

Using Deep Learning to Classify Lymphoma Types

Extraterrestrial Confounding Cats

Nicole Firestone, Keerthana Jegatheesan, Elisa Tau, Alejandro Cartes and Zachary Richards

TA: Daniela Grandón

Mentor: Mauricio Cerda

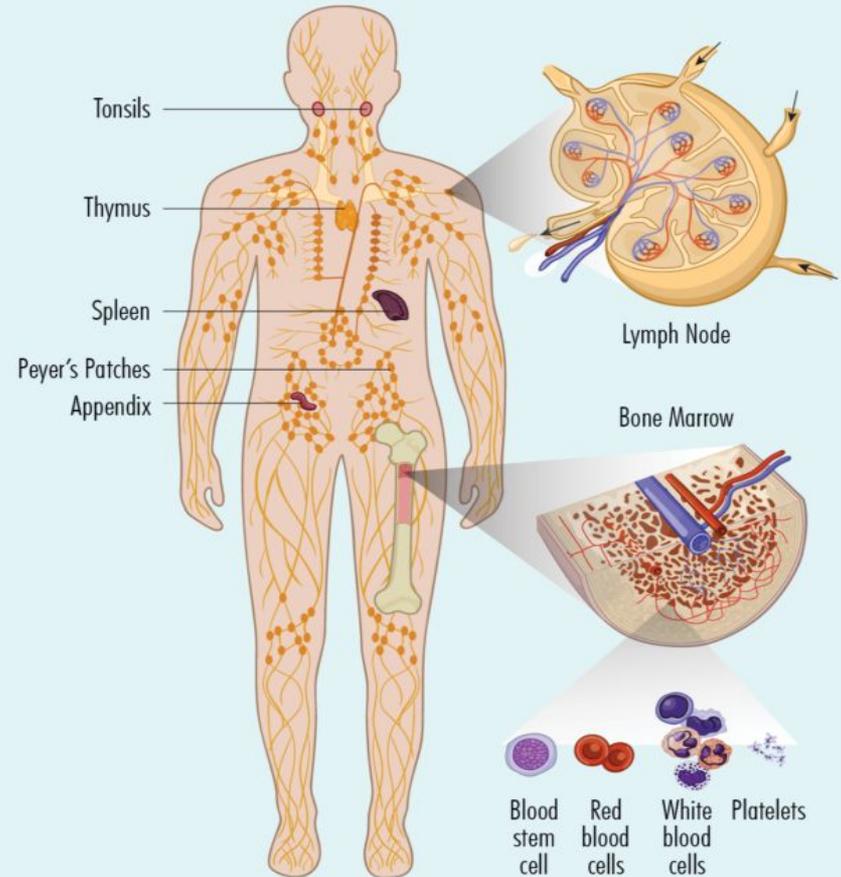
Outline

1. Scientific Context
2. Our data
3. Methods
4. Results
5. Conclusions
6. Future work and connections to astronomy

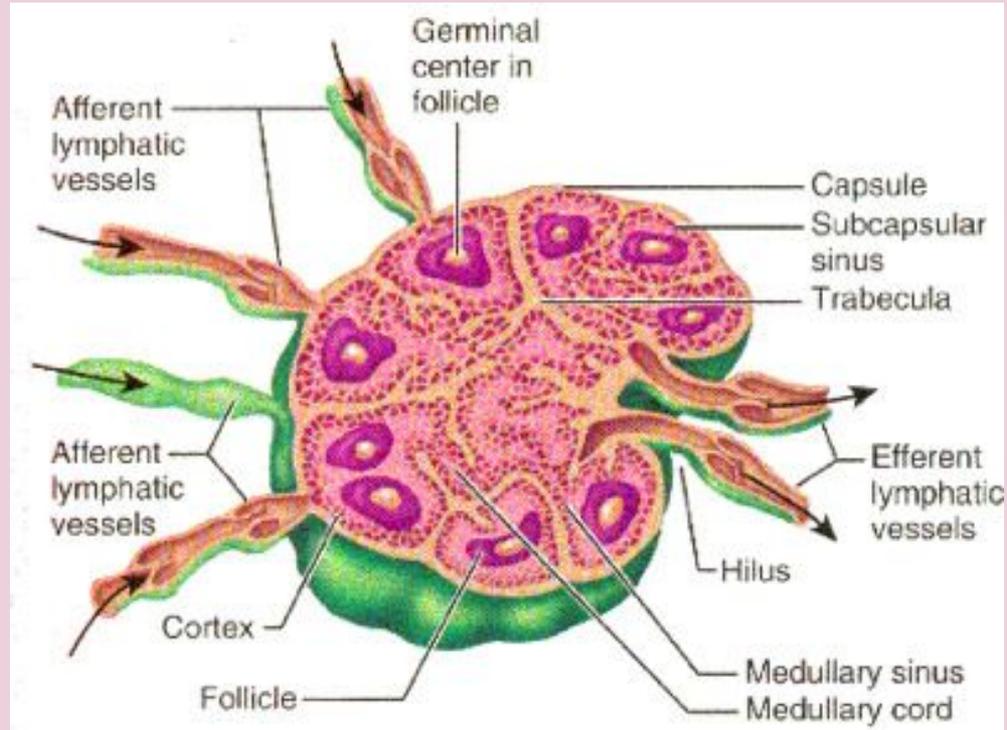
Scientific Context: What is lymphoma?

Lymphoma is a cancer of the lymphatic system.

Lymphatic System

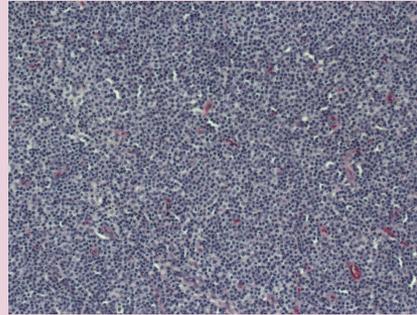


Scientific Context: Anatomy of a Lymph Node

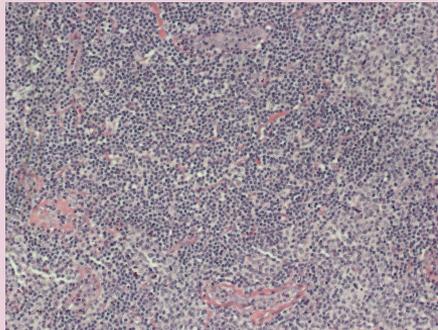


Scientific Context: Three Classes of Lymphoma

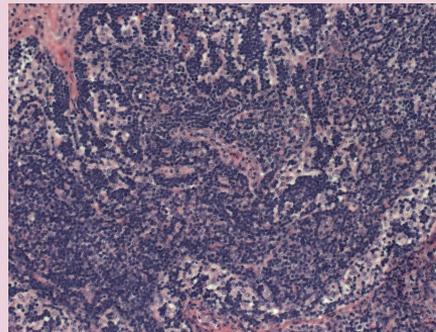
Chronic Lymphocytic Leukemia (CLL)



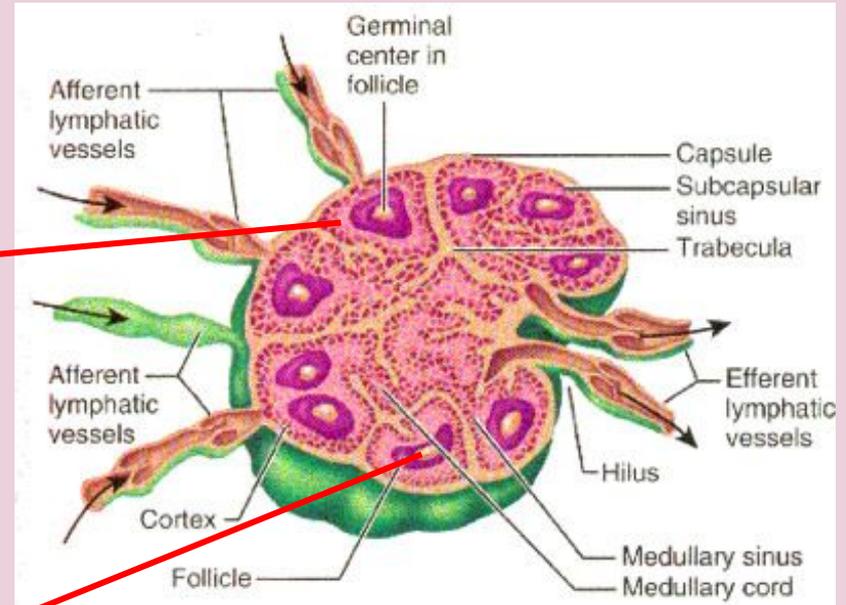
Mantle Cell Lymphoma (MCL)



Follicular Lymphoma (FL)

