CS 181AI
Lecture 5

Training cont. + Object Detectors

Arthi Padmanabhan

Feb 1, 2023

Logistics

- Assignment 2 is out

Today's Plan

- Group B leads reading discussion
- Finish discussion about debugging low accuracy when training a neural network
- Object Detection models

Course Path: The First Half

Intro to ML models and hands-on experience working with them

- Run a pre-trained model over your own images (1/23)
- Learn about training process and train a small model on a common dataset (1/25)
- Debug common issues with low accuracy when training (1/30)
- Learn about different types of models and their training processes (2/1)

Ethical AI

- Assess goals in building models
- Assess methods of data collection
- Study how various types of bias affect model results
- Study real use-cases

Computational Resources to Run and Train Models

- Learn about properties of machines that are particularly good for ML
- Study state-of-the-art ML devices + some promising future computing methods (e.g., quantum)
- Learn about how using multiple machines can enable scaling of ML tasks
- Study methods for lowering computational resources needed

Other Resources to Run and Train Models: Memory, Energy

- Compare memory and energy usage of running and training various types of models
- Study strategies for lowering memory and energy usage
- Case study: ML for monitoring agriculture

Improving Model Accuracy

- Common culprits:
  - Data
  - Model
  - Training Process

Labeling Example

- How many **entities** are in the following sentence?

Darth Sidious, known simply as the Emperor, was a Dark Lord of the Sith
who reigned over the galaxy as Galactic Emperor of the First Galactic Empire.

