Mondays with MATLAB

Hands on workshop

Deep Learning on Alvis - the AI Cluster

Emelie Andersson, *Application Engineer*, **MathWorks**
Sagar Zade, *Customer Success Engineer*, **MathWorks**
Rohit Agrawal, *Customer Success Engineer*, **MathWorks**
Anders Sjöström, *National Co-ordinator and Application Expert at SNIC*, **Lund University**

How can you participate?

You have three options:

1. Run exercises yourself on Alvis (you need to have SNIC and Alvis credentials)

2. Run exercises on your own computer (if you don't have access to Alvis, but access to MATLAB)

   Download exercise material: https://tinyurl.com/4je3v6vm

   Username: AppleHill
   Password: fjr5n8nf

3. Get your popcorn and just follow along

What's Alvis?

"The Alvis cluster is a national SNIC resource dedicated for **Artificial Intelligence** and **Machine Learning** research. The system is built around Graphical Processing Units (GPUs) accelerator cards and consists of several types of compute nodes with multiple Nvidia GPUs."

Set-up Instructions for Alvis

- Login to Chalmers VPN

- Login to Alvis using Thinlinc
  - alvis1.c3se.chalmers.se
  - Alvis specific username and password

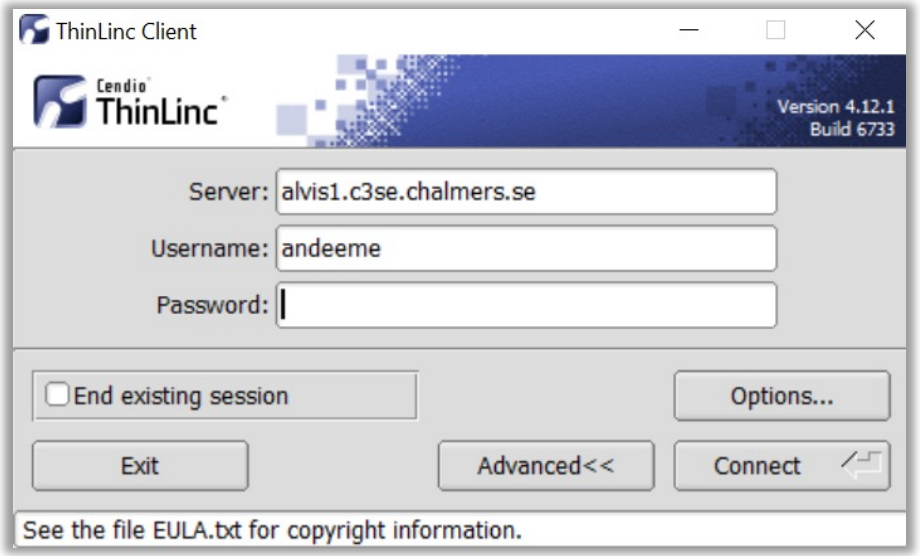

- Open terminal window
  - Copy workshop material:
  - `cp -r /cephyr/users/andeeme/ExerciseFilesDLWorkshop .`

  - Get on a compute node:
  - `interactive --reservation=matlab-monday -n8 -t 03:00:00`

- Start MATLAB (from terminal)
  - `module load MATLAB/R2021a`
  - `srun -A SNIC2021-7-61 -t 03:00:00 -gpus-per-node=T4:1 --x11 -pty matlab`

Agenda

Introduction

 Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Break

 Exercise 2 & 3: Classify blood smear images

Break

 Exercise 4: Improving Network Accuracy

Conclusion

Agenda

Introduction
 Exercise 1: Deep learning in 6 lines of code
Deep Learning Fundamentals
Break
 Exercise 2 & 3: Classify blood smear images
Break
 Exercise 4: Improving Network Accuracy
Conclusion

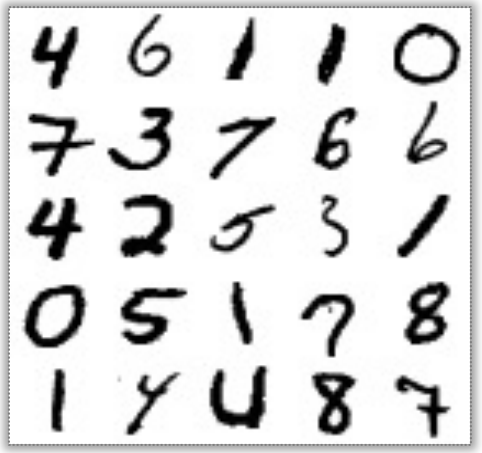

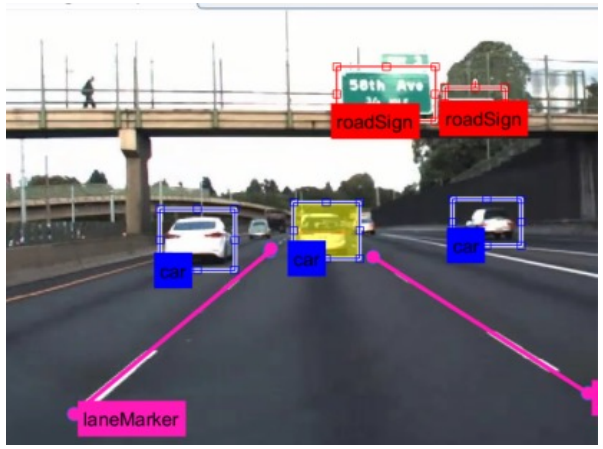

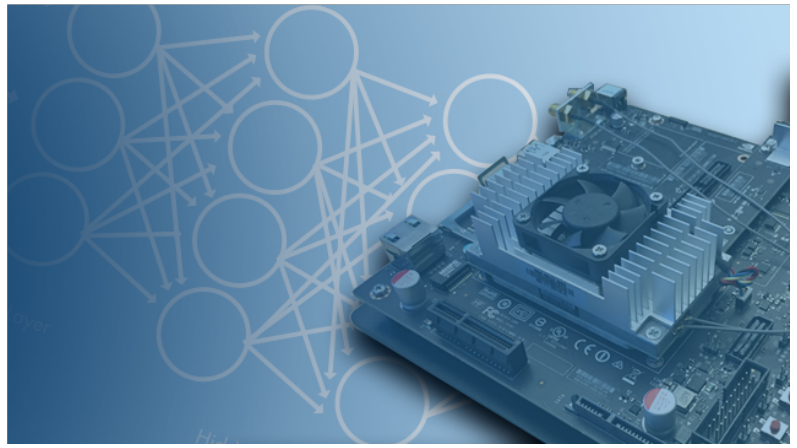

Deep Learning?

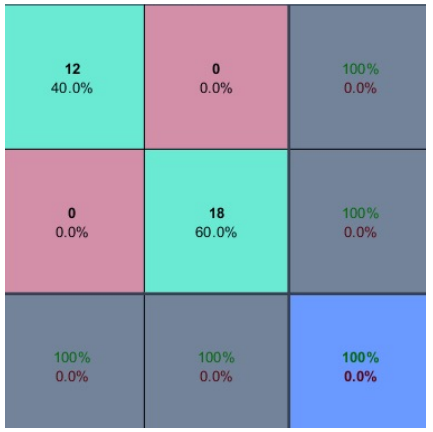

Some Basic Definitions

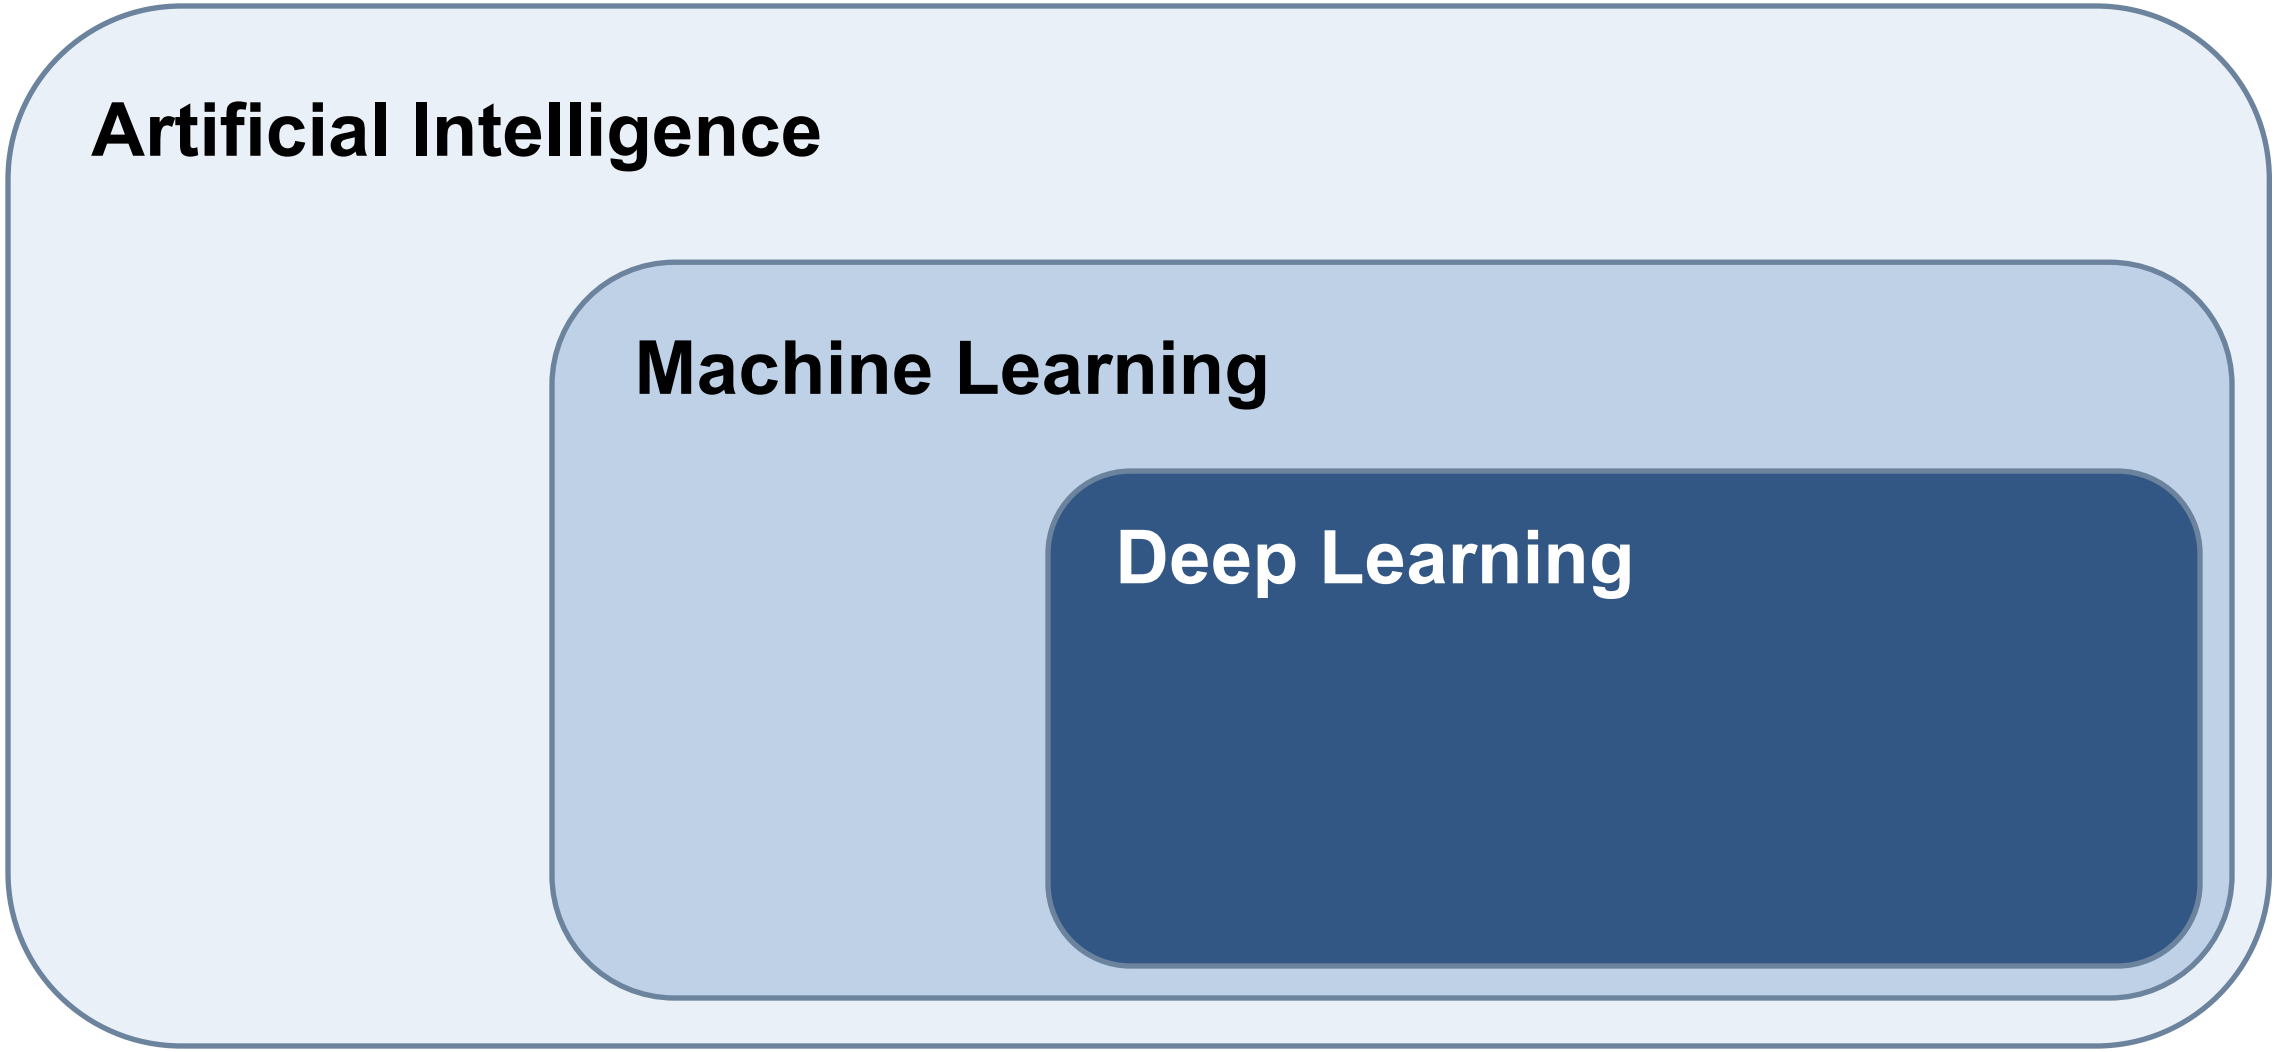

Some Basic Definitions

Artificial Intelligence

The ability of

Machine Learning

The practice of **learning a task from data** **without** relying on a predetermined equation or model

