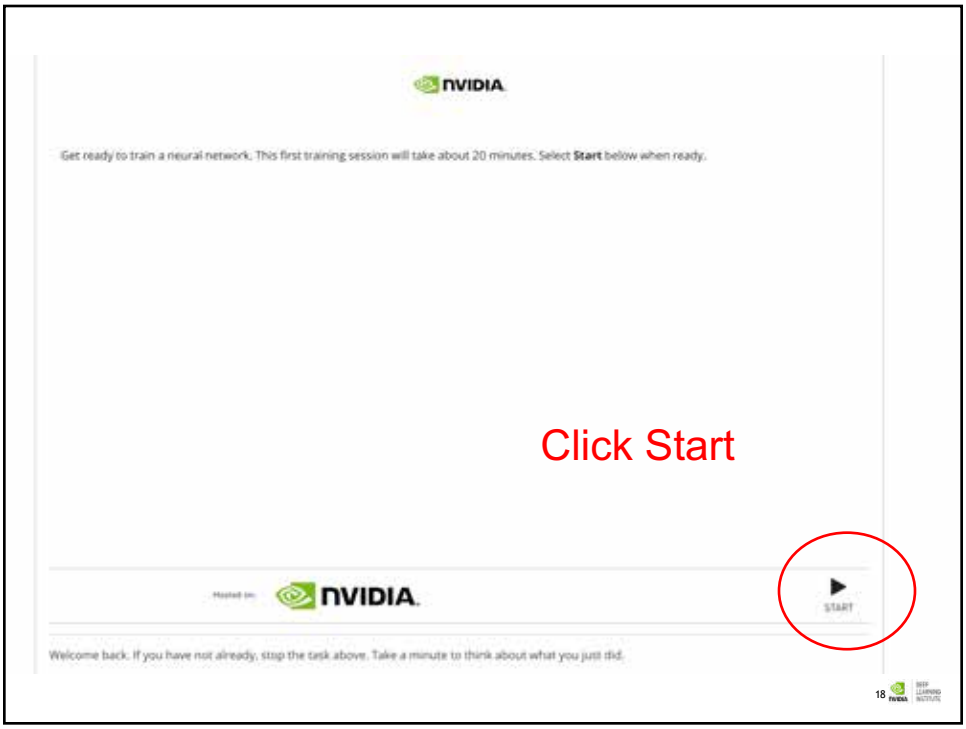

18

It will take a few minutes
While it is loading, read intro material…

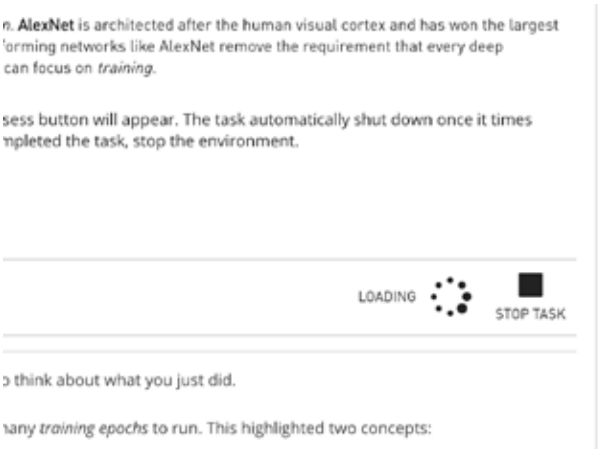

19

Click Launch Task when ready

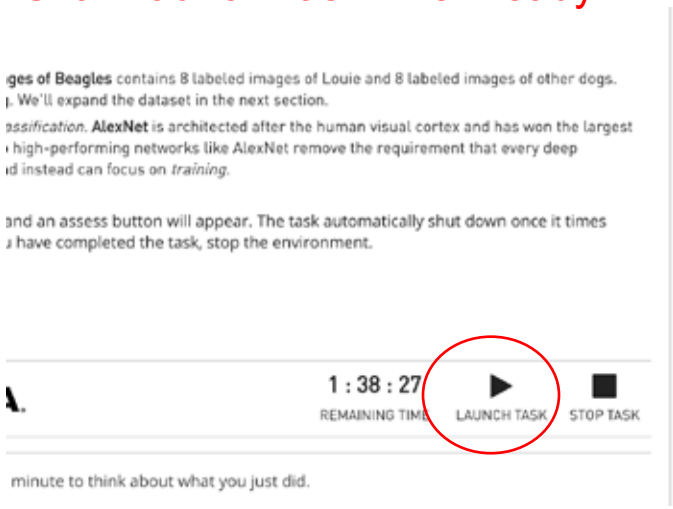

20

This will open up a new tab.
You can go back and forth between original and this new tab