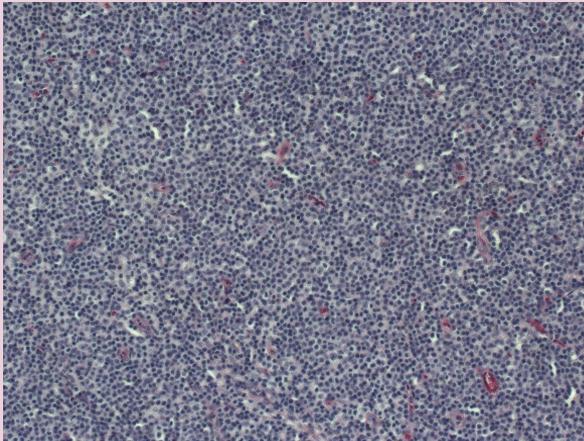


Anatomy of a Lymph Node

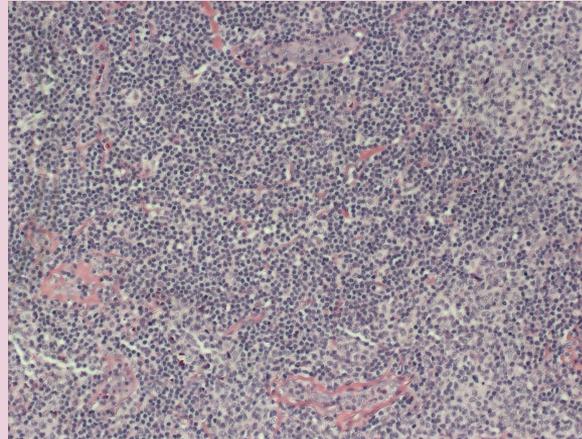


Scientific Context: Three Classes of Lymphoma

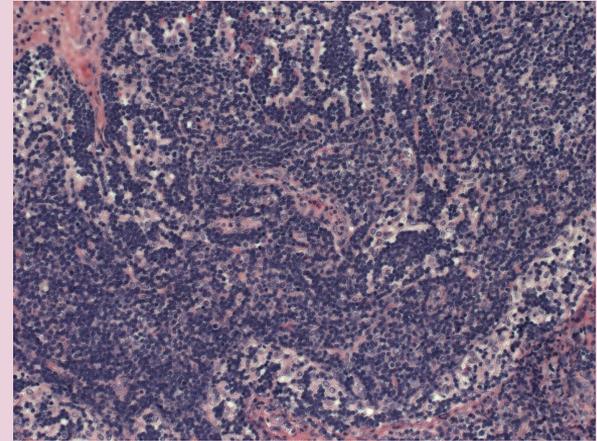
Chronic Lymphocytic
Leukemia (CLL)



Mantle Cell Lymphoma (MCL)



Follicular Lymphoma (FL)



Biopsies stained with hematoxylin and Eosin → Lymphocytes appear blue

Motivation

Challenge:

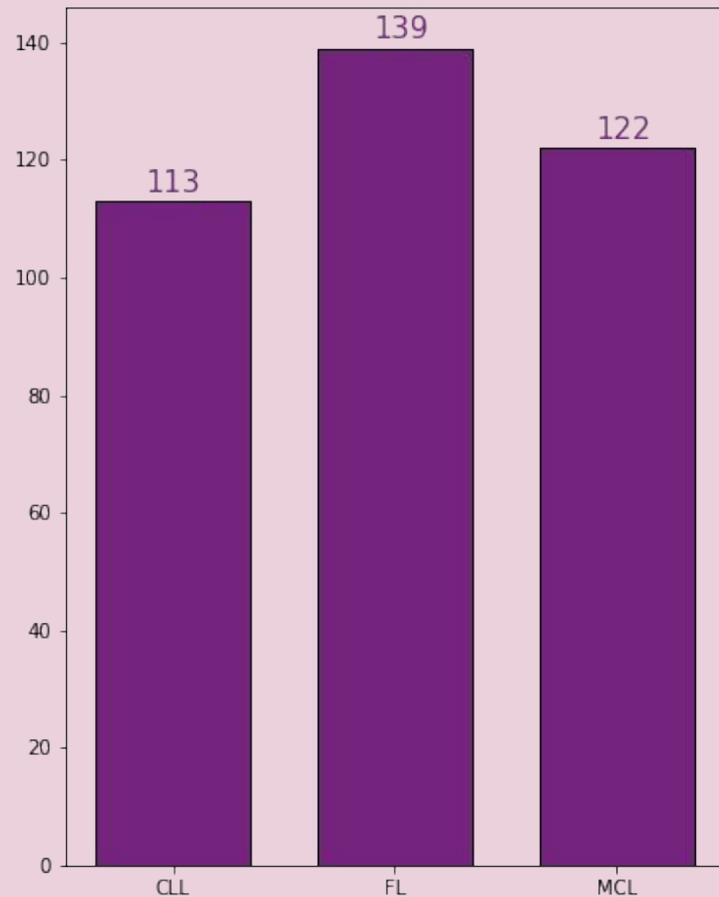
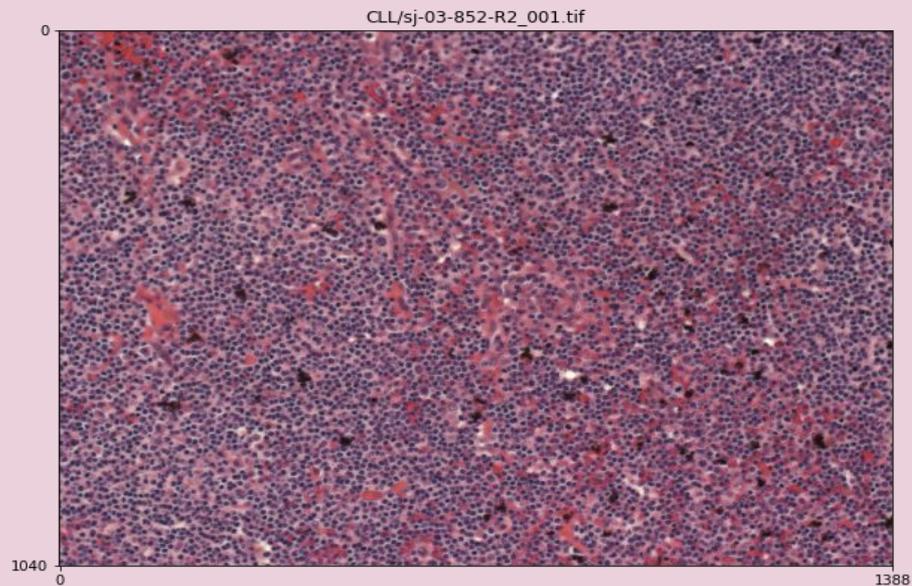
Lymphoma classification is inaccessible, unstandardized, and time sensitive- requires efficient and consistent classification to assign most promising treatment options.

Solution:

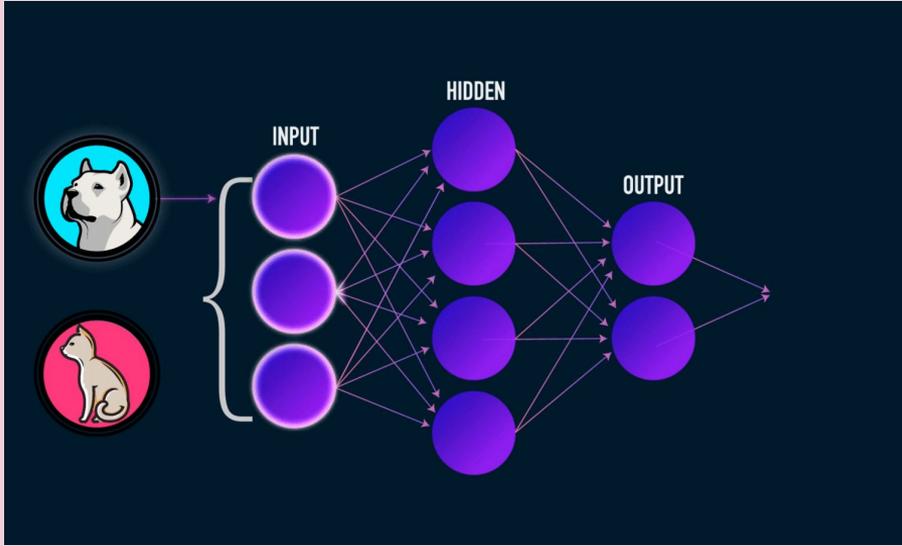
Develop a convolution neural network to properly classify 3 types of lymphoma in order to speed up and standardize the diagnostic process, making diagnostics tools more accessible.