Some Basic Definitions

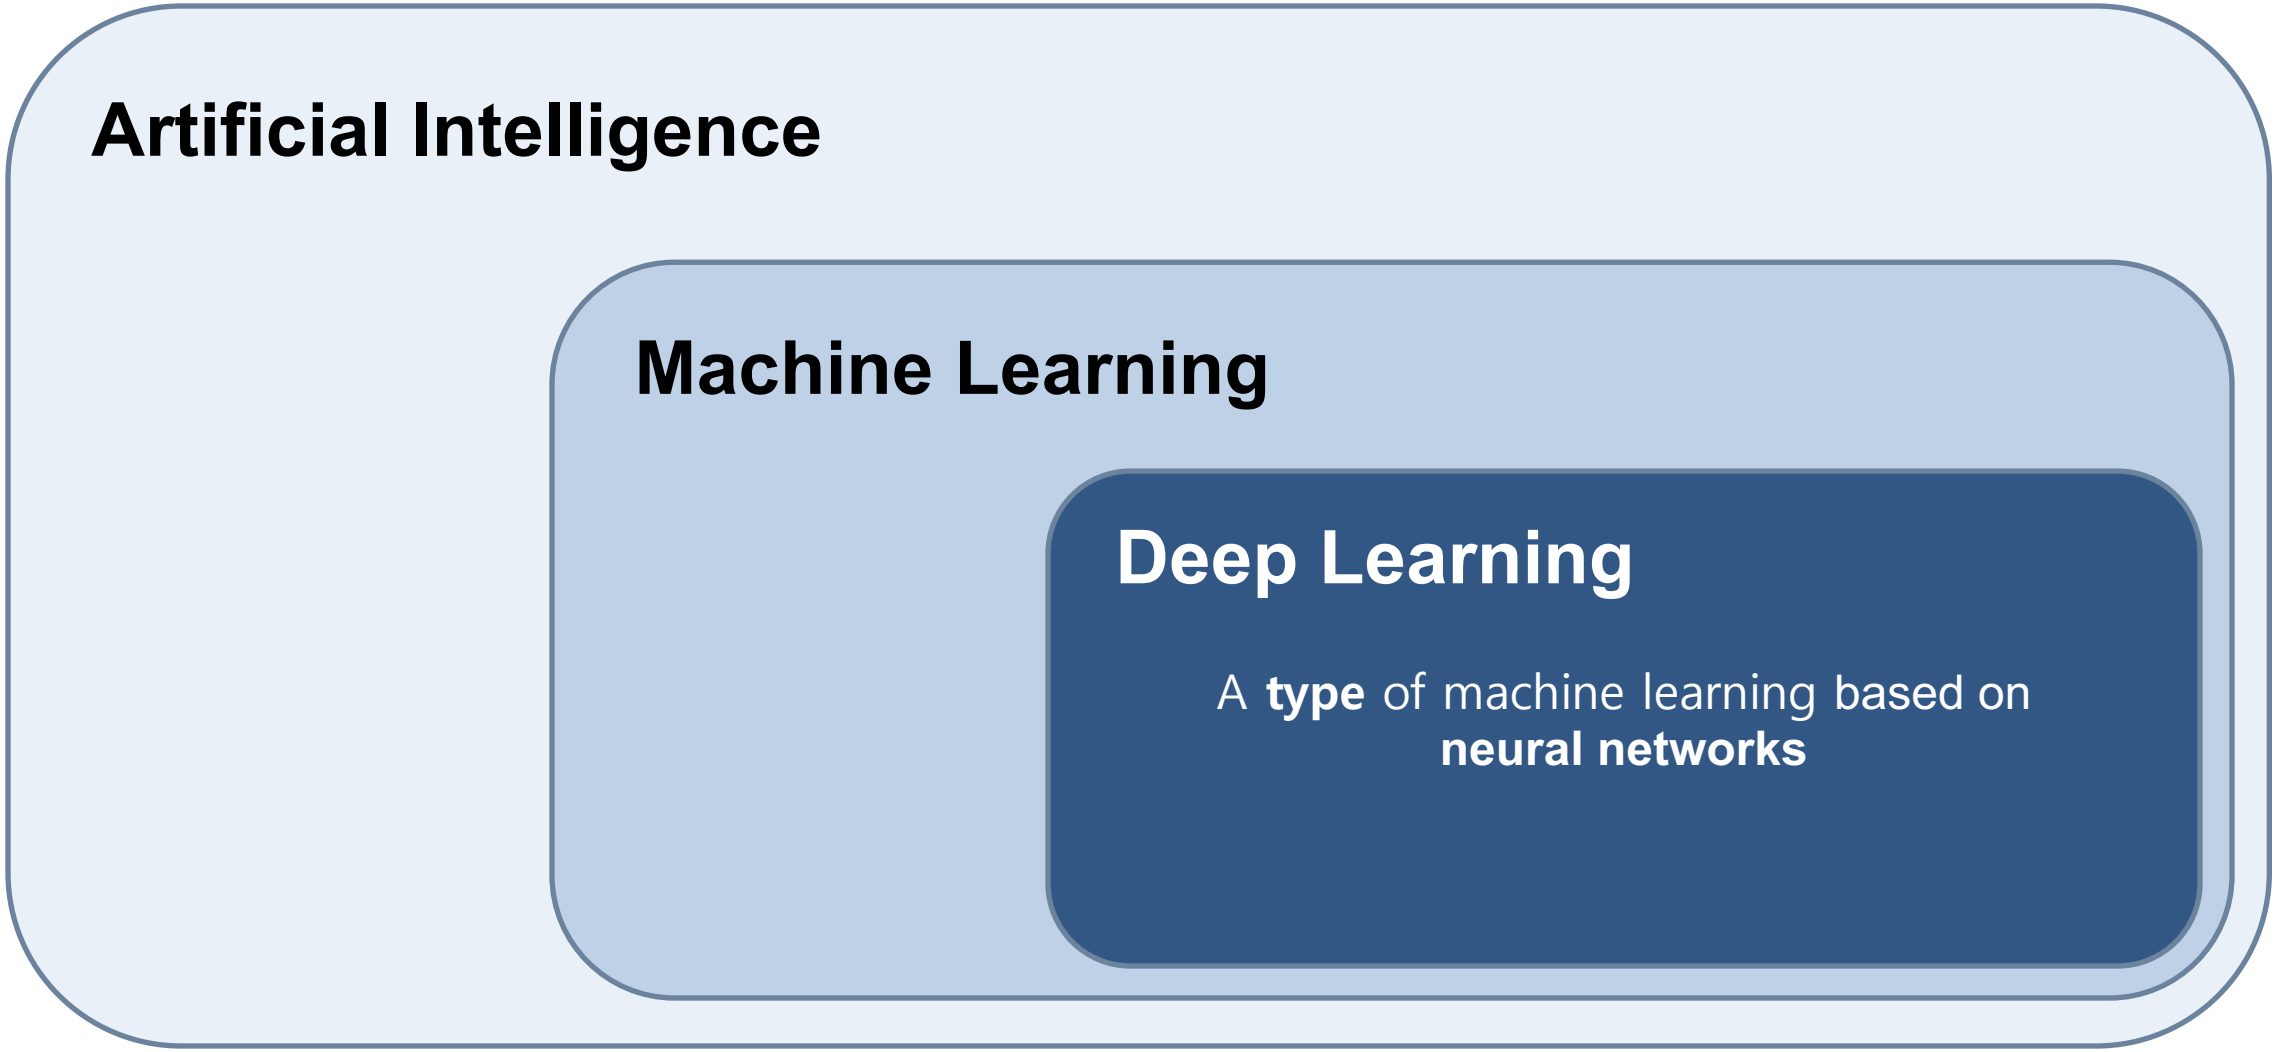

Machine Learning & Deep Learning

- Ability to learn from data inherently without being explicitly programmed
  - Learns complex non-linear relationship
  - More Data = better model

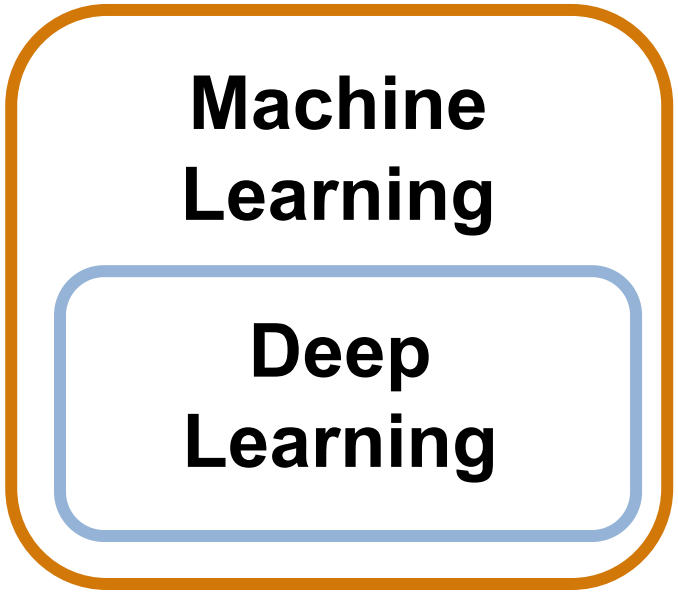

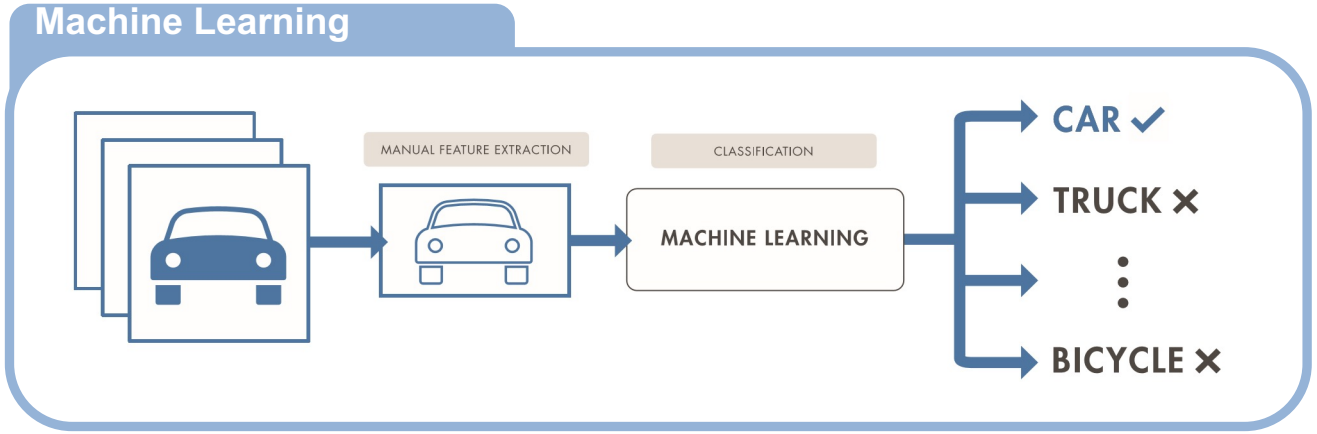

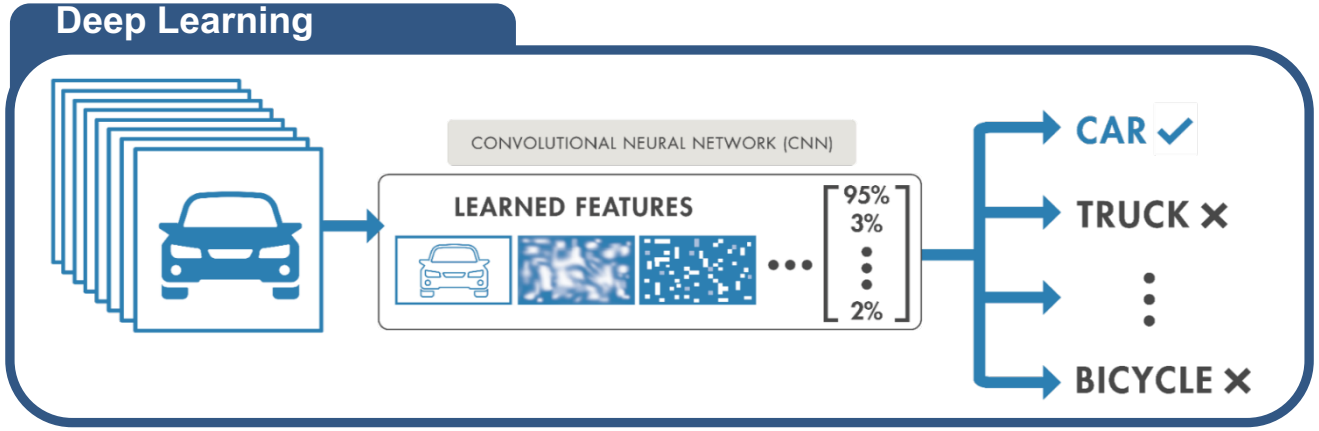

Neural Network Structure

- Deep neural networks have many layers
- Data is passed through the network, and the layer parameters are updated (training)