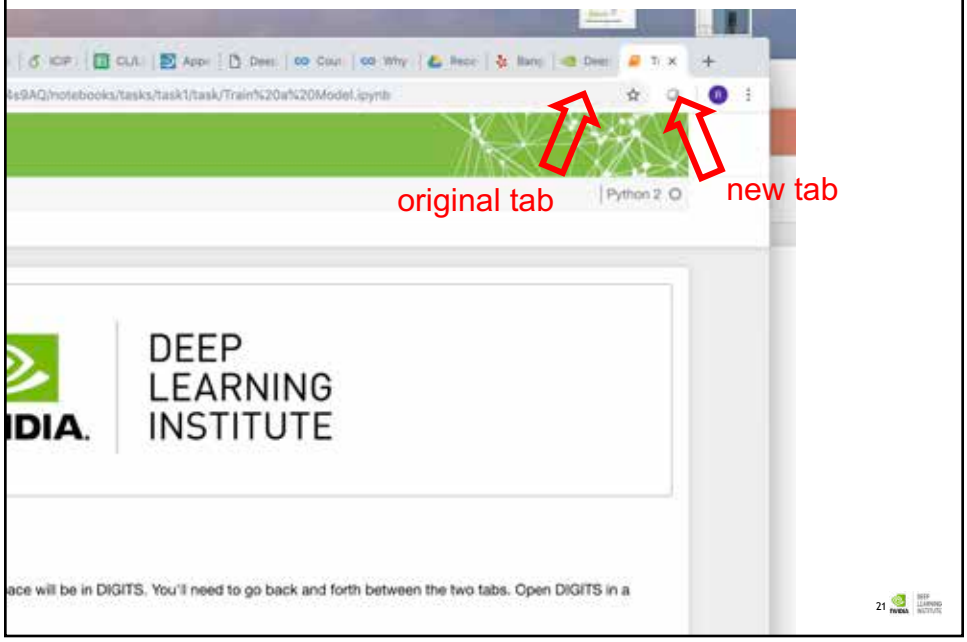

21

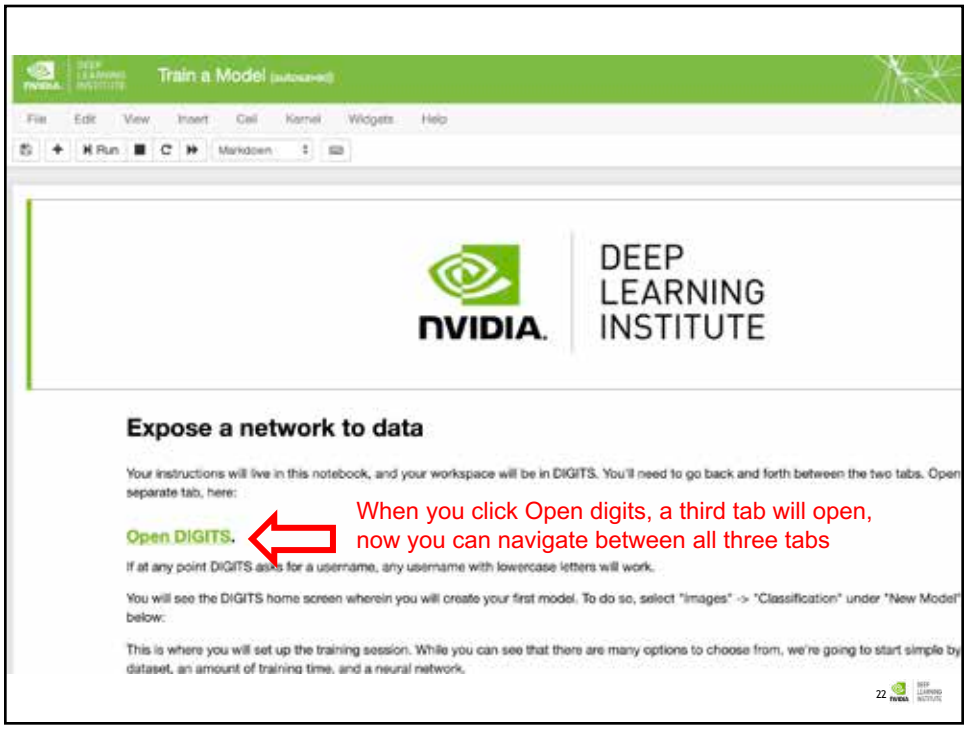

22

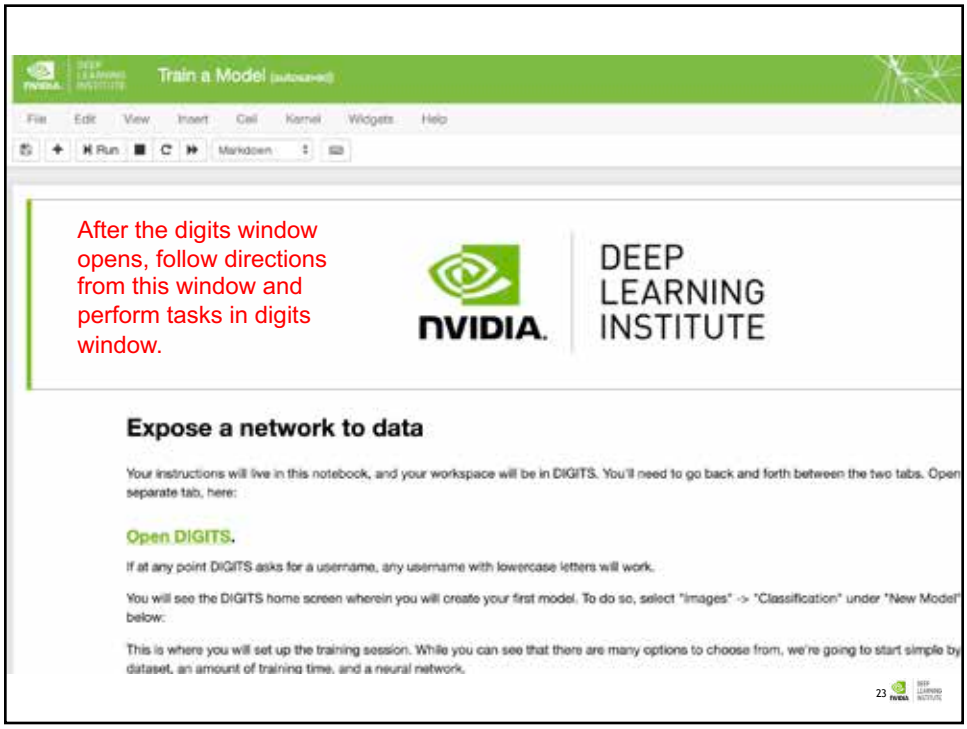
Train a Model
File Edit View Insert Cell Kernel Widgets Help
Run
After the digits window opens, follow directions from this window and perform tasks in digits window.
NVIDIA
DEEP LEARNING INSTITUTE
Expose a network to data
Your instructions will live in this notebook, and your workspace will be in DIGITS. You'll need to go back and forth between the two tabs. Open
separate tab, here:
Open DIGITS.
If at any point DIGITS asks for a username, any username with lowercase letters will work.
You will see the DIGITS home screen wherein you will create your first model. To do so, select "Images" -> "Classification" under "New Model"
below:
This is where you will set up the training session. While you can see that there are many options to choose from, we're going to start simple by
dataset, an amount of training time, and a neural network.
23