Labeling

- More data isn’t always better

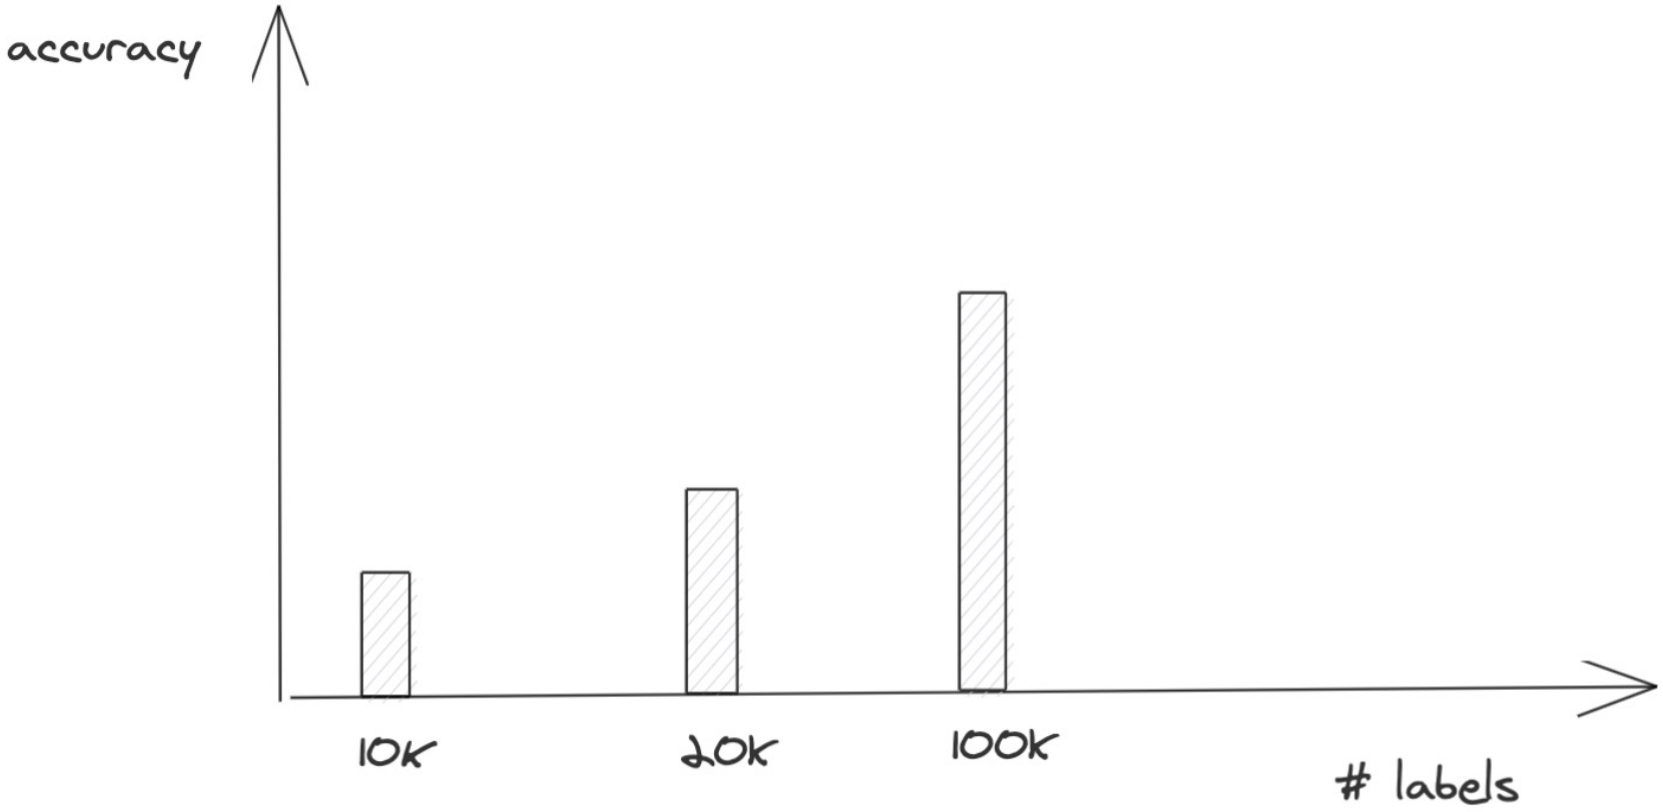

Labeling

- More data isn’t always better

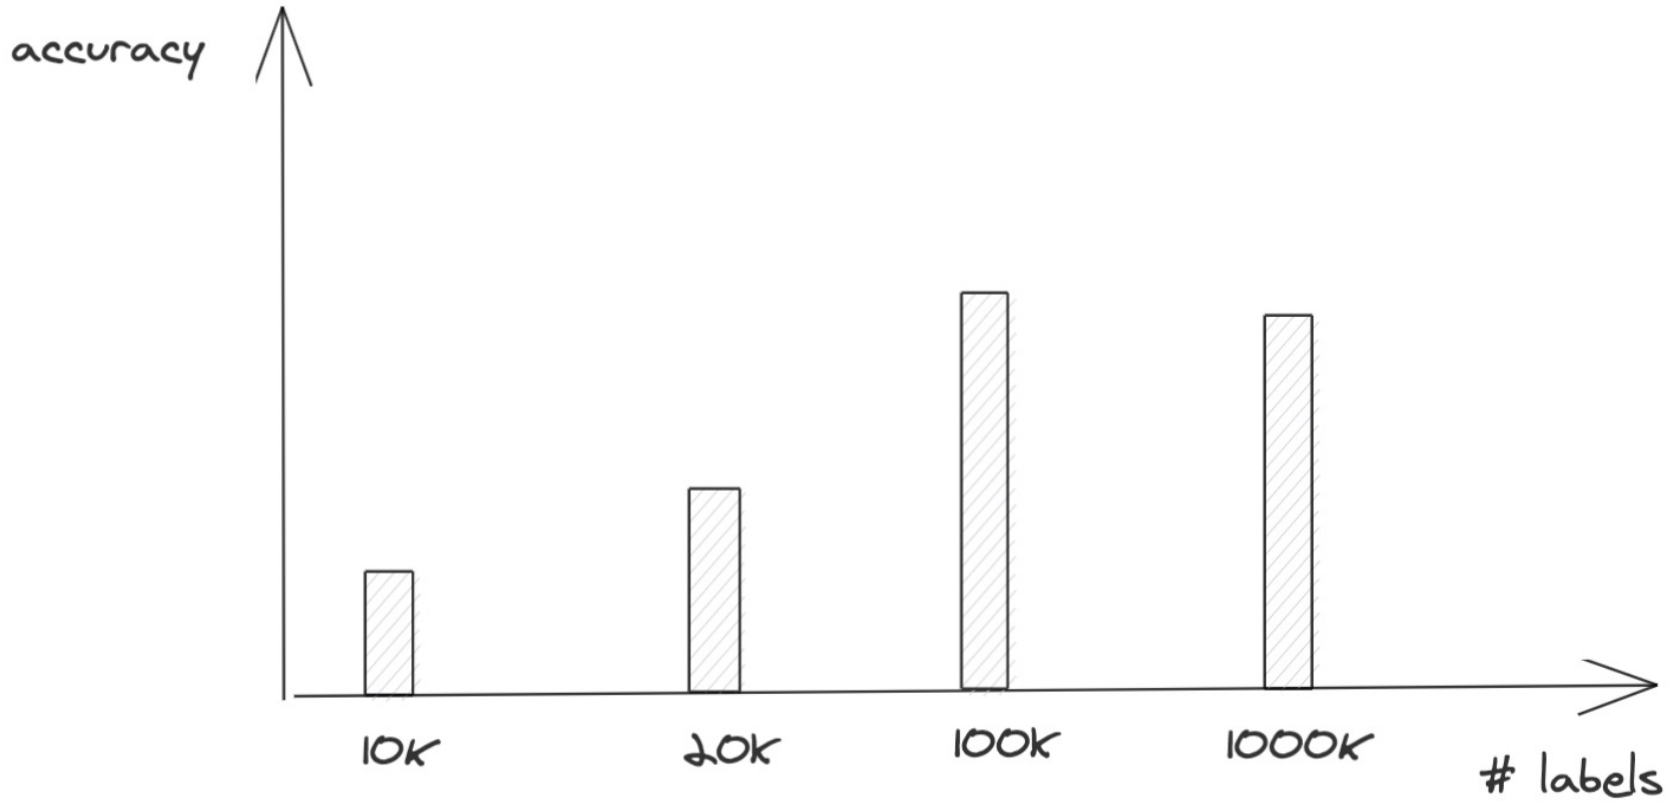

Labeling

- Programmatic Labeling can be messy but much faster and more adaptive
- Semi-supervision, or self-training: start with initial set of labels and train on new samples - only include the ones with a high score
- Active learning – only label those that model is uncertain about

Class Imbalance

- Which model do you hope your bank uses?

Model A	Actual Fraud	Actual Normal
Predicted Fraud	10	10
Predicted Normal	90	890

Model B	Actual Fraud	Actual Normal
Predicted Fraud	90	90
Predicted Normal	10	810

Class Imbalance

- Which model do you hope your bank uses?
  - Just going by accuracy, they both have an accuracy of 90%

Model A	Actual Fraud	Actual Normal
Predicted Fraud	10	10
Predicted Normal	90	890

Model B	Actual Fraud	Actual Normal
Predicted Fraud	90	90
Predicted Normal	10	810

Class Imbalance

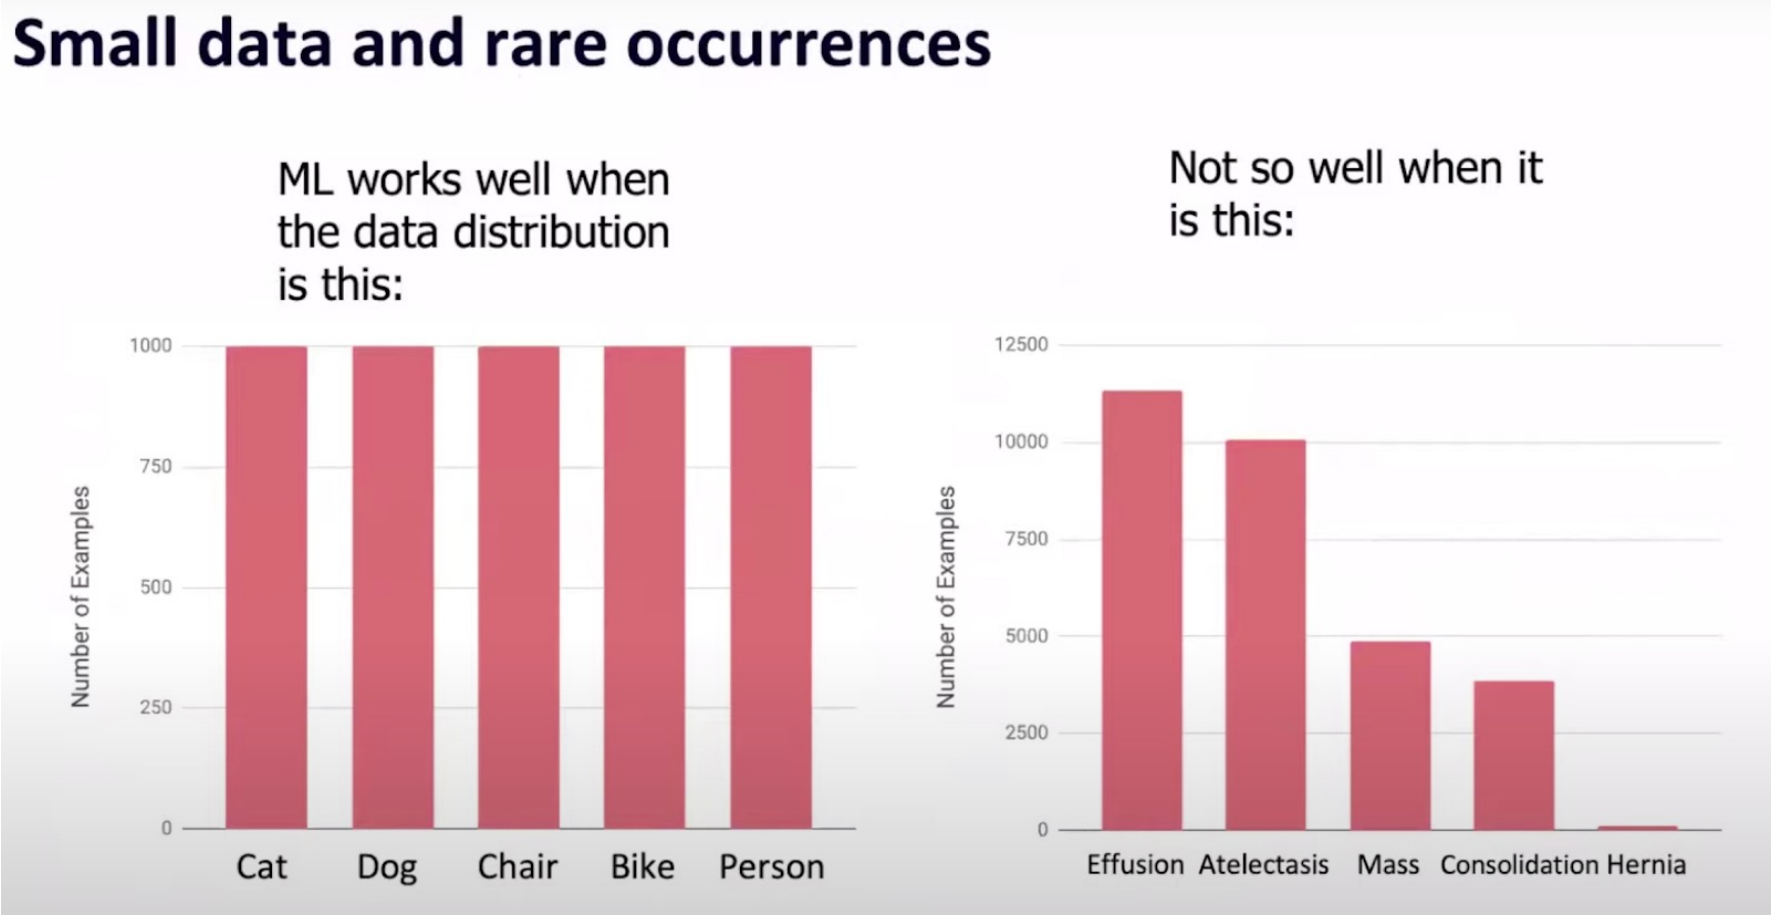

Class Imbalance

- Statistically, predicting the majority class has a higher chance of being right
- Asymmetric cost of error: different cost of wrong prediction

Class Imbalance

- Overall accuracy is most commonly used, but it's insufficient when there's class imbalance because it treats all classes the same
- F1: 2*(precision)*(recall) / (precision + recall)

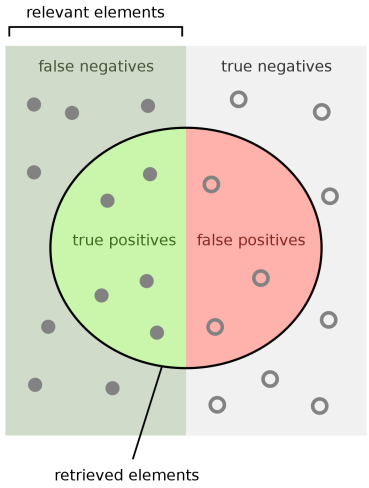

Class Imbalance

- Precision, Recall, and F1?

Model A	Actual Fraud	Actual Normal
Predicted Fraud	10	10
Predicted Normal	90	890

Model B	Actual Fraud	Actual Normal
Predicted Fraud	90	90
Predicted Normal	10	810

Class Imbalance

- Model A:
  - Precision = 50% (10/20)
  - Recall = 10% (10/100)
  - F1 = 17%
- Model B:
  - Precision = 50% (90/180)
  - Recall = 90% (90/100)
  - F1 = 64%

Model A	Actual Fraud	Actual Normal
Predicted Fraud	10	10
Predicted Normal	90	890

Model B	Actual Fraud	Actual Normal
Predicted Fraud	90	90
Predicted Normal	10	810

Not Enough Data

- Transfer Learning

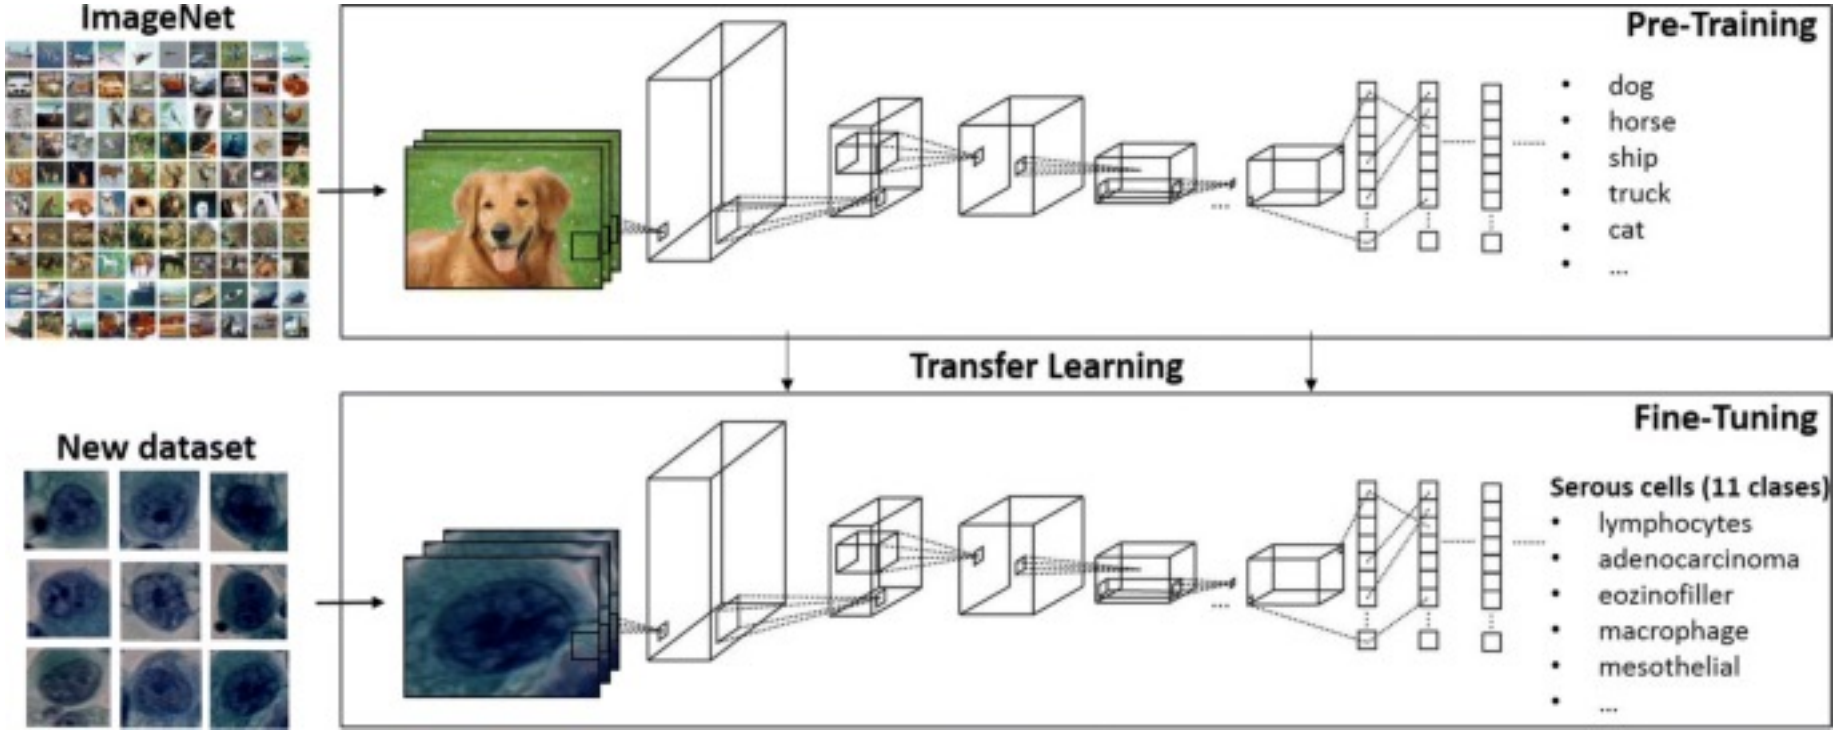

To Freeze or Not to Freeze

- While fine-tuning, ”freezing” some layers can make training go faster at the expense of lower accuracy