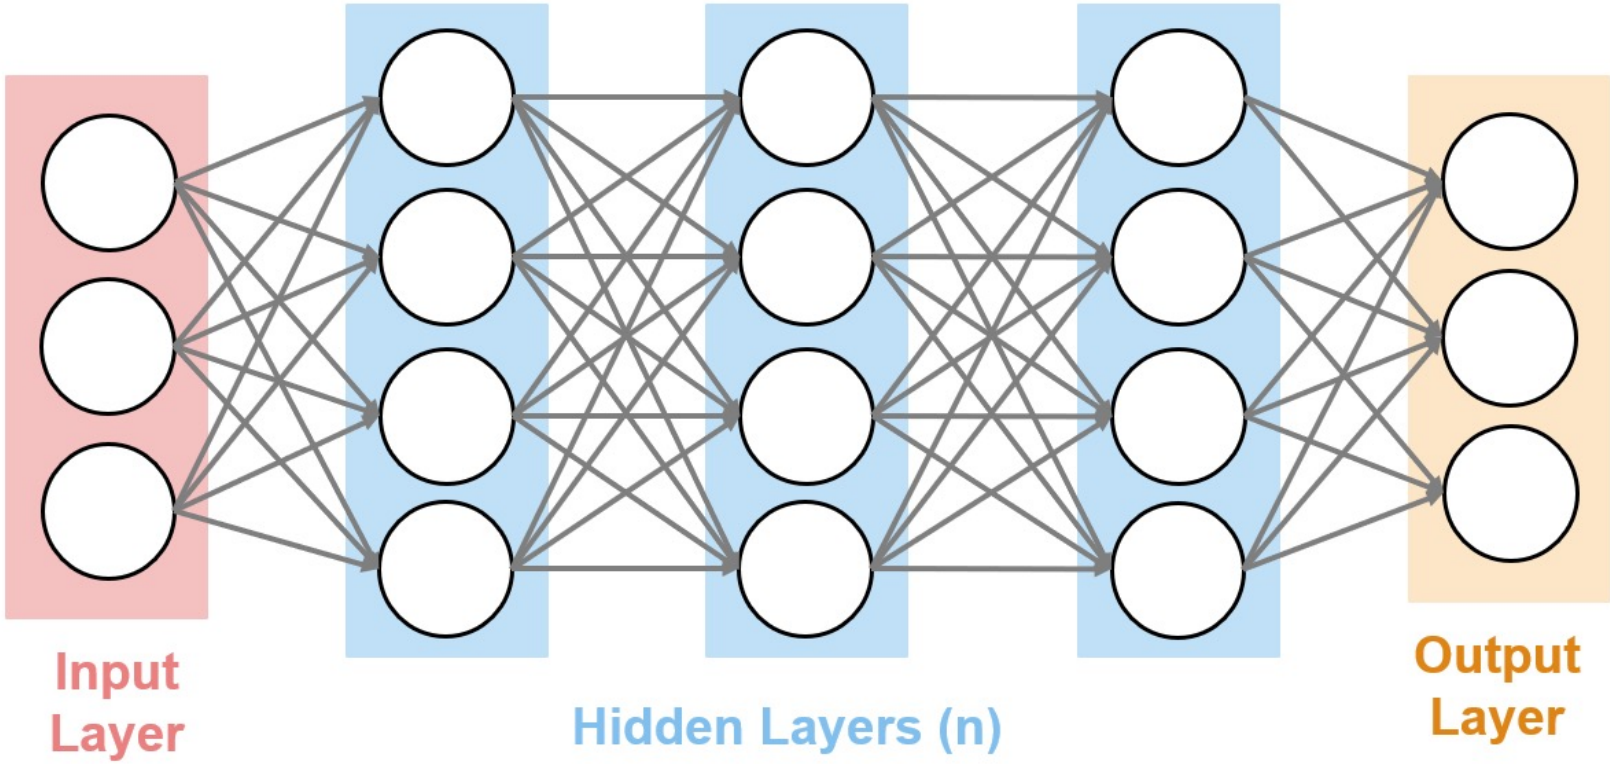

Machine Learning and Deep Learning Datatypes

Image

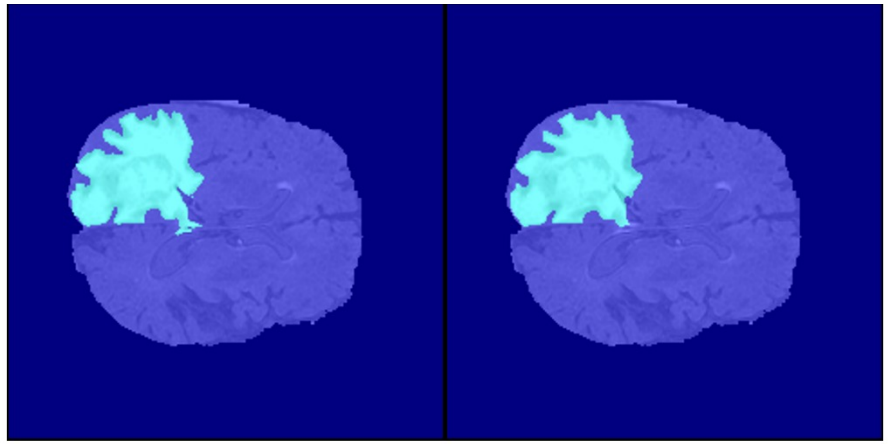

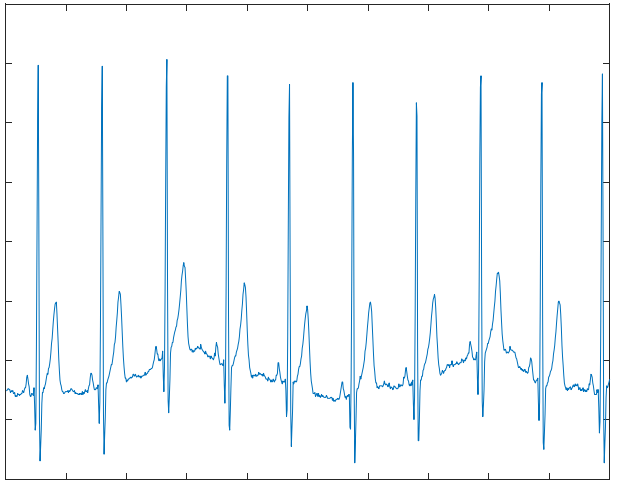

Signal

Numeric

AgeCat	WeightQ	GroupCount	mean_BloodPressure	
Under 30	Q1	6	123.17	79.667
Under 30	Q2	3	120.33	79.667
Under 30	Q3	2	127.5	86.5
Under 30	Q4	4	122	78
30-39	Q1	12	121.75	81.75
30-39	Q2	9	119.56	82.556
30-39	Q3	9	121	83.222
30-39	Q4	11	125.55	87.273
Over 40	Q1	7	122.14	84.714
Over 40	Q2	13	123.38	79.385
Over 40	Q3	14	123.07	84.643
Over 40	Q4	10	124.6	85.1

Text

Deep Learning and AI in Research

University of Southern California

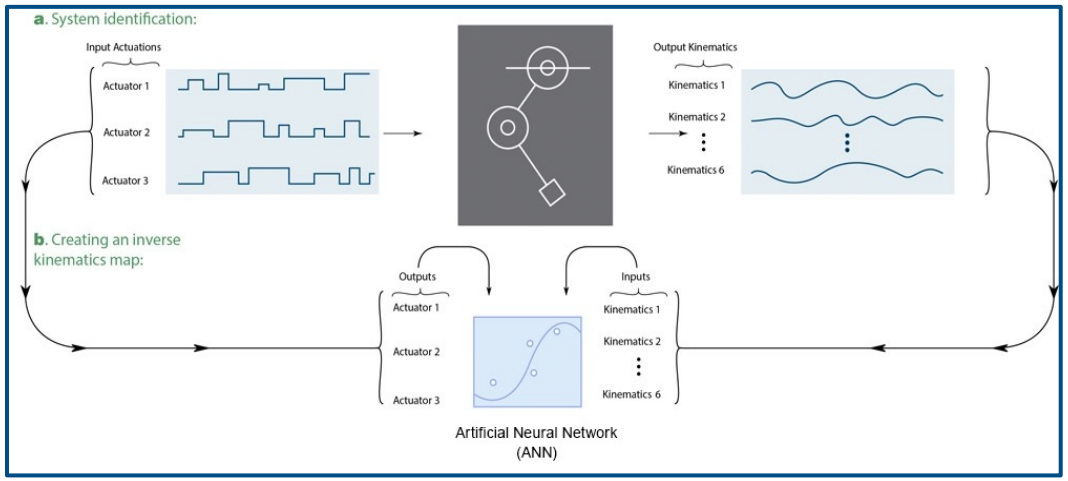

[Reinforcement Learning for Robotic Arm]

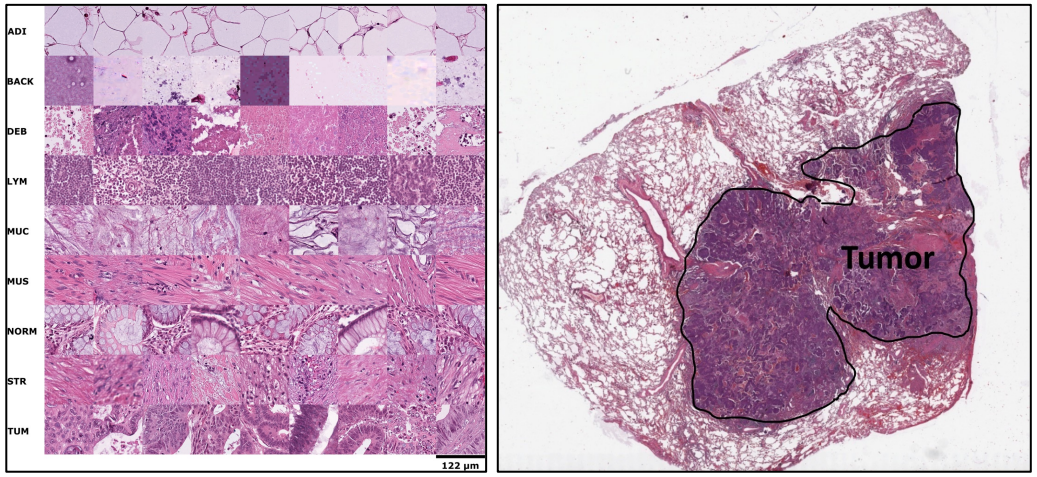

[Deep Learning for Tumor Detection]

DKFZ Heidelberg

University of Twente

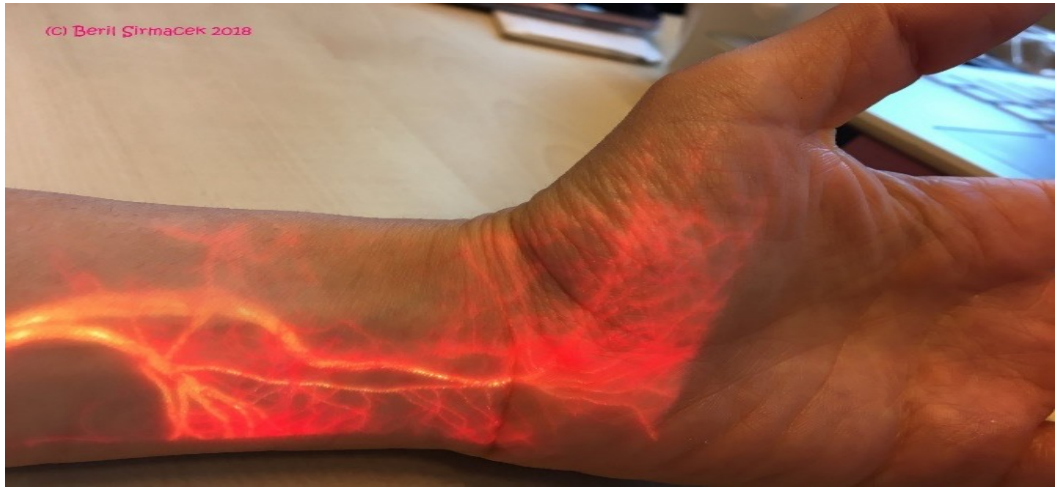

[Augmented Reality of blood flow]