23

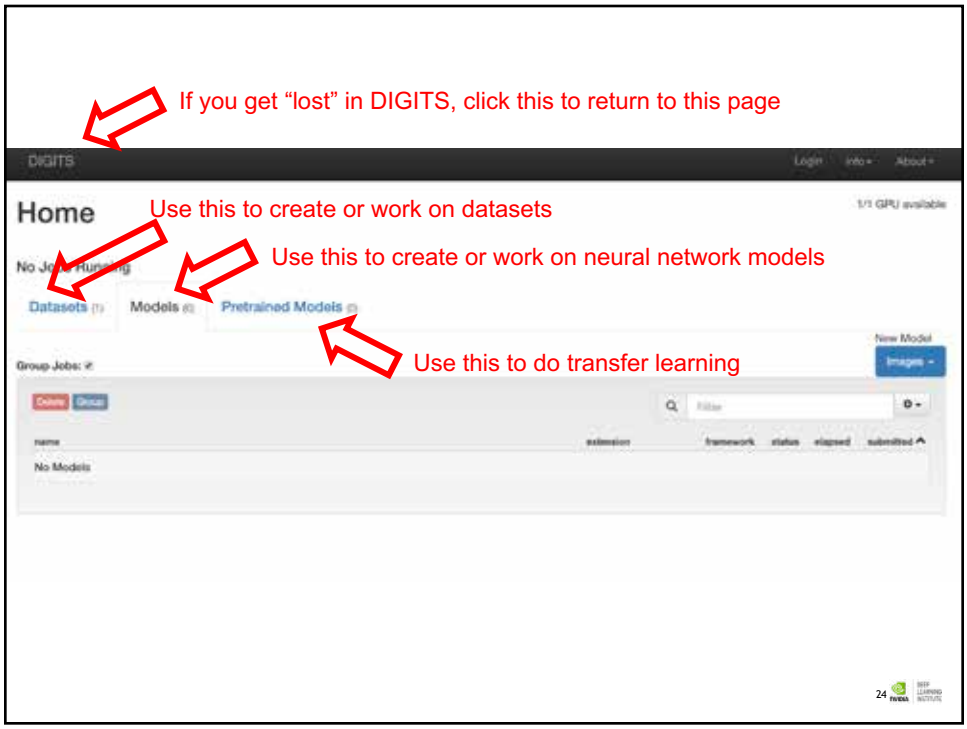
If you get "lost" in DIGITS, click this to return to this page
DIGITS
Login
Info
About
Home
Use this to create or work on datasets
1/1 GPU available
No Jobs Running
Use this to create or work on neural network models
Datasets
Models
Pretrained Models
Use this to do transfer learning
New Model
Images
Group Jobs:
Delete
Group
name
extension
framework
status
elapsed
submitted
No Models
24

24

When you get to the end of the instructions, you are done with that module.
You can then close the digitis tab and this module tab
Do NOT close the original Deep Neural Networks tab
Return to the original Deep Neural Networks tab