The Image Data

- Format: .tif
- Count: 374
- Classes 3
- Dimensions: 1040x1388 pixels
- Channels: 3 (rgb)



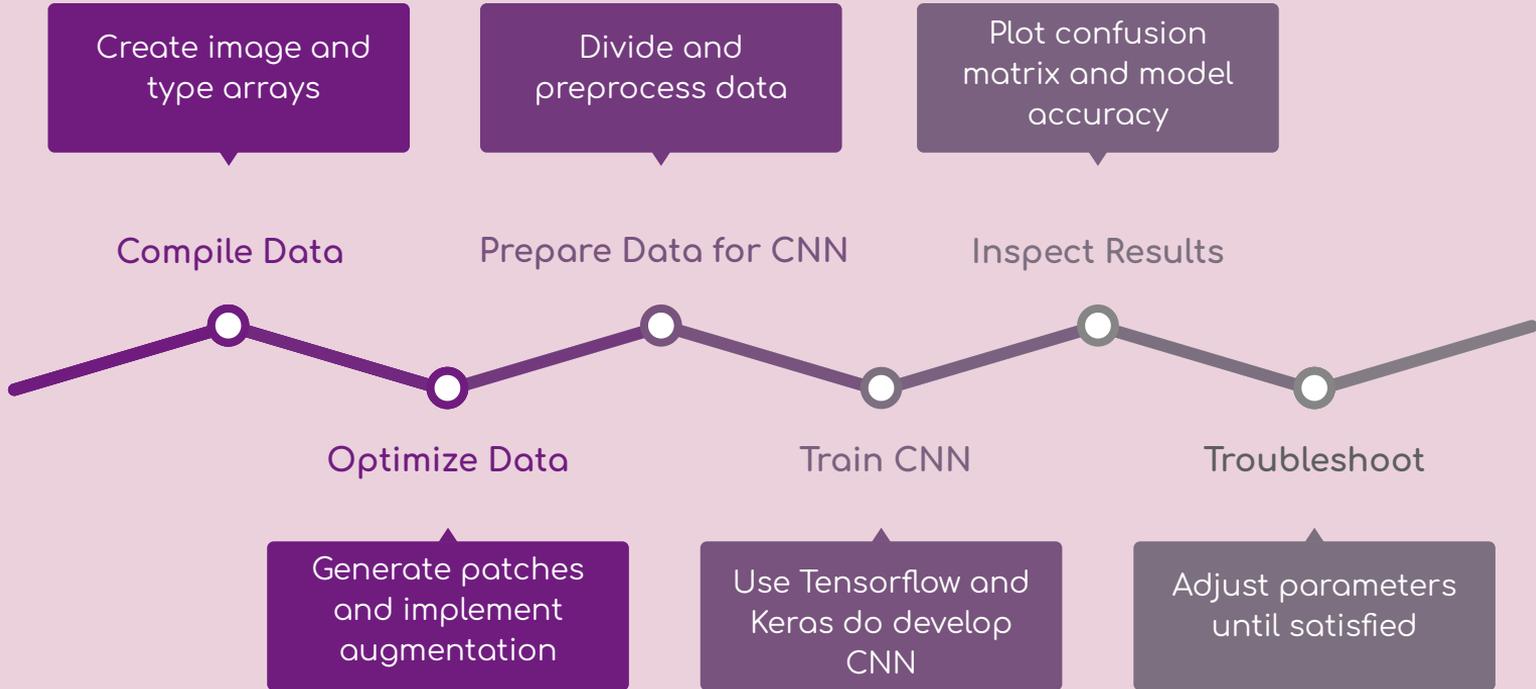
Why use a convolutional neural network?



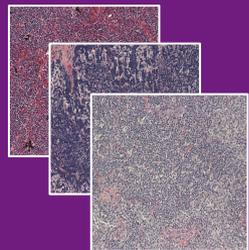
- Useful for processing images (as opposed to data sets)
- Robust and computationally efficient with large sample of data
- Does not require human supervision

Some common architectures: LeNet-5 (1998), AlexNet (2012), VGG-16 (2014)

Overview of Analysis Process



Concatenate Image Files Into Array



lymphoma_imgs

```
[[[122 90 114] [ 71 54 105]
 [122 90 115] [ 70 54 107]
 [120 96 113] [ 70 55 109]
 ...
 [102 64 78] ...
 [105 69 85] [141 142 194]
 [115 70 96] [120 120 109]
 ...
 [117 109 169]]]
 [[123 90 121]
 [122 90 119] [[ 66 49 99]
 [119 96 114] [ 66 49 103]
 ...
 [ 69 52 110]
 [101 64 78] ...
 [109 69 83] [110 112 169]
 [115 70 92] [108 93 147]
 ...
 [ 91 75 133]]]
 [[134 106 125]
 [131 106 125] [[ 56 41 88]
 [129 109 125] [ 56 41 93]
 ...
 [ 62 45 99]
 [106 71 80] ...
 [118 82 90] [ 90 80 135]
 ...
 [ 86 68 119]
 ...
 [ 84 62 106]]]]
```

Construct Image Types Array



CLL



FL



MCL

Integer One-Hot Encoding



0



[1 0 0]



1



[0 1 0]



2

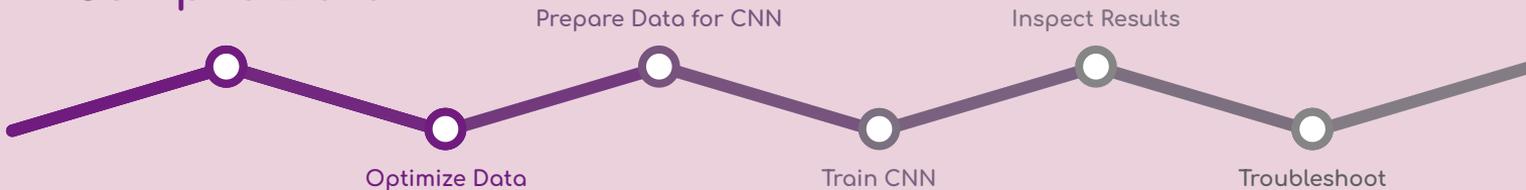


[0 0 1]

lymphoma_types

```
[[1. 0. 0.]
 [1. 0. 0.]
 [1. 0. 0.]
 ...
 [0. 0. 1.]
 [0. 0. 1.]
 [0. 0. 1.]
```

Compile Data



Optimize Data

Train CNN

Troubleshoot

Generate 32 x 32 Pixel Patches



Apply Image Augmentations



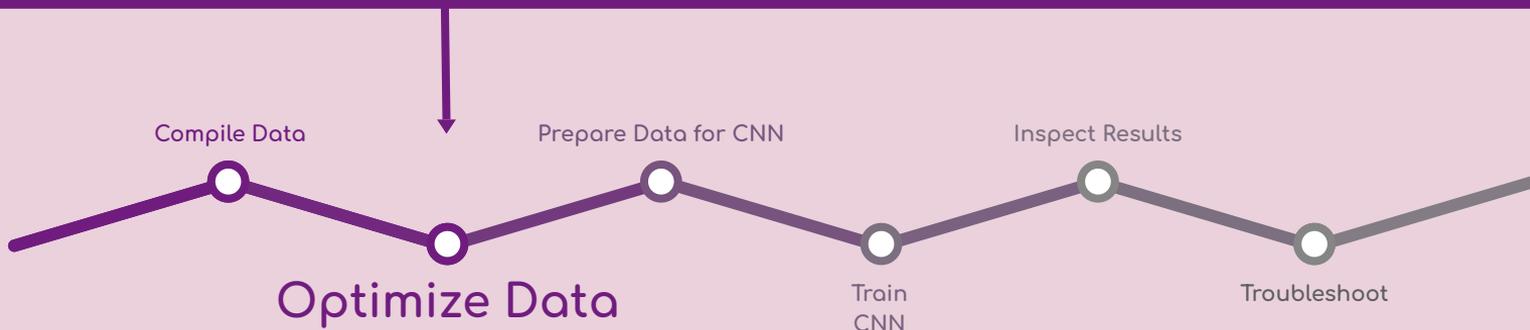
Horizontal Flip



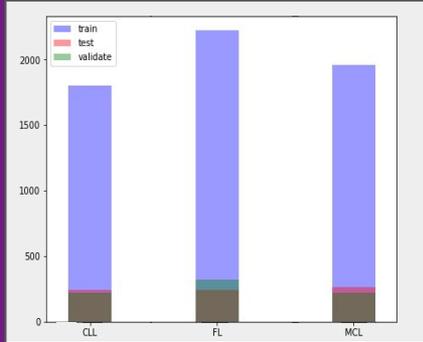
Vertical Flip



45° Rotation

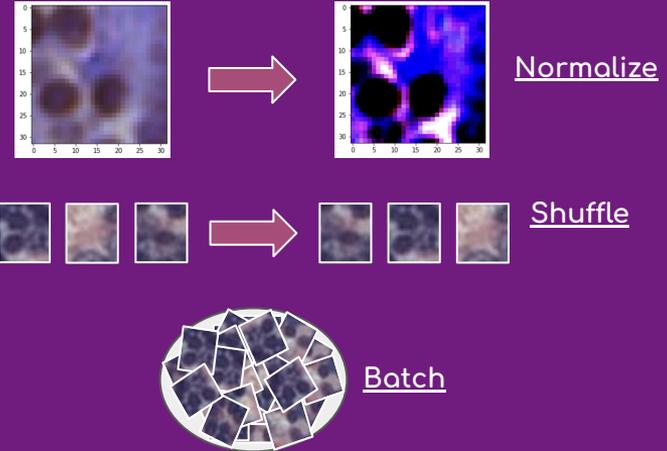


Divide Data Into Subsets



Training: 80% ; Validation: 10% ; Test: 10%

Preprocess Data



Prepare Data for CNN



Building Blocks of a CNN



Convolution

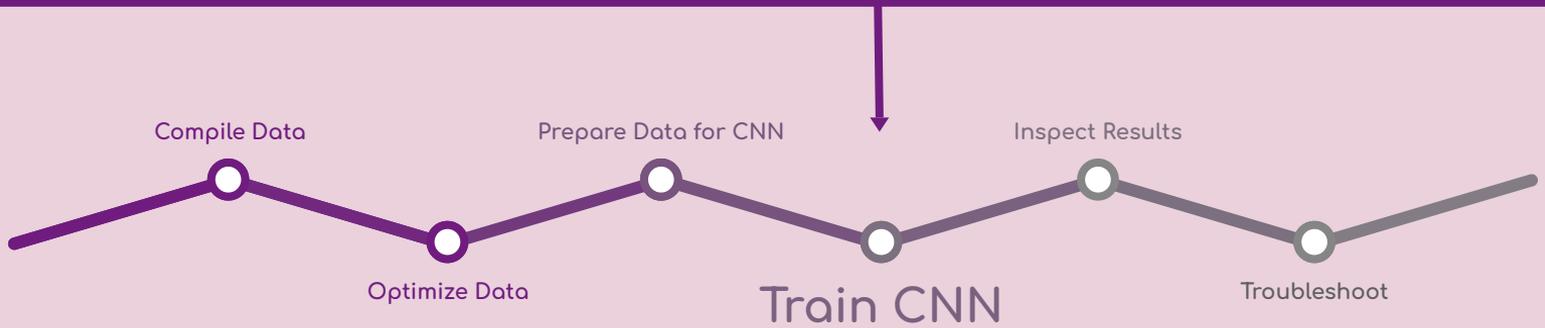
Activation

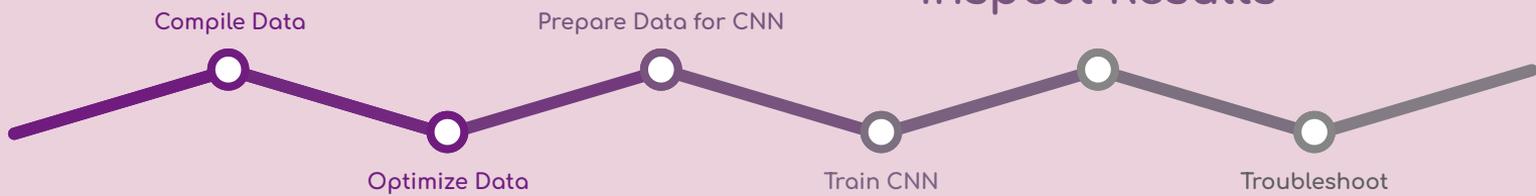
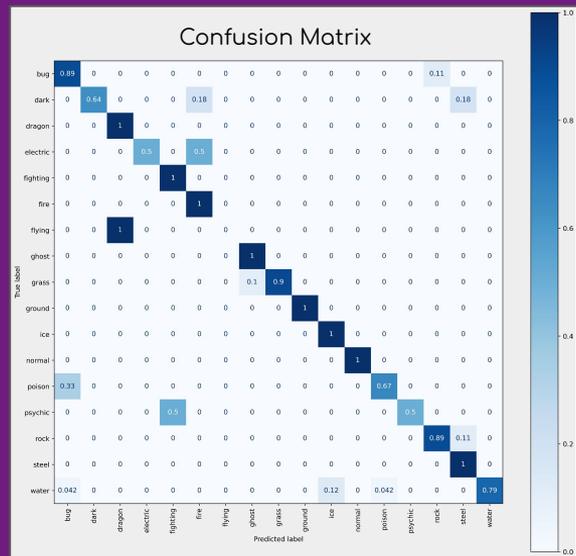
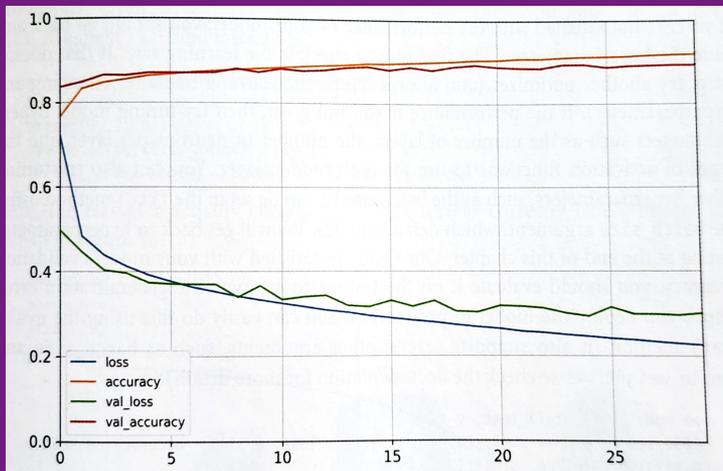
Max Pooling

Average Pooling

Dense

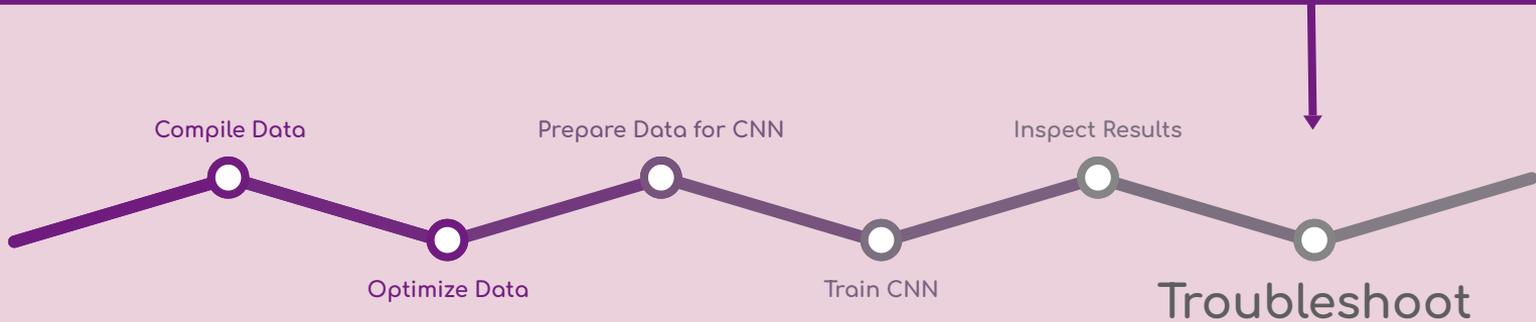
Flatten



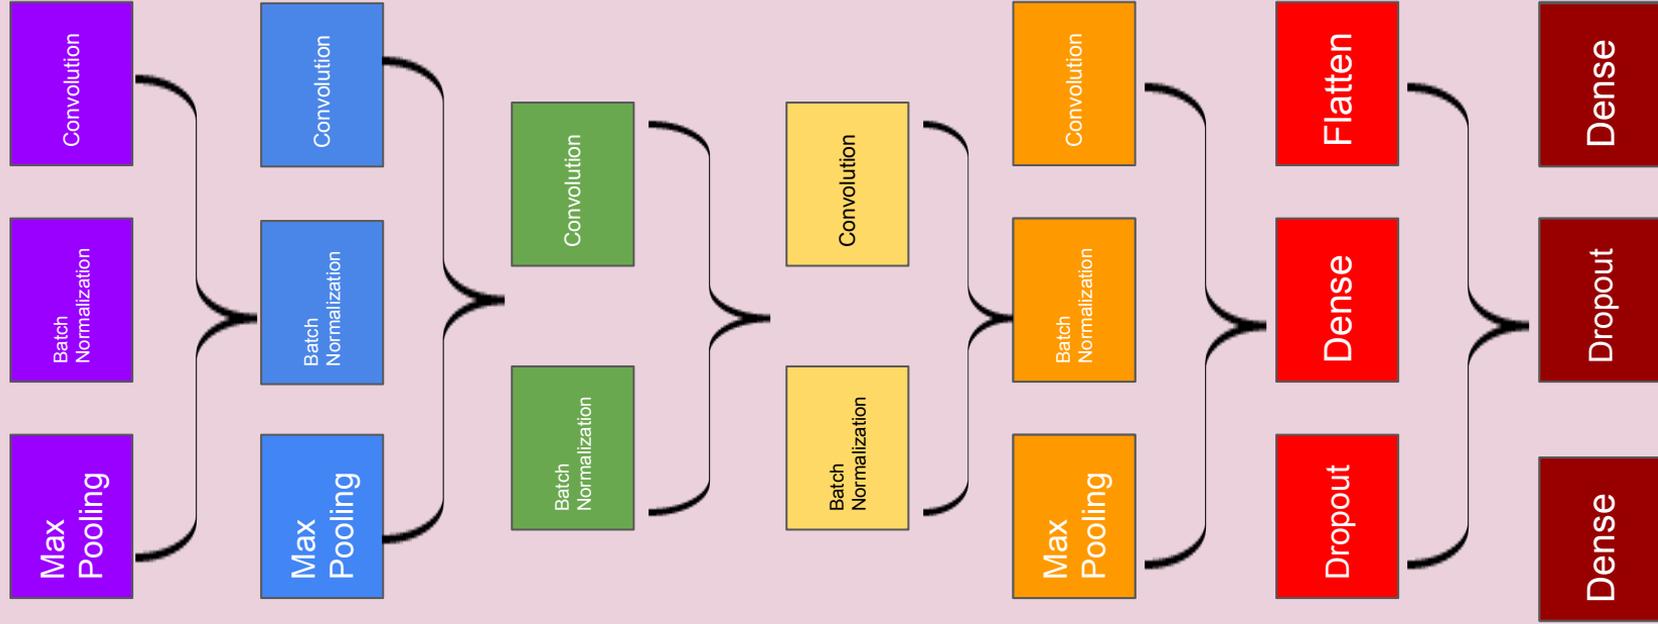


CNN Output

```
Epoch 72/150
8/8 [=====] - 5s 496ms/step - loss: 1.4068 - accuracy: 0.5348 - val_loss: 1.0313 - val_accuracy: 0.5321
Epoch 73/150
8/8 [=====] - 5s 577ms/step - loss: 1.4446 - accuracy: 0.5223 - val_loss: 1.0227 - val_accuracy: 0.5334
Epoch 74/150
8/8 [=====] - 5s 595ms/step - loss: 1.3753 - accuracy: 0.5307 - val_loss: 1.0302 - val_accuracy: 0.5374
Epoch 75/150
8/8 [=====] - 4s 492ms/step - loss: 1.3943 - accuracy: 0.5343 - val_loss: 1.0244 - val_accuracy: 0.5361
Epoch 76/150
8/8 [=====] - 4s 488ms/step - loss: 1.3497 - accuracy: 0.5365 - val_loss: 1.0122 - val_accuracy: 0.5428
Epoch 77/150
8/8 [=====] - 4s 492ms/step - loss: 1.3699 - accuracy: 0.5348 - val_loss: 0.9895 - val_accuracy: 0.5548
Epoch 78/150
8/8 [=====] - 4s 490ms/step - loss: 1.4060 - accuracy: 0.5285 - val_loss: 1.0254 - val_accuracy: 0.5388
Epoch 79/150
8/8 [=====] - 5s 493ms/step - loss: 1.3207 - accuracy: 0.5521 - val_loss: 1.0096 - val_accuracy: 0.5468
Epoch 80/150
8/8 [=====] - 4s 491ms/step - loss: 1.3166 - accuracy: 0.5521 - val_loss: 0.9780 - val_accuracy: 0.5602
Epoch 81/150
```



Extraterrestrial Confounding Cat Architecture



Results and Discussion

CNN Architectures

Convolution



Activation



Max Pooling



Average Pooling



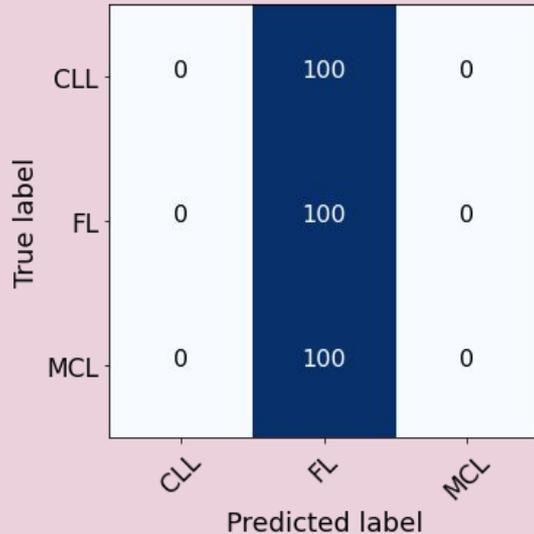
Dense



Flatten



First CNN implementation: goes wrong



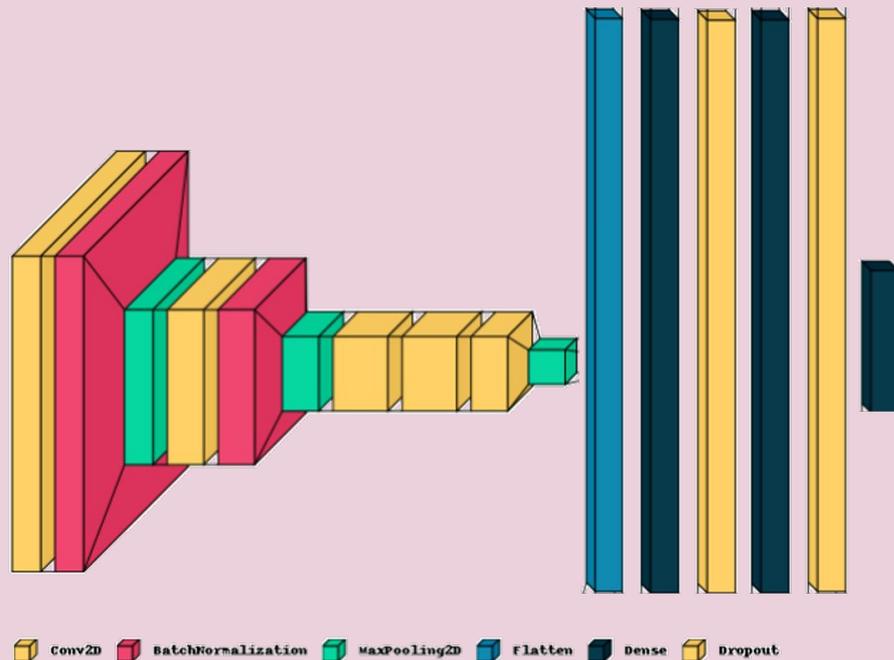
- Architecture adapted from morphological classification of SDSS galaxies activity
- A lot of changes, same performance



Solution: AlexNet CNN architecture

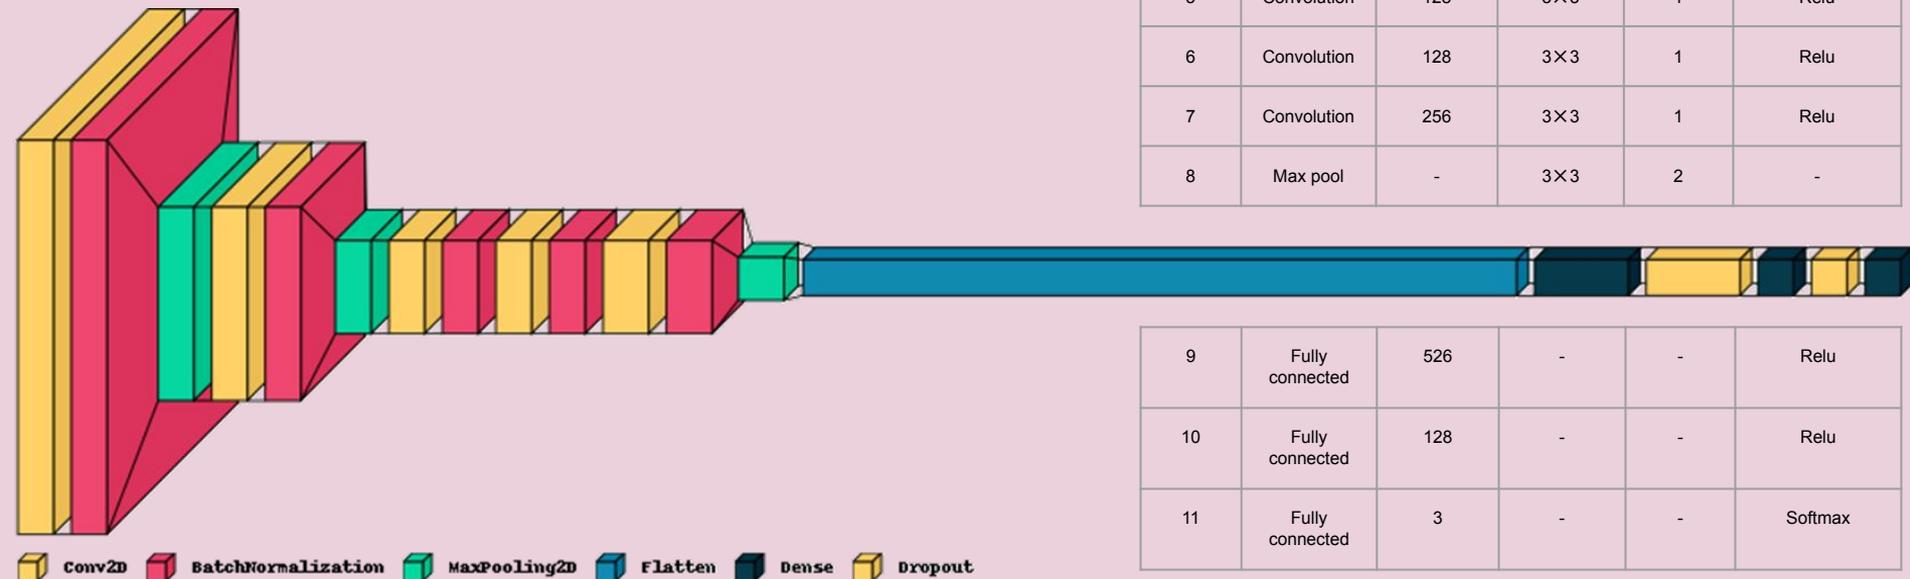
It consists of:

- 5 convolutional layers
- 3 max pooling layers
- 2 normalization layers
- 2 fully connected layers
- 1 softmax layer



Extraterrestrial Confounding Cat Architecture

- Slightly different of the original AlexNet architecture



Layer	Type	Num Kernels	Kernel Size	Stride	Activation
0	Input	3	227 × 227	-	-
1	Convolution	64	11 × 11	4	Relu
2	Max pool	-	3 × 3	2	-
3	Convolution	128	5 × 5	1	Relu
4	Max pool	-	3 × 3	2	-
5	Convolution	128	3 × 3	1	Relu
6	Convolution	128	3 × 3	1	Relu
7	Convolution	256	3 × 3	1	Relu
8	Max pool	-	3 × 3	2	-

9	Fully connected	526	-	-	Relu
10	Fully connected	128	-	-	Relu
11	Fully connected	3	-	-	Softmax

Results and Discussion (continued)

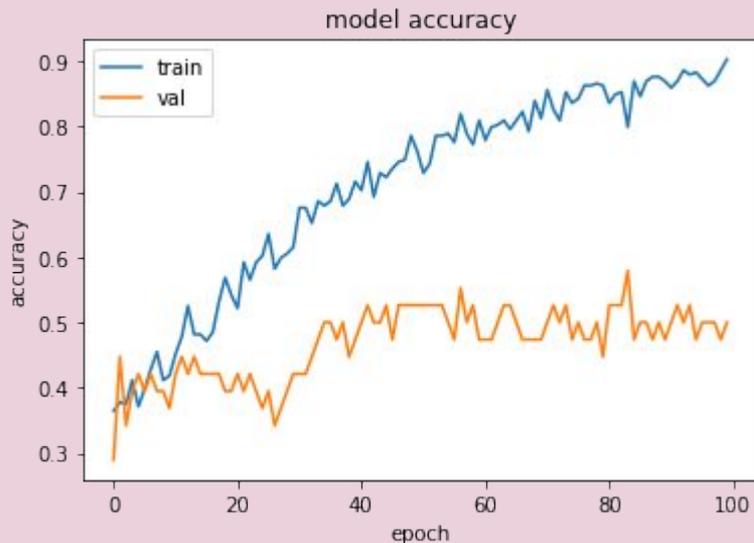
Learning curves - evaluate model performance

Confusion matrix - performance measurement for classification algorithms

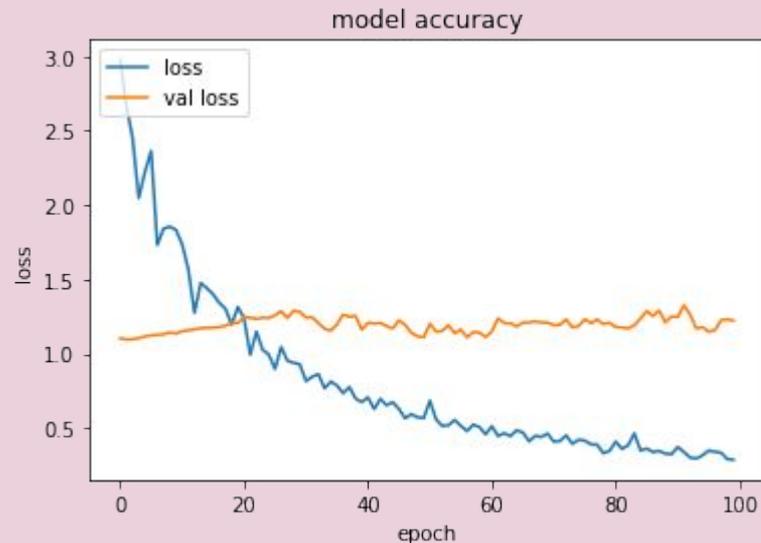
Simplest case: 1 patch per image

training	validation	test
299	37	38

- Training accuracy increases
- Validation accuracy constant



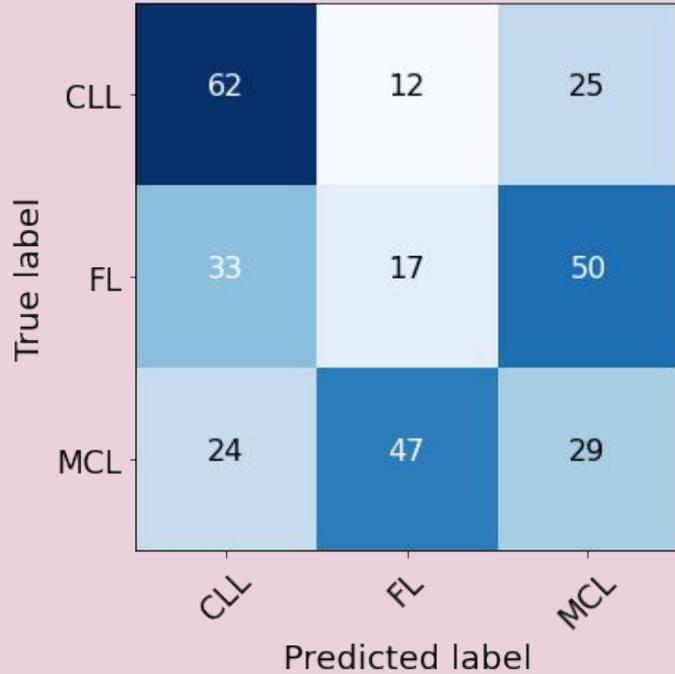
Accuracy curve over 100 epochs



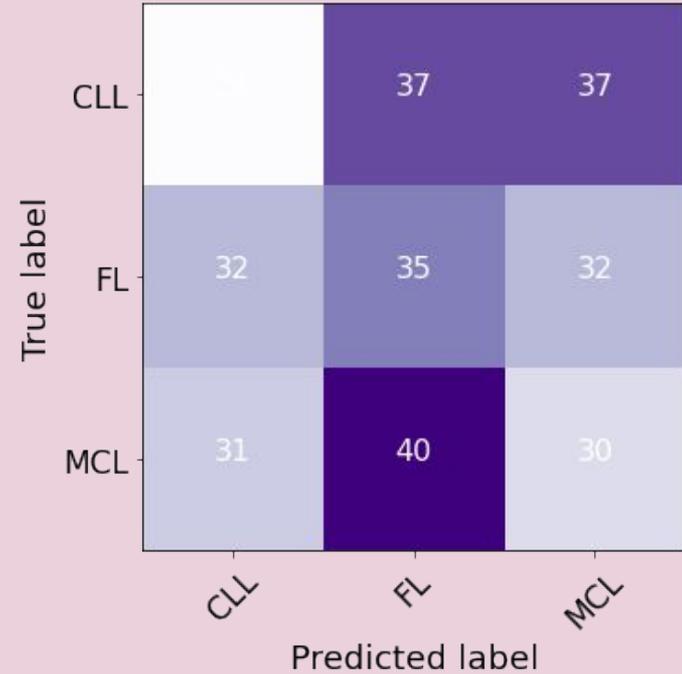
Loss curve over 100 epochs

Confusion matrices for test and training set

- Predictions for test set not very accurate

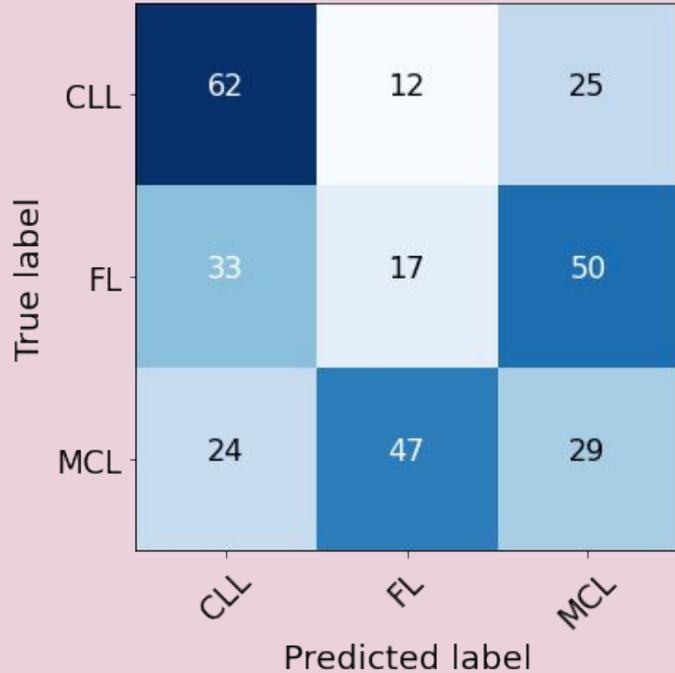


- Sanity check: predictions for training set

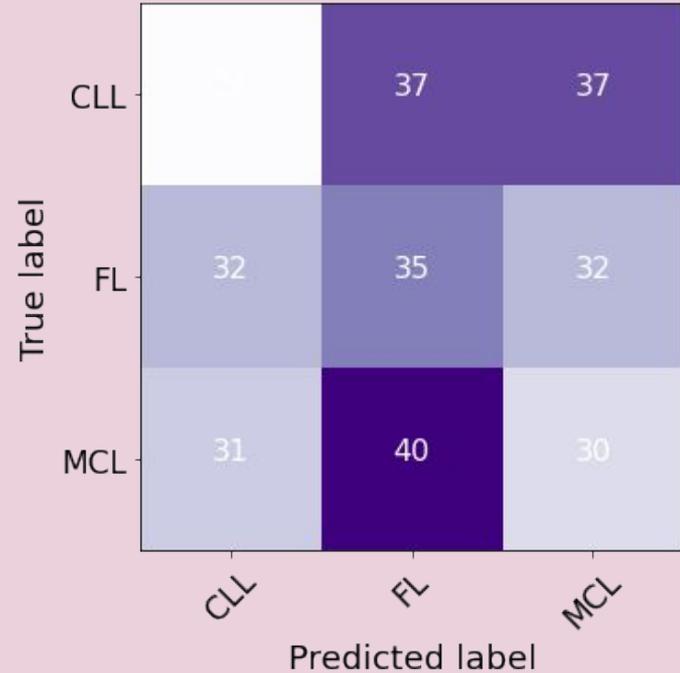


Is the model performance bad, or are other factors in play?

- Predictions for test set not very accurate



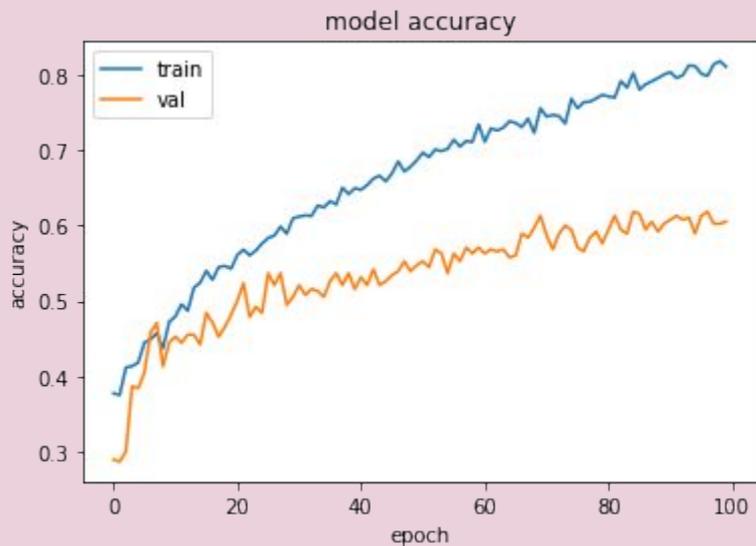
- Sanity check: predictions for training set



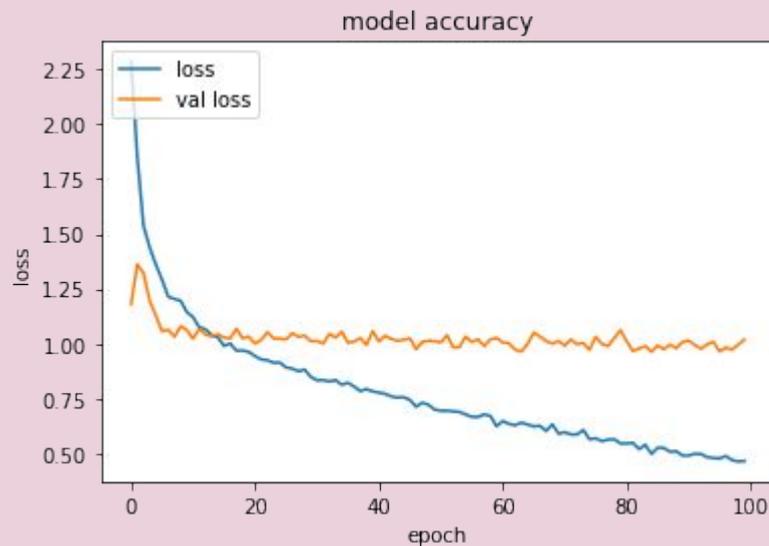
Increase sample size: 10 patch per image

training	validation	test
2990	370	380

Learning curves show indication of overfitting



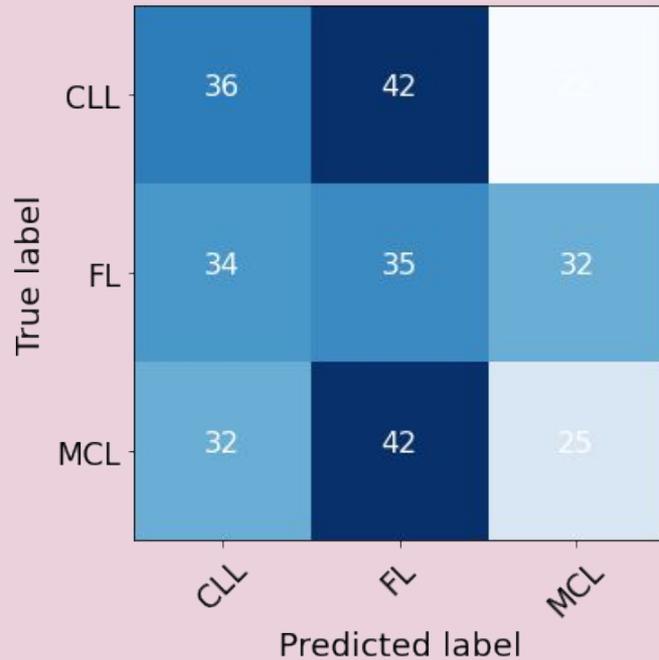
Accuracy curve over 100 epochs



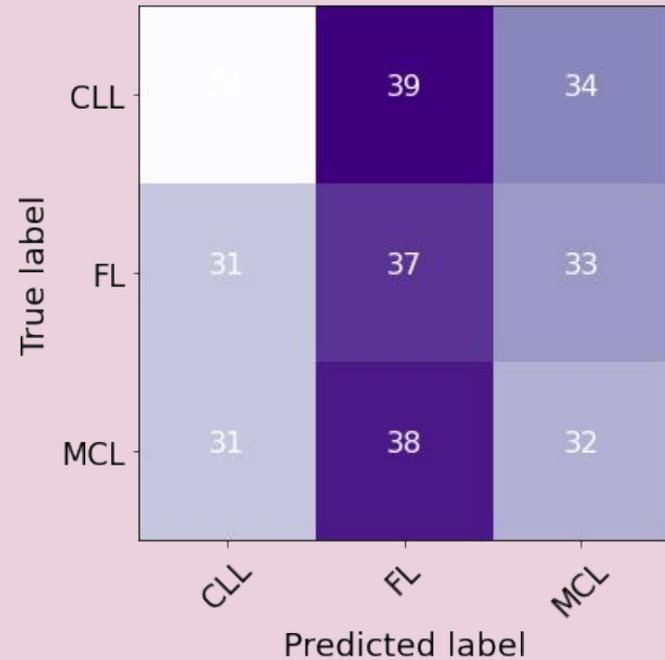
Loss curve over 100 epochs

Confusion matrices for test and training set

- Predictions for test set have not improved



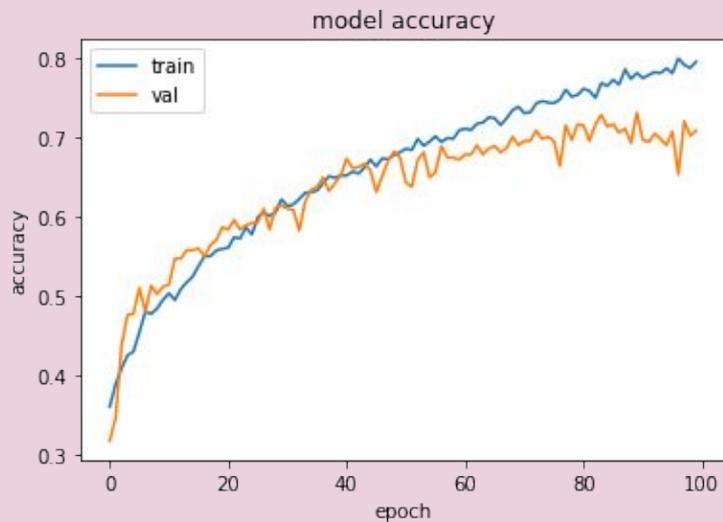
- Sanity check not so sane?



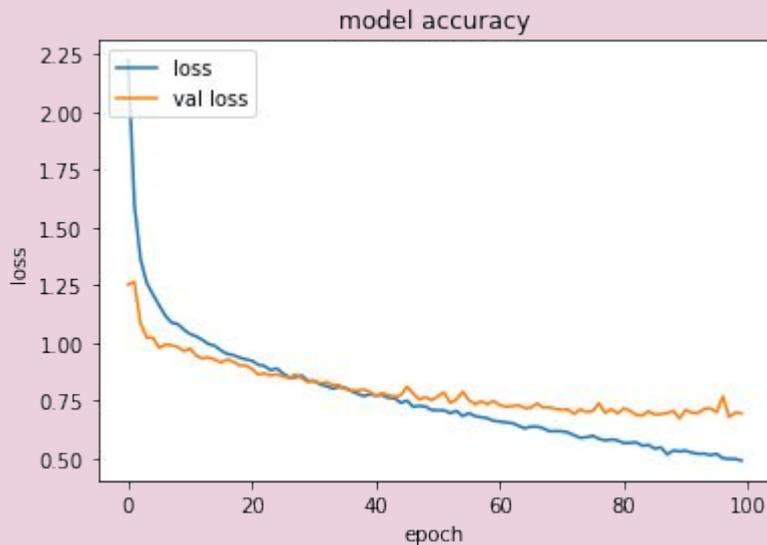
Increase sample size: 20 patch per image

training	validation	test
5980	740	760

- Validation accuracy curve shows slight improvement
- Training accuracy curve continues to rise



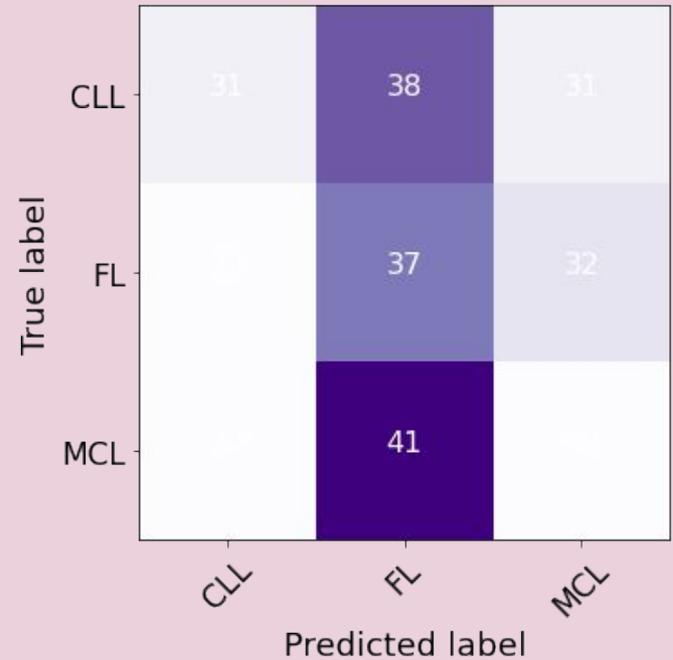
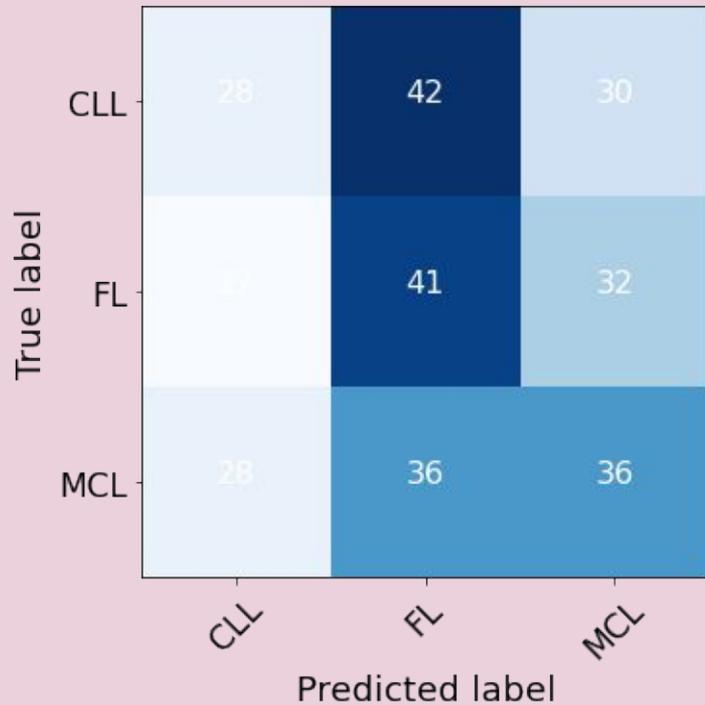
Accuracy curve over 100 epochs



Loss curve over 100 epochs

Confusion matrices for test and training set

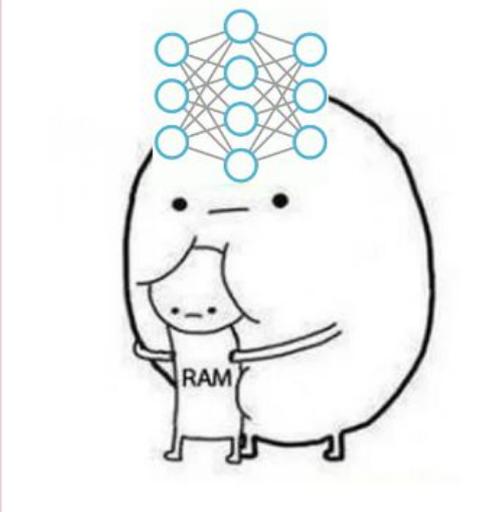
- Predictions for test set have not improved



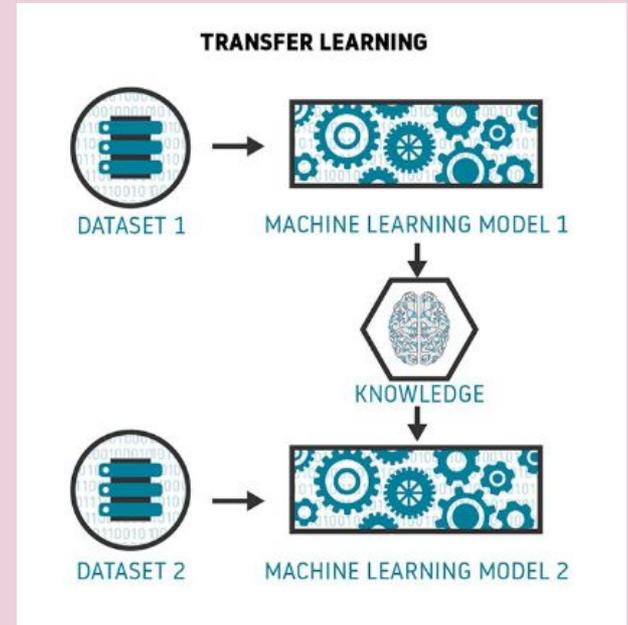
Tests with other parameters

- Increasing sample size by x20 → *no significant improvement*
- Decreasing the complexity of CNN → different architectures
- Dropout parameter → 0.1 - 0.5
- Image augmentation in training set → rotation - horizontal and vertical flipping - zooming
- Activation function, optimizer, kernels, learning rate
- Increasing batch size → *made things worse*

Optimization improvements for classification



- Better computing power - increase training sample without memory issues
- Implement AlexNet with transfer learning
- Segmentation/feature extraction



Connections to Astronomy

The techniques implemented in this work can also be used in astronomy!

- CNNs: very useful for image processing; applicable for limiting the necessity of visual inspections in photometric works such as Lyman Alpha studies.
- Data augmentation: useful for those cases in which we do not have many observations (example: rotating galaxies with particular features that we want to study).
- Nowadays surveys can serve as training samples.
- Catalogs can be used as seeds for transfer learning techniques (e.g. JWST)

...and to other fields!

Summary

- Created CNN using TensorFlow and Keras in order to classify 3 different types of Lymphoma
- Improved results significantly from initial methods, but ultimately could achieve better results with less limitations in RAM related to the homogeneity in images
- Gained an exciting experience of working out of area
- Learned many valuable methods in deep learning, data augmentation, and image processing, which we hope to apply to our research in astronomy

Thank you for your attention!

Room 2 Zoom Meeting

A screenshot of a Zoom meeting window titled "Room 2 Zoom Meeting". The window displays a grid of five video thumbnails. The top row contains two thumbnails: the left one shows a woman with glasses and a red hoodie, labeled "Elisa Tau"; the right one shows a woman with long brown hair, labeled "Nicole M Firestone (she/her)". The middle row contains two thumbnails: the left one shows a woman with glasses and a patterned top, labeled "Keerthana"; the right one shows a man with glasses and a black hoodie, labeled "Alejandro Cartes". The bottom row contains a single, larger thumbnail centered below the others, showing a man with a beard and glasses, labeled "Zachary Richards". The Zoom interface includes a green checkmark icon in the top-left corner and standard window controls (minimize, maximize, close) in the top-right corner.

Elisa Tau

Nicole M Firestone (she/her)

Keerthana

Alejandro Cartes

Zachary Richards

References

J Pathol Inform

Editor-in-Chief:
Anil V Parwani, Columbus, OH, USA
Loren Pantanowitz, Pittsburgh, PA, USA

OPEN ACCESS
HTML format

For entire Editorial Board visit : www.jpathinformatics.org/editorialboard.asp

Original Article

Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases

Andrew Janowczyk¹, Anant Madabhushi¹

¹Department of Biomedical Engineering, Case Western Reserve University, Cleveland, OH 44106, USA
E-mail: *Dr. Andrew Janowczyk - andrew.janowczyk@case.edu
*Corresponding author

Received: 18 November 2015 Accepted: 18 March 2016 Published: 26 Jul 2016

Abstract

Background: Deep learning (DL) is a representation learning approach ideally suited for image analysis challenges in digital pathology (DP). The variety of image analysis tasks in the context of DP includes detection and counting (e.g., mitotic events),

Histopathology

Histopathology 2018, 72, 227–238. DOI: 10.1111/his.13333

HER2 challenge contest: a detailed assessment of automated HER2 scoring algorithms in whole slide images of breast cancer tissues

Talha Qaiser,^{1,*} Abhik Mukherjee,^{2,*} Chaitanya Reddy PB,³ Sai D Munugoti,³ Vamsi Tallam,³ Tomi Pitkääaho,⁴ Taina Lehtimäki,⁴ Thomas Naughton,⁴ Matt Berseth,⁵ Anibal Pedraza,⁶ Ramakrishnan Mukundan,⁷ Matthew Smith,⁸ Abhir Bhalerao,¹ Erik Rodner,⁹ Marcel Simon,⁹ Joachim Denzler,⁹ Chao-Hui Huang,^{10,11} Gloria Bueno,⁶ David Snead,^{1,2} Ian O Ellis,² Mohammad Ilyas^{2,13} & Nasir Rajpoot^{1,12}

¹Department of Computer Science, University of Warwick, Coventry, UK, ²Department of Histopathology, Division of Cancer and Stem Cells, School of Medicine, University of Nottingham, Nottingham, UK, ³Department of Electronics and Electrical Engineering, Indian Institute of Technology, Guwahati, India, ⁴Department of Computer Science, Maynooth

NOTES

Started to segmentation

Slide of failed approaches

Explain Alexnet and our architecture

Plots

Talk about one hot encoder

Image normalization

What did work

Diagram of steps in analysis - arrays, preprocessing, training etc (Flow chart)

Why using neural networks

Future works - application to astronomy