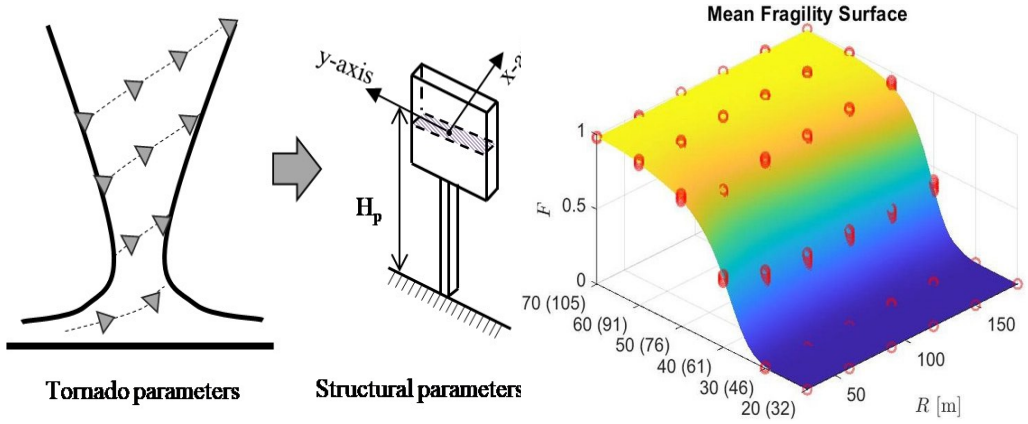

[Neural Networks simulate tornadic wind load]

Northeaster University


Deep Learning Demo
Image Classification

Deep Learning Networks Take in Numeric Data

199	206	208	201	188	178	165	164	180
202	205	202	188	176	169	178	186	183
203	206	189	178	181	183	182	154	87
203	192	184	186	177	167	153	181	192
191	182	176	166	153	141	136	180	227
166	165	154	154	138	137	169	170	211
158	150	145	183	144	156	158	154	179
143	51	98	144	129	130	143	178	123
107	50	33	95	152	173	192	159	87
104	100	84	120	132	172	131	64	94
119	101	97	81	90	109	87	106	111
127	122	110	97	108	120	133	131	134
111	117	108	119	131	143	146	141	156
126	122	113	119	139	142	155	161	151
129	126	130	111	103	130	149	149	156
138	128	136	144	136	129	134	122	145
154	133	134	141	168	150	126	127	151

Images are a numeric matrix

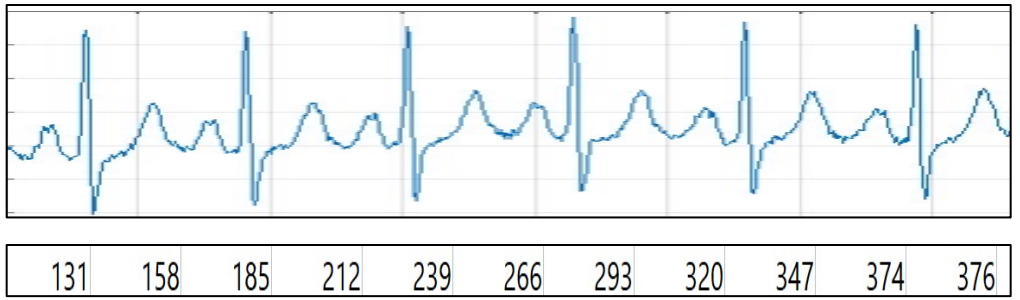

Signals are numeric vectors

The Bird Flies = [0 13 5 6]

The Leaf Is Brown = [13 3 11 2]

Text is processed as numeric vectors

[Deep Learning Documentation]

Agenda

Introduction

Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Break

Exercise 2 & 3: Classify blood smear images

Break

Exercise 4: Improving Network Accuracy

Conclusion

Exercise 1 – Deep Learning in 6 Lines of Code

- Run **runMeFirst.mlx**
- Open **work_deeplearningin6lines.mlx** in **01-DeepLearning101** folder
- Add `load alexnet.mat` at the top to read network
- If you have an account on Alvis, or running on your own computer – Go to the **breakout room** and finish the exercise
- If you don't have a MATLAB in front of you – follow along in Main Room. Let's do the exercise together!

```
% Load the AlexNet neural network
load alexnet.mat;
net = alexnet;
```

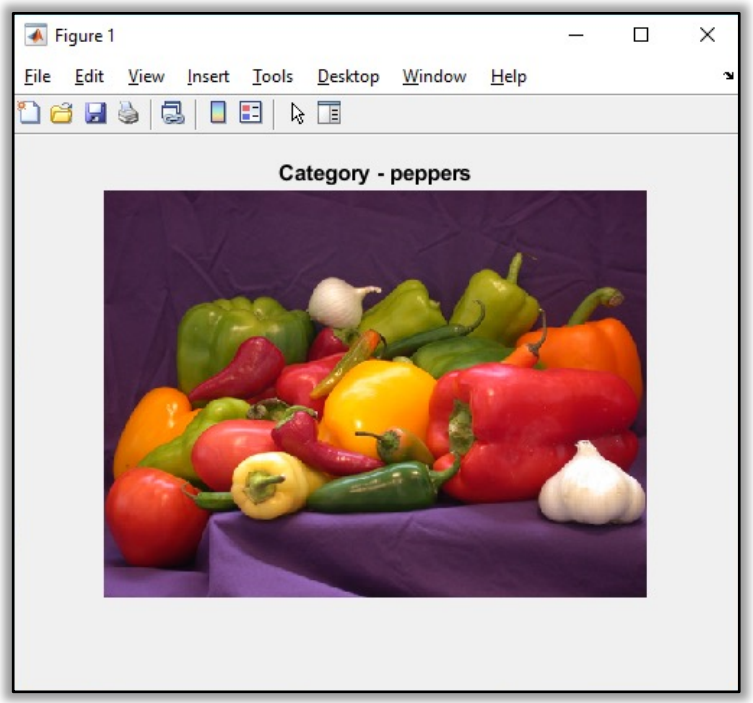

Agenda

Introduction


Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Break


Exercise 2 & 3: Classify blood smear images

Break


Exercise 4: Improving Network Accuracy

Conclusion

Creating Layer Architectures

- Convolutional Neural Networks – CNN
- Special layer combinations which learn parameters to classify images

- Convolution Layer
- ReLU Layer
- Max Pooling Layer

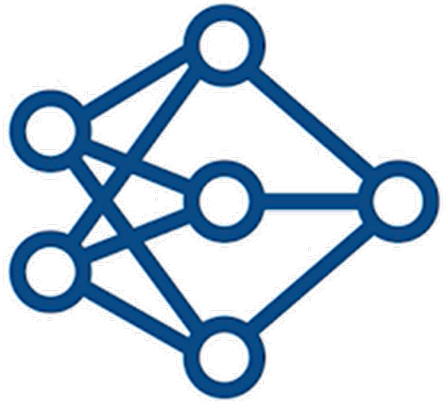

Convolution Layers Search for Patterns

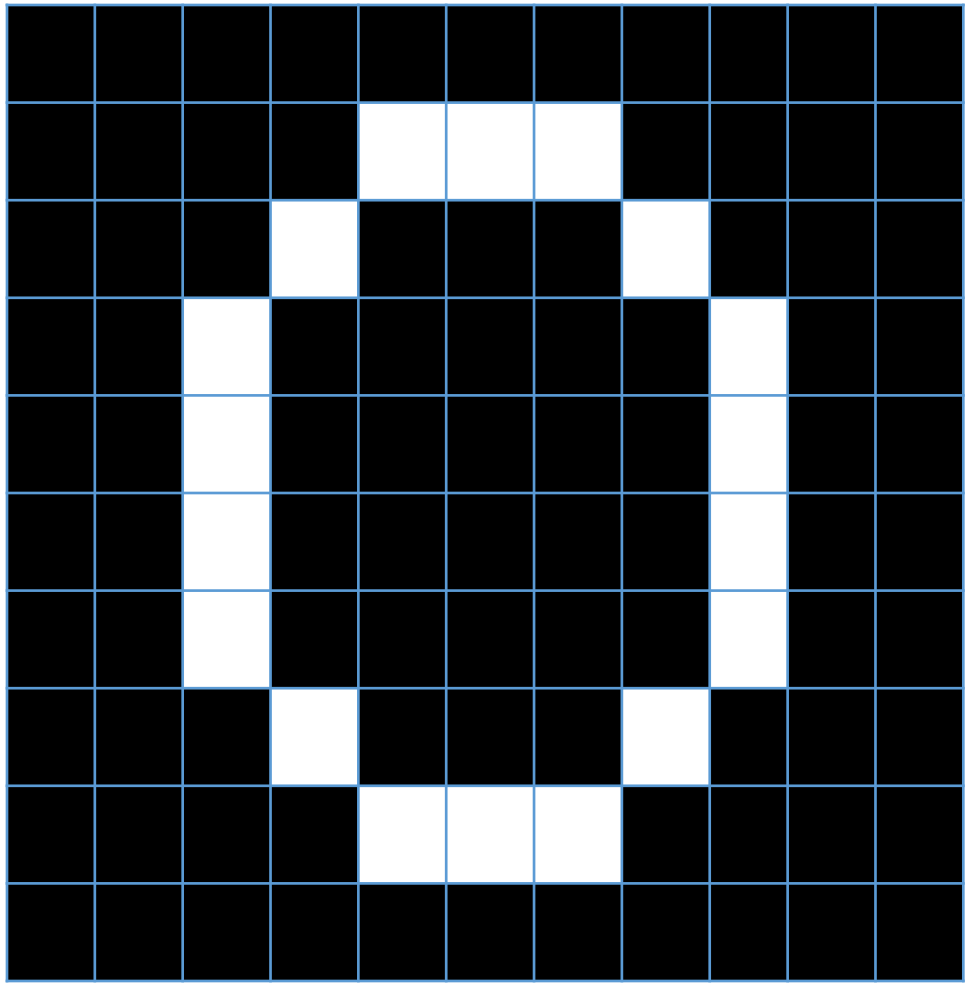

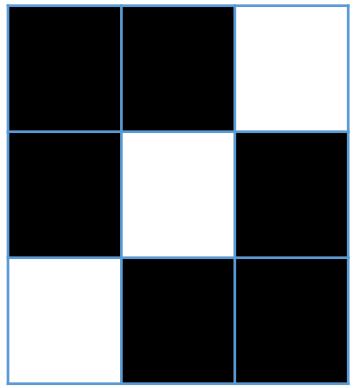

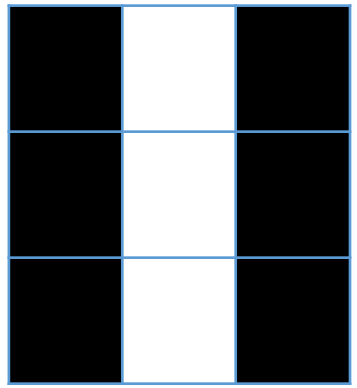

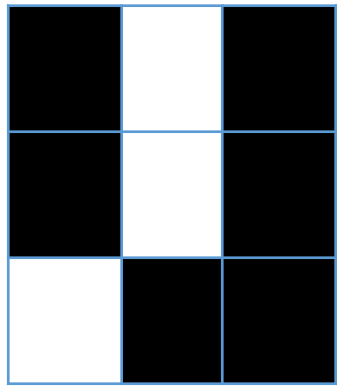

These patterns would be common in the number 0

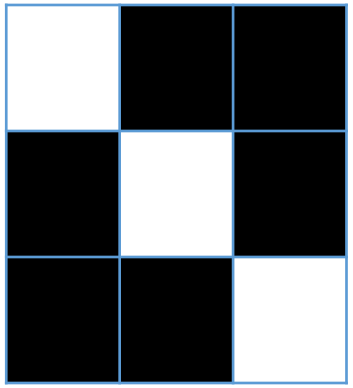

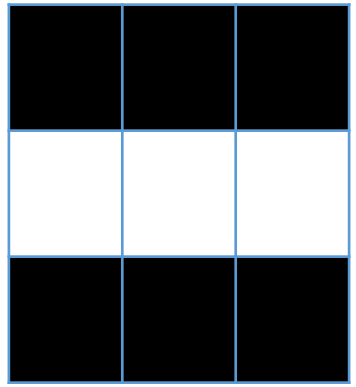

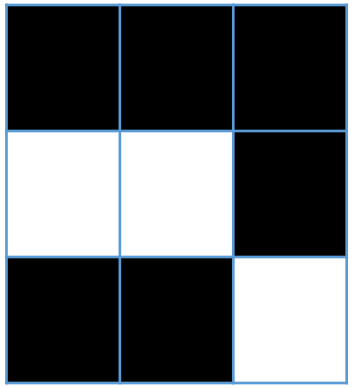

All patterns are compared to the patterns on a new image

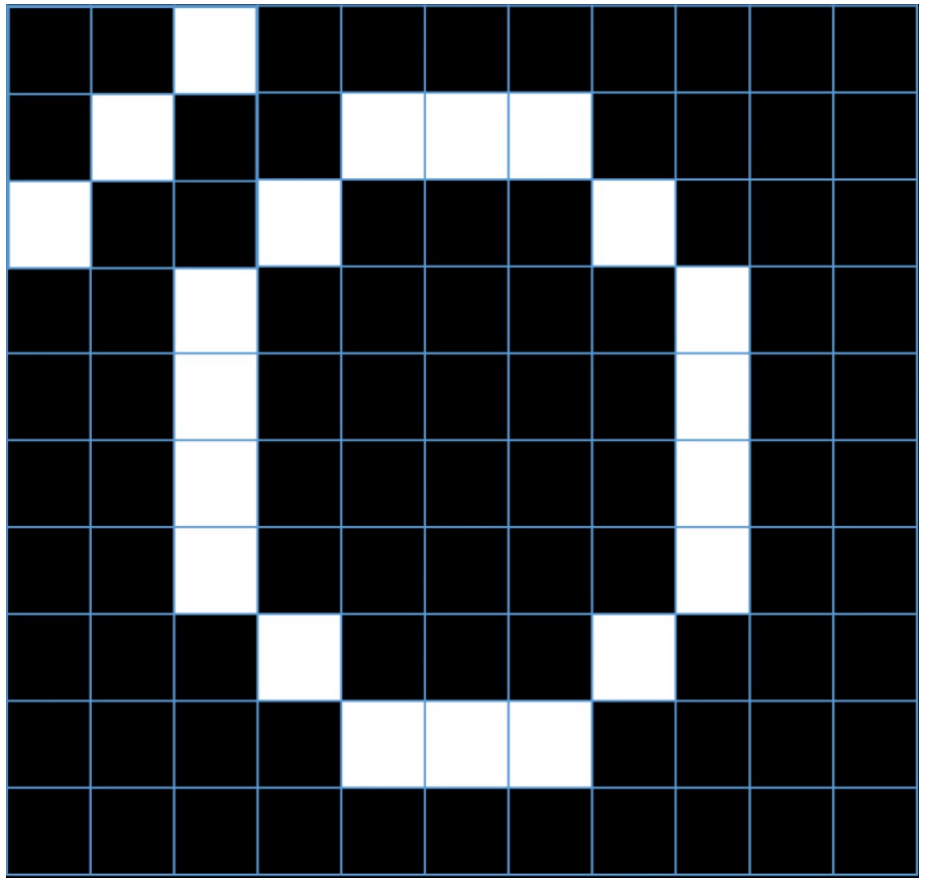

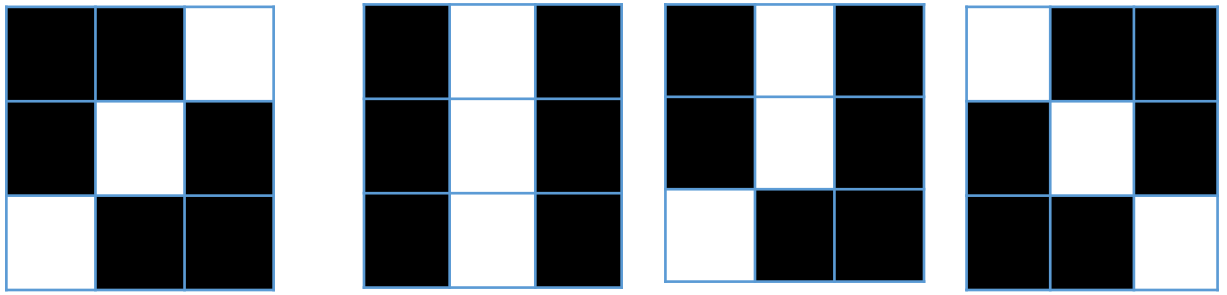

- Pattern starts at left corner
  - Perform comparison
  - Slide over one pixel
- Reach end of image
- Repeat for next pattern

Good pattern matching in convolution improves chances that object will classify properly

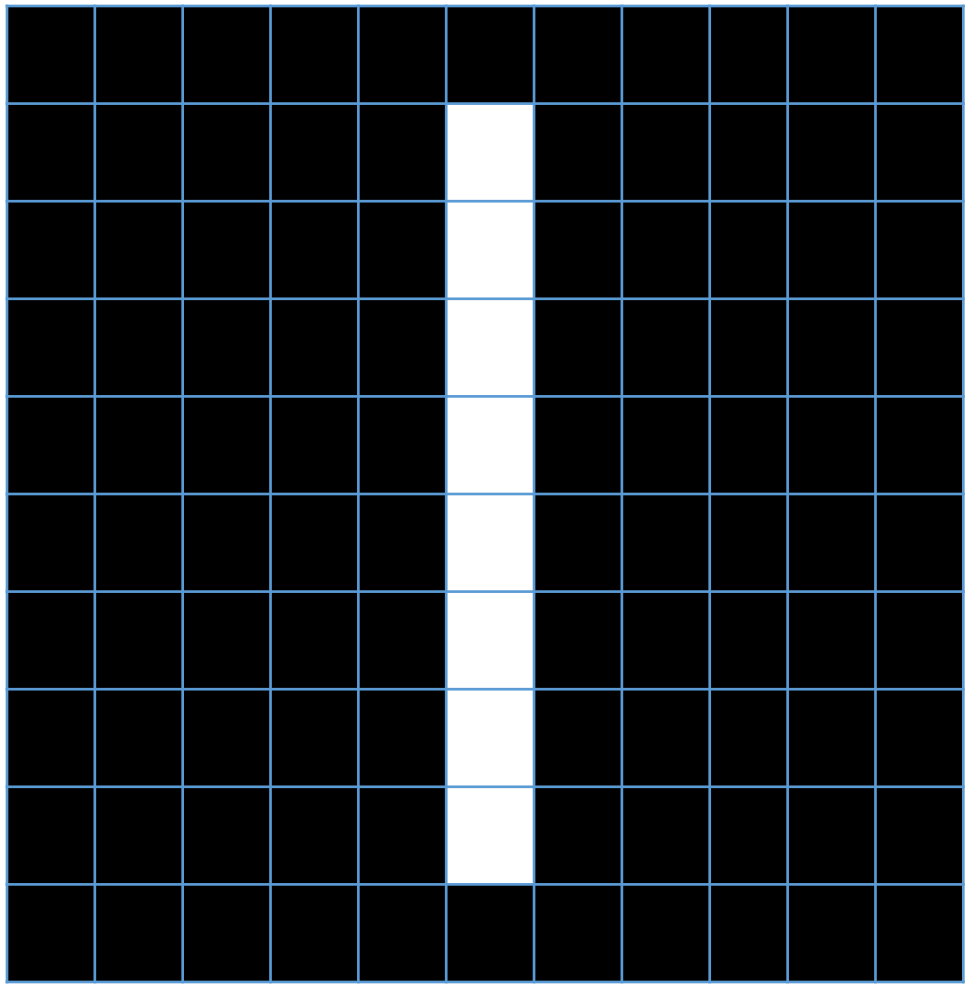

- This image would not match well against the patterns for the number zero
- It would only do very well against this pattern

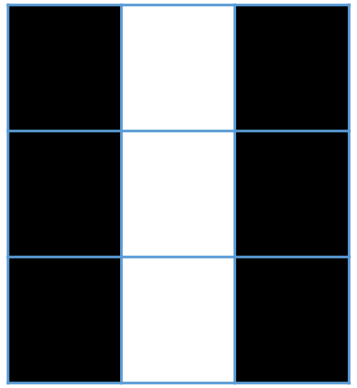

Activation: Rectified Linear Units Layer (ReLU)

Converts negative numbers to zero

-1	0	5	4
3	-4	-8	3
1	4	6	-5
-2	-5	4	1

0	0	5	4
3	0	0	3
1	4	6	0
0	0	4	1

Max Pooling is a down-sampling operation

Shrink large images while preserving important information

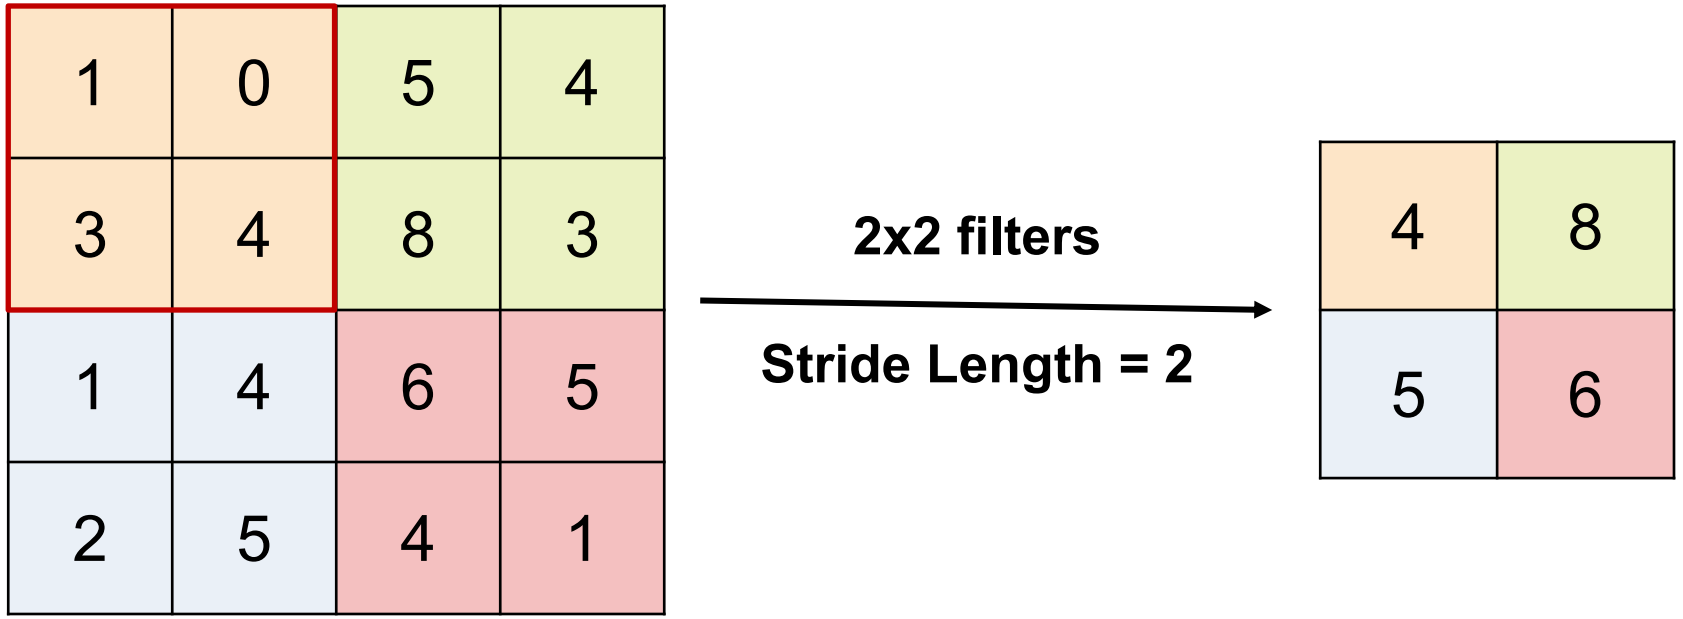

Classification Problems End with 3 Layers

- Fully Connected Layer
  - Looks at which high-level features correspond to a specific category
  - Calculates scores for each category (highest score wins)

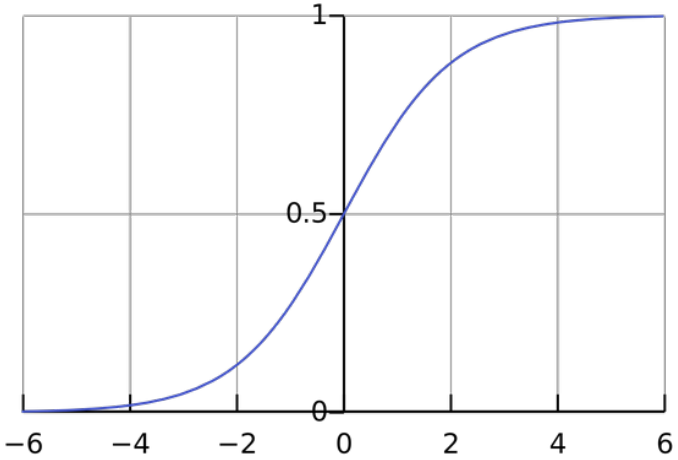

- Softmax Layer
  - Turns scores into probabilities.
- Classification Layer
  - Categorizes image into one of the classes that the network is trained on

Note: Regression problems end with a fully connected layer and regression layer

3 Components to Train any Network

"How much data do I need?"

It depends on the complexity of the problem…
Is it complex?
Then **a lot**

Define inputs and layers for deep learning model

Influence training time and accuracy

- Solver Type
- Initial Learn Rate
- Minibatch Size
- Max Epochs
- …

Agenda

Introduction

 Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Break

 Exercise 2 & 3: Classify blood smear images

Break

 Exercise 4: Improving Network Accuracy

Conclusion

Agenda

Introduction


Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Break


Exercise 2 & 3: Classify blood smear images

Break


Exercise 4: Improving Network Accuracy

Conclusion

Exercise 2 – Create a network from scratch

- Open **work_DeepLearningFormScratch.mlx** in **02-DeepLearningFromScratchUsingCNN** folder
- If you have an account on Alvis, or running on your own computer – Go to the **breakout room** and finish the exercise
- If you don't have a MATLAB in front of you – follow along in Main Room. Let's do the exercise together!

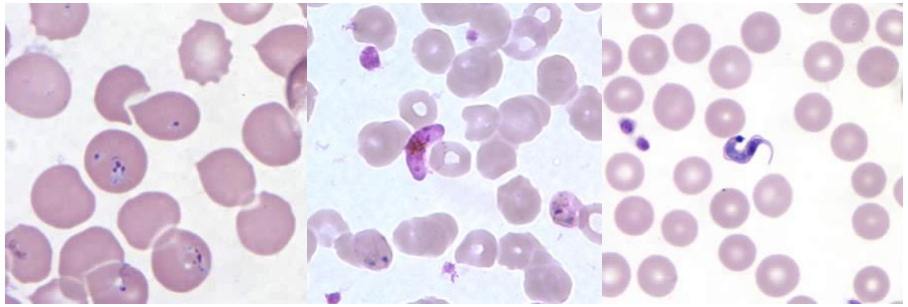

Task:
Create a network that can differentiate between different blood parasites

Workflow:
- Read in data from a folder
- Analyze the data
- Set up a CNN network
- Train network
- Test network

Transfer Learning Workflow

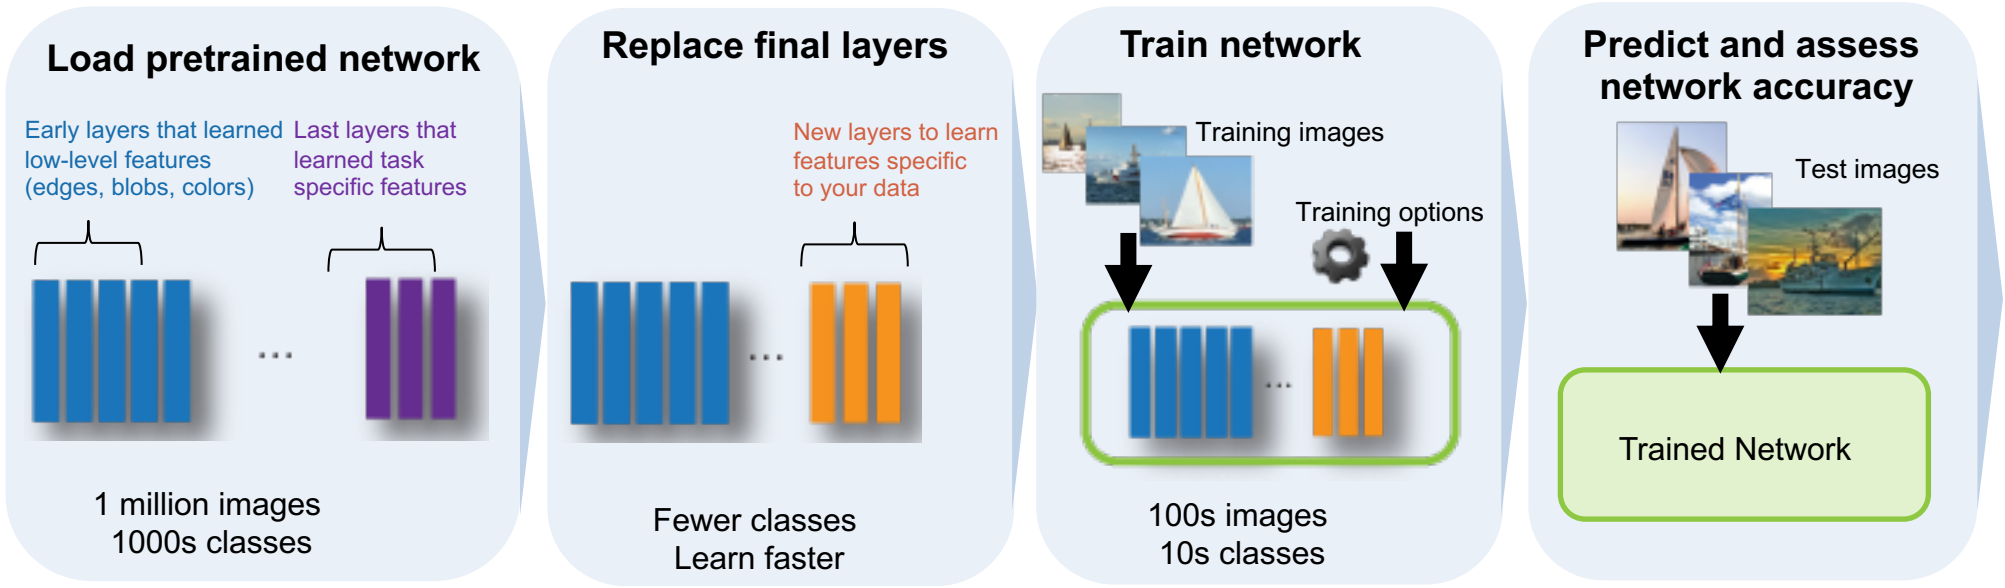

Transfer Learning Workflow – Step 1

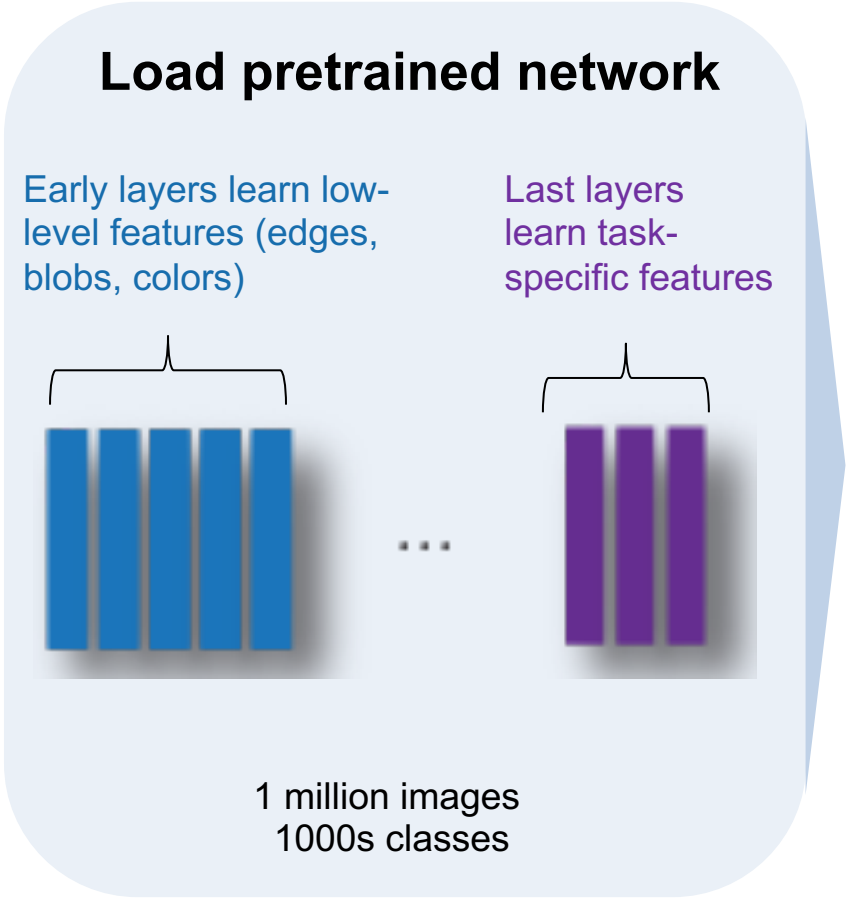

Transfer Learning Workflow – Step 2

Load pretrained network

Early layers that learned low-level features (edges, blobs, colors)

Last layers that learned task specific features

1 million images
1000s classes

Replace final layers

New layers learn features specific to your data

Fewer classes
Learn faster

Transfer Learning Workflow – Step 3

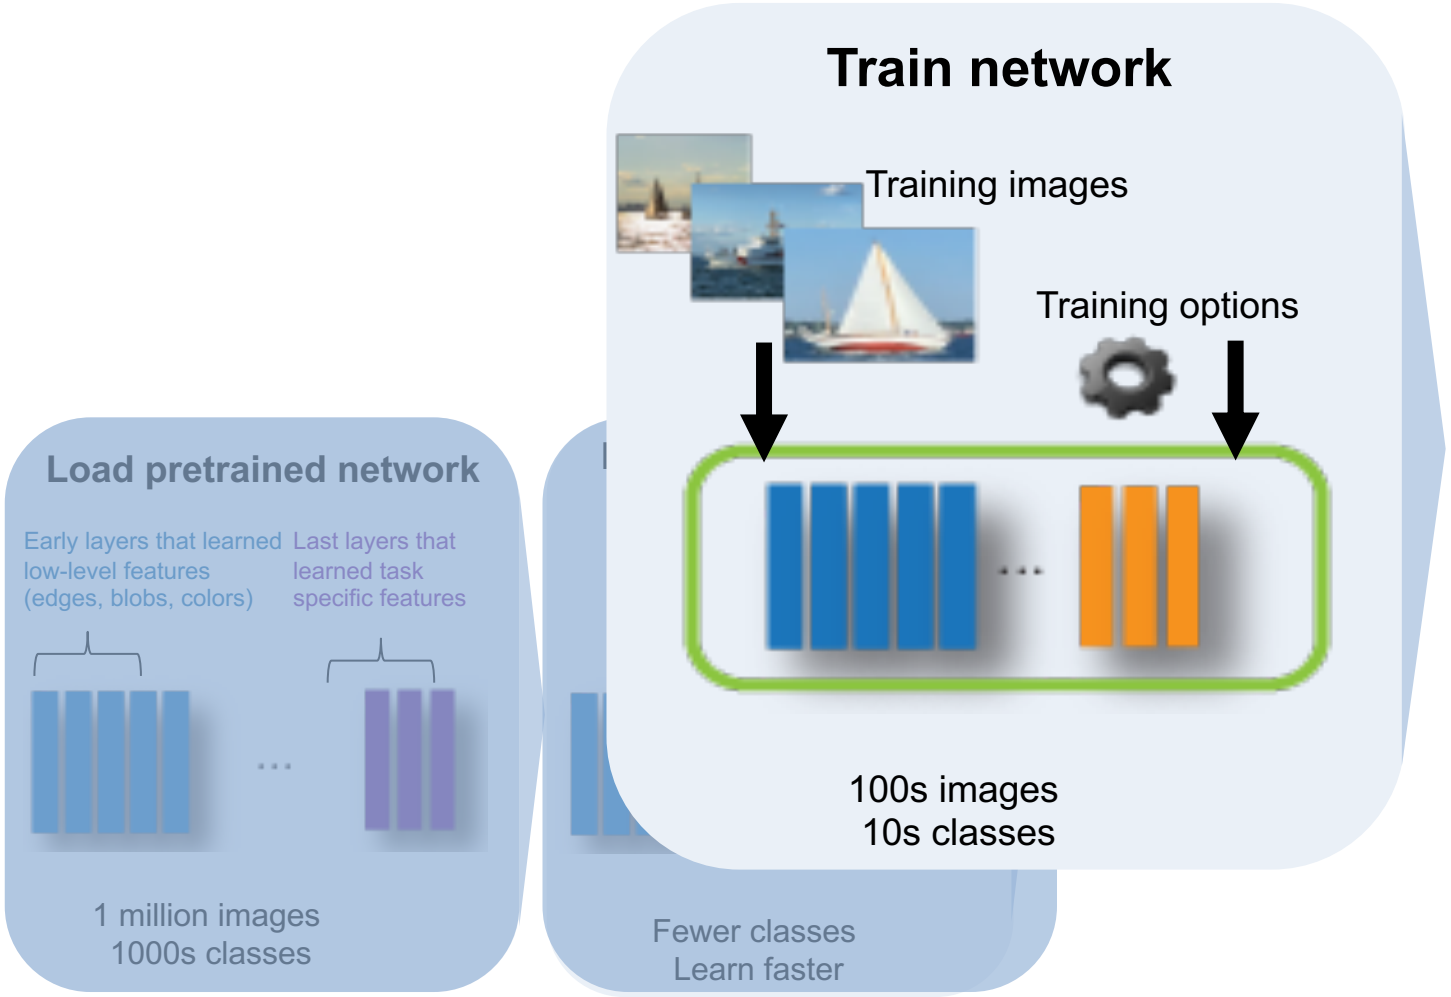

Transfer Learning Workflow – Step 4

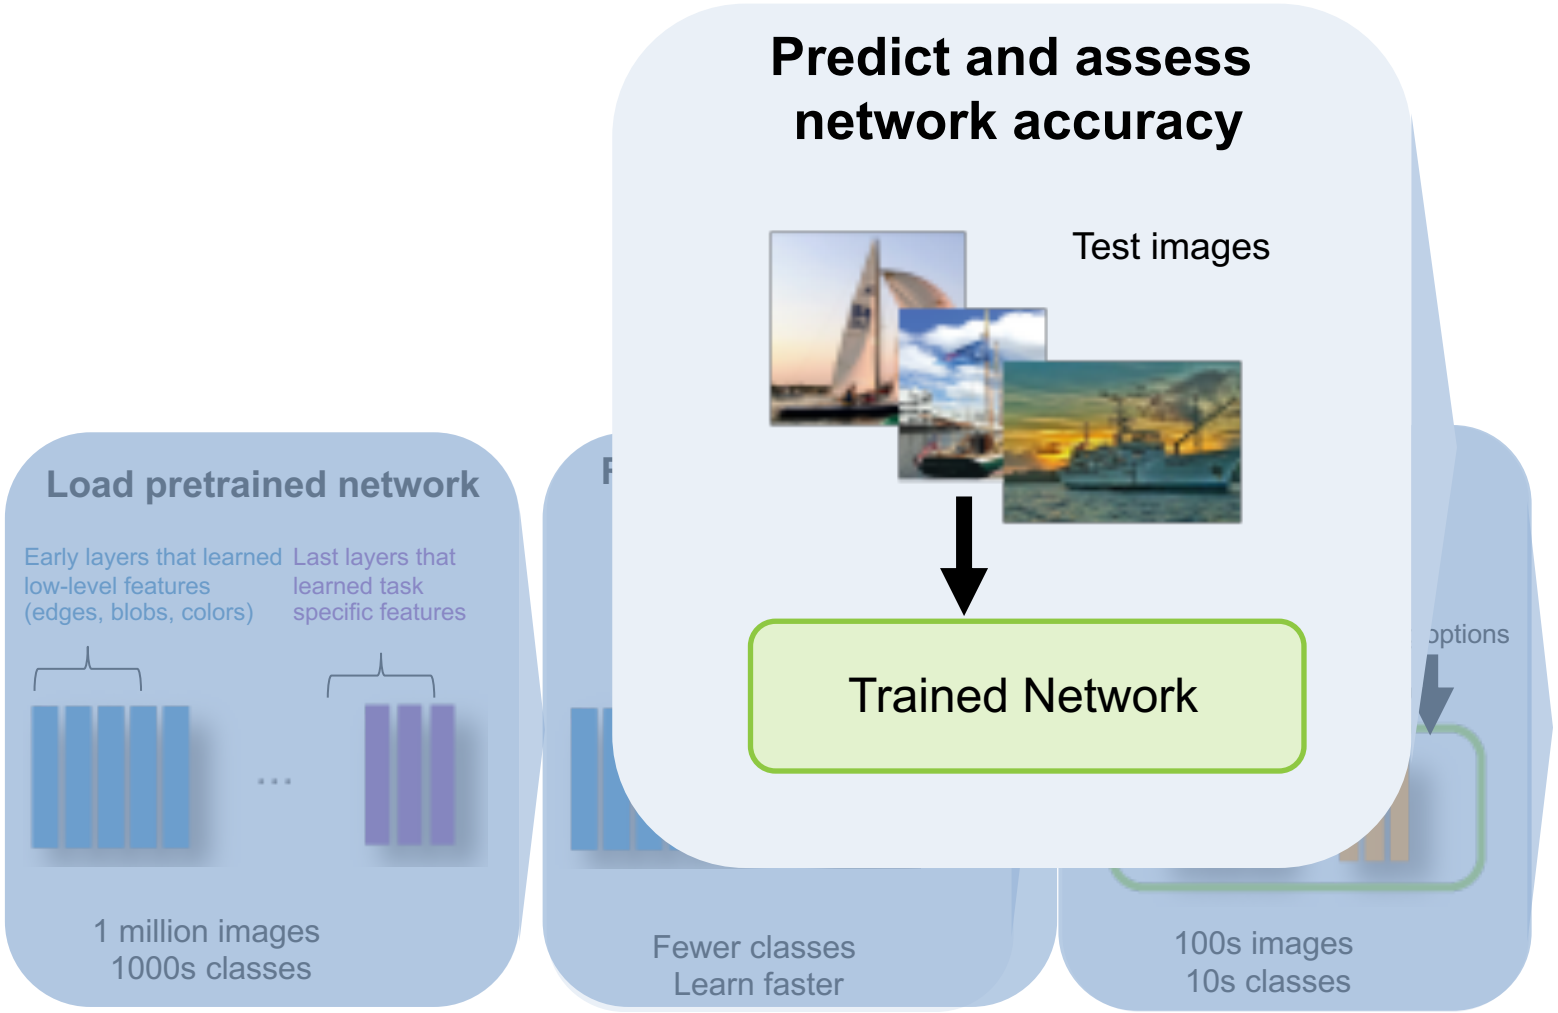

Transfer Learning Workflow

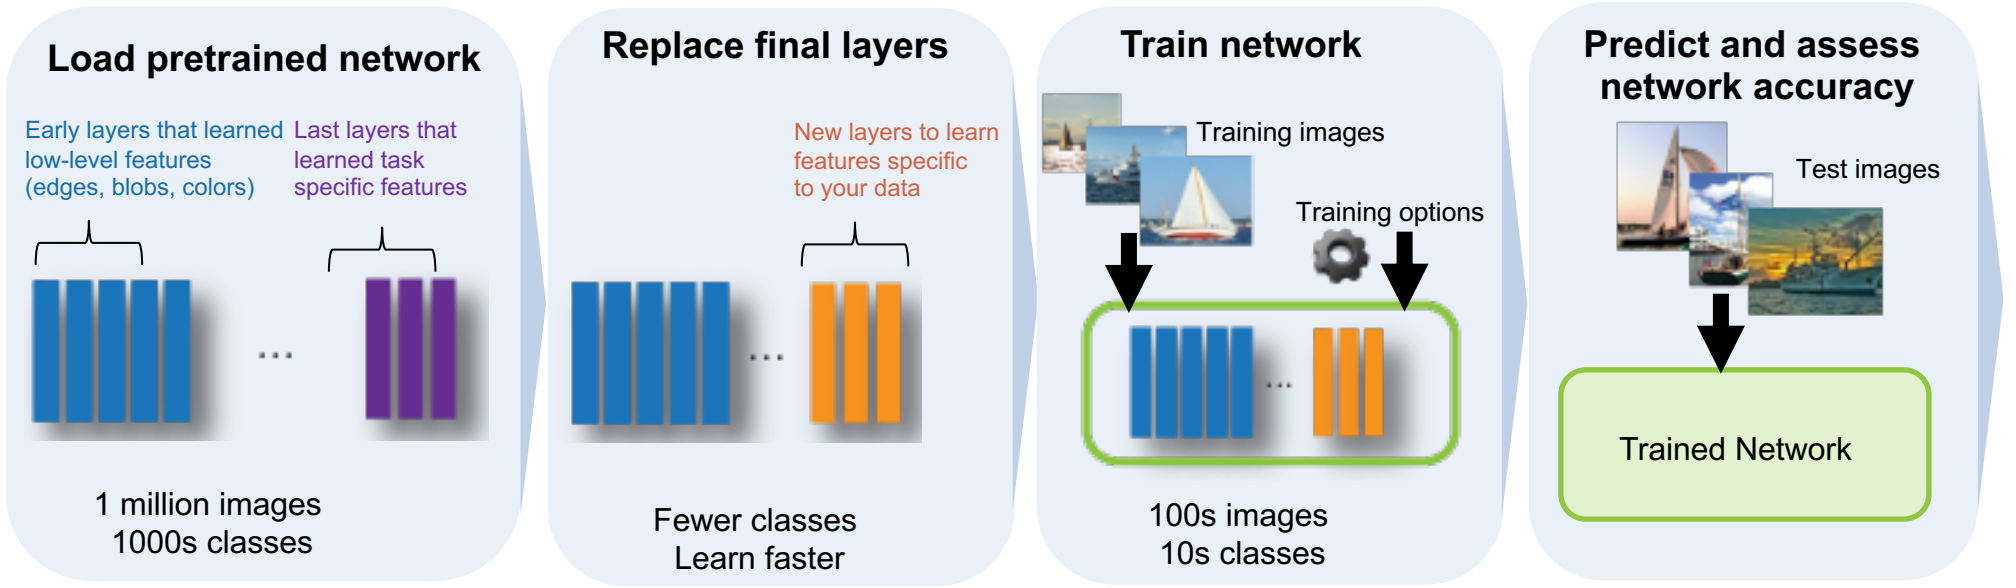

Exercise 3 – Transfer Learning

- Open **work_DeepLearningTransferLearning.mlx** in **03-TransferLearningWithCNN** folder
- If you have an account on Alvis, or running on your own computer – Go to the **breakout room** and finish the exercise
- If you don't have a MATLAB in front of you – follow along in Main Room. Let's do the exercise together!

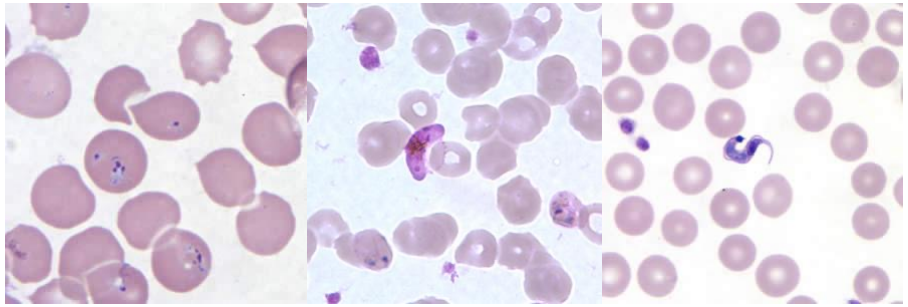

Task:
Create a network that can differentiate between different blood parasites using a pretrained network

Workflow:
- Read in data from a folder
- Read in pretrained network
- Change layers in the network
- Train network
- Test network

Agenda

Introduction

 Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Break

 Exercise 2 & 3: Classify blood smear images

Break

 Exercise 4: Improving Network Accuracy

Conclusion

Agenda

Introduction

 Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Break

 Exercise 2 & 3: Classify blood smear images

Break

 ***Exercise 4: Improving Network Accuracy***

Conclusion

Why is my neural network giving me “incorrect” results?

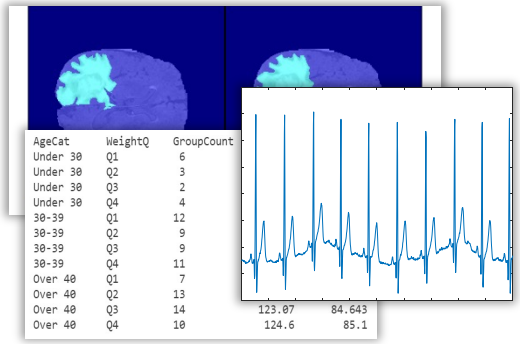

Data

- Is there **enough** data?
- Do we have the **right** data for the application?
- What about the **quality** of the data?
- Is data **augmentation** in use?

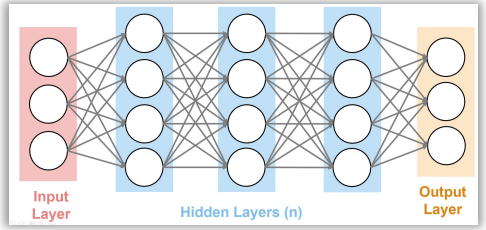

Neural Network Architecture

- What layers are being used and why?
- Are the **appropriate layers** being used for the application?
- Is there **normalization** and/or **regularization**?
- Are you validating during training to check for **overfitting**?

Parameters and Their Settings

- What parameters are present?
- Were the batch size, learn rate, weights, and other values on “default”?

Exercise 4 – Improving Network Accuracy

- Open **work_DeepLearningImprovingAccuracy.mlx** in **04-ImprovingNetworkAccuracy** folder
- If you have an account on Alvis, or running on your own computer – Go to the **breakout room** and play around with different settings
- If you don't have a MATLAB in front of you – follow along in Main Room. Let's do the exercise together!

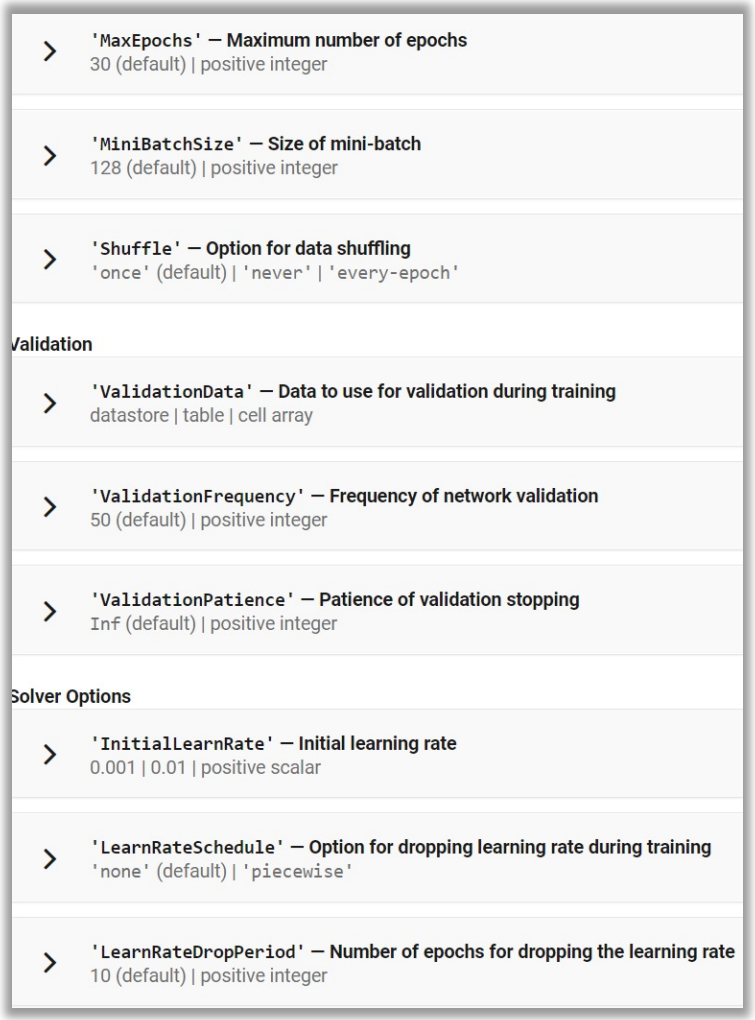

Hyperparameter Tuning in Deep Learning

"The challenge with hyperparameters is that there are no magic number that works everywhere. The best numbers depend on each task and each dataset." Source

Hyperparameter tuning using Bayesian Optimization

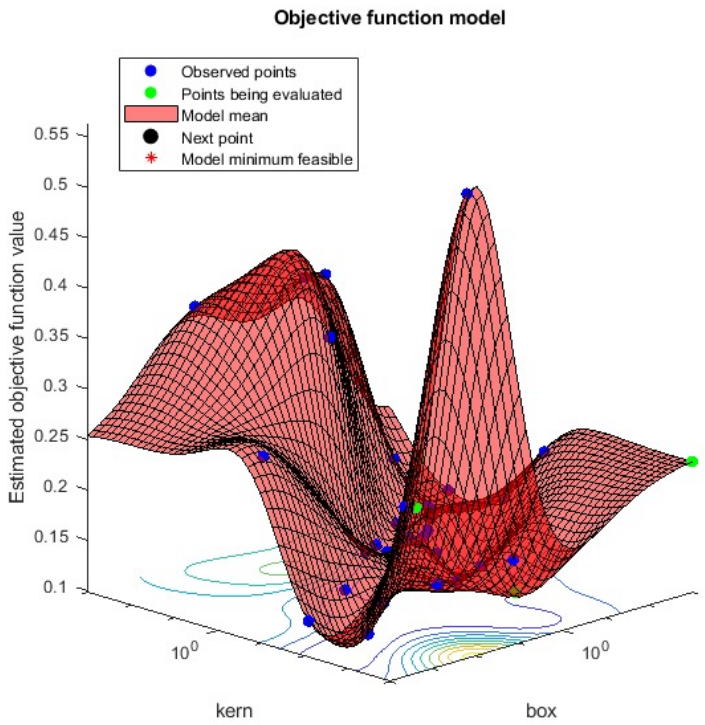

Experiment Manager

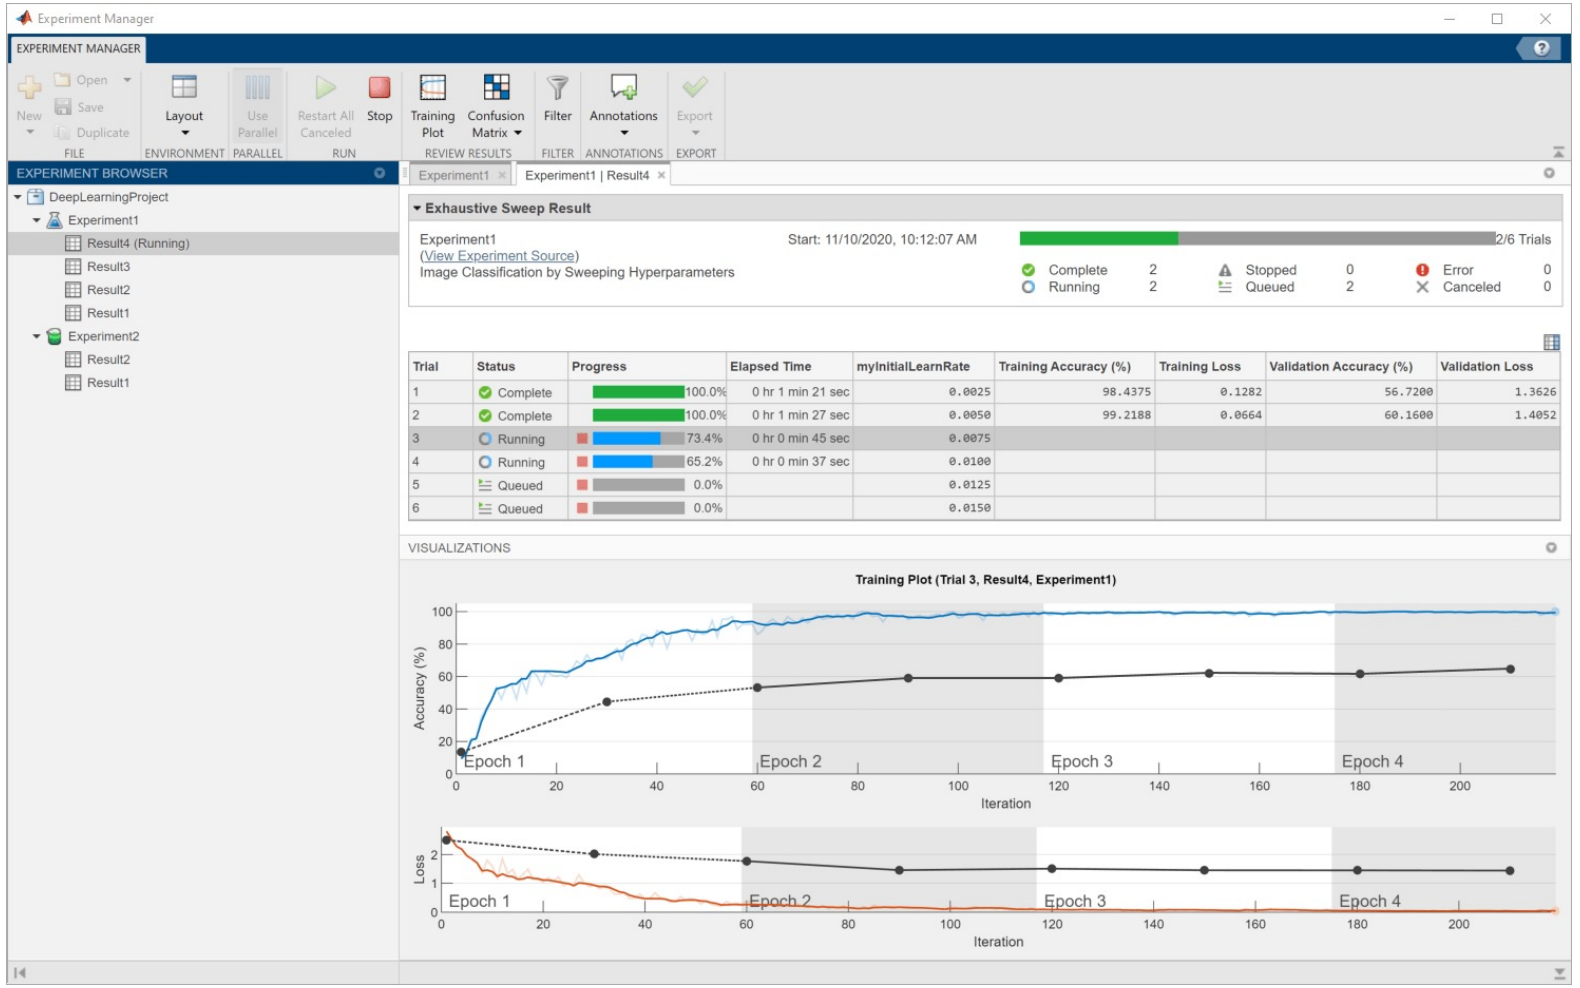

Agenda

Introduction

 Exercise 1: Deep learning in 6 lines of code

Deep Learning Fundamentals

Break

 Exercise 2 & 3: Classify blood smear images

Break

 Exercise 4: Improving Network Accuracy

Conclusion

Home Exercises

Try LSTM models for time series data:

- Classify heart status based on ECG signals
- work_ClassifyECGSignals.mlx
- https://sft.mathworks.com/human.aspx?r=245144200&arg06=776492647-15764732a8bb11675e2f947023798fa0&arg12=filelist
- Username: AppleHill Password: fjr5n8nf

Try a larger dataset for GPU speed up on Alvis:

- Classify pixels in a Semantic Segmentation problem
- https://se.mathworks.com/help/vision/ug/semantic-segmentation-using-deep-learning.html
- ~5 hours training time on a low performing GPU. What can you get on the Tesla v100??

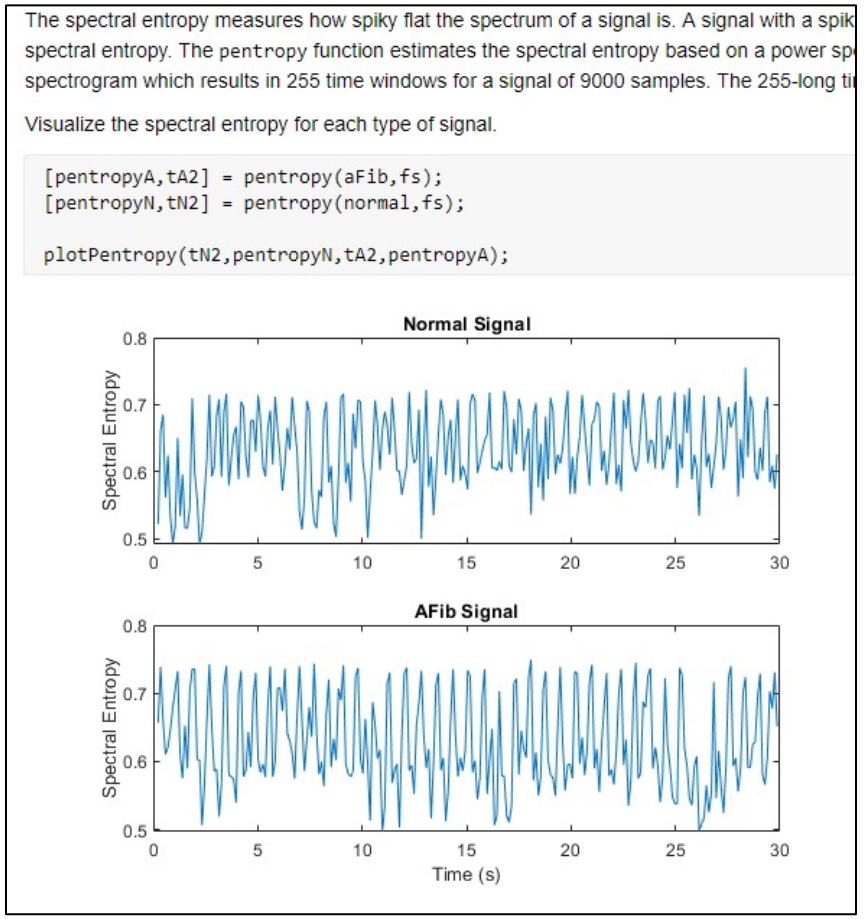

Further Learning and Teaching

- Deep Learning Onramp
  - 2 hr online tutorial
- Deep Learning with MATLAB self-paced course
  - >14 hours interactive tutorial
- Deep Learning Workshop
  - 3 hr hands on session
  - Contact us to schedule
- Deep Learning Training
  - 16 hr in depth course
  - Online or Instructor Lead
- Teaching Deep Learning with MATLAB
  - Curriculum support

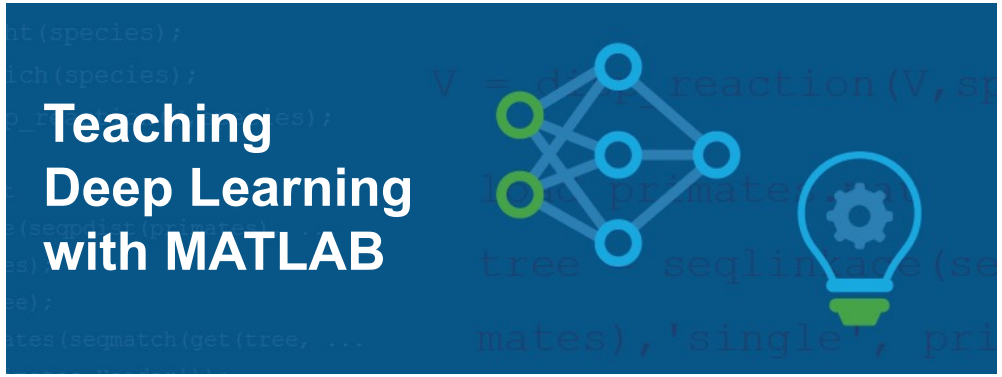

Good Luck with the home exercises!