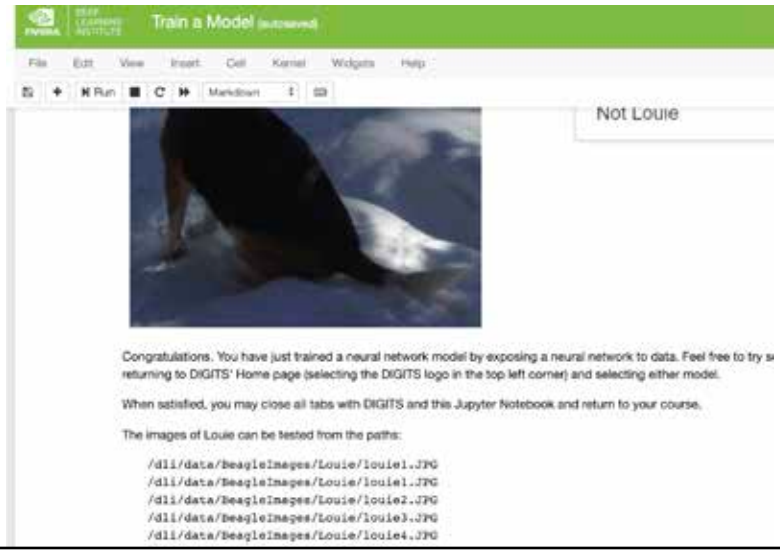

25

When you get to the original Deep Neural Networks tab, click Stop task, and continue to end of page,.

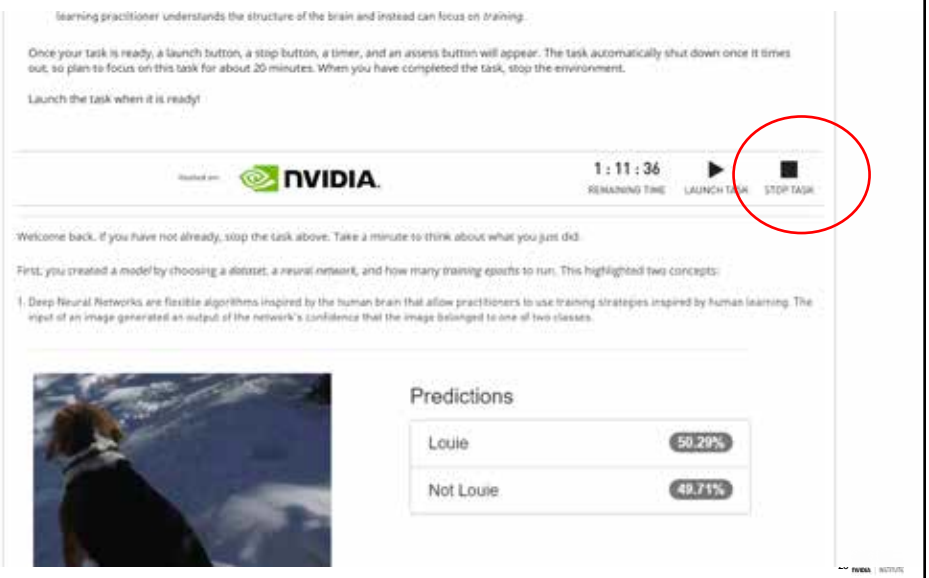

26

When you get to the end of page, click Next and continue to the next module

Predictions

Louie 100.0%

Not Louie 0.0%

Something clearly changed between the first *epoch* and the 100th. Let's take a look at what is actually changing as a neural network trains and intr
our next challenge.

Previous Next

27

For Each Module

- Step through the slides
- Scroll down
- Click “start” to launch a GPU session
- When ready, click “launch”
- This opens a new window for a new module…you go back and forth between tabs
- This new module will have you open DIGITS, creating a third window…you can go back and forth between all three tabs
- When done with module, hit “stop task”
- And continue to bottom of page
- Close module and Digits windows, then click Next…

28

NVIDIA DEEP LEARNING INSTITUTE

For More Information: http://www.rit.edu/mil

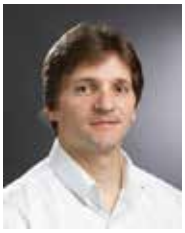

Raymond W. Ptucha
Assistant Professor, Computer Engineering
Director, Machine Intelligence Laboratory
Rochester Institute of Technology

Email: rwpeec@rit.edu
Phone: +1 (585) 797-5561
Office: GLE (09) 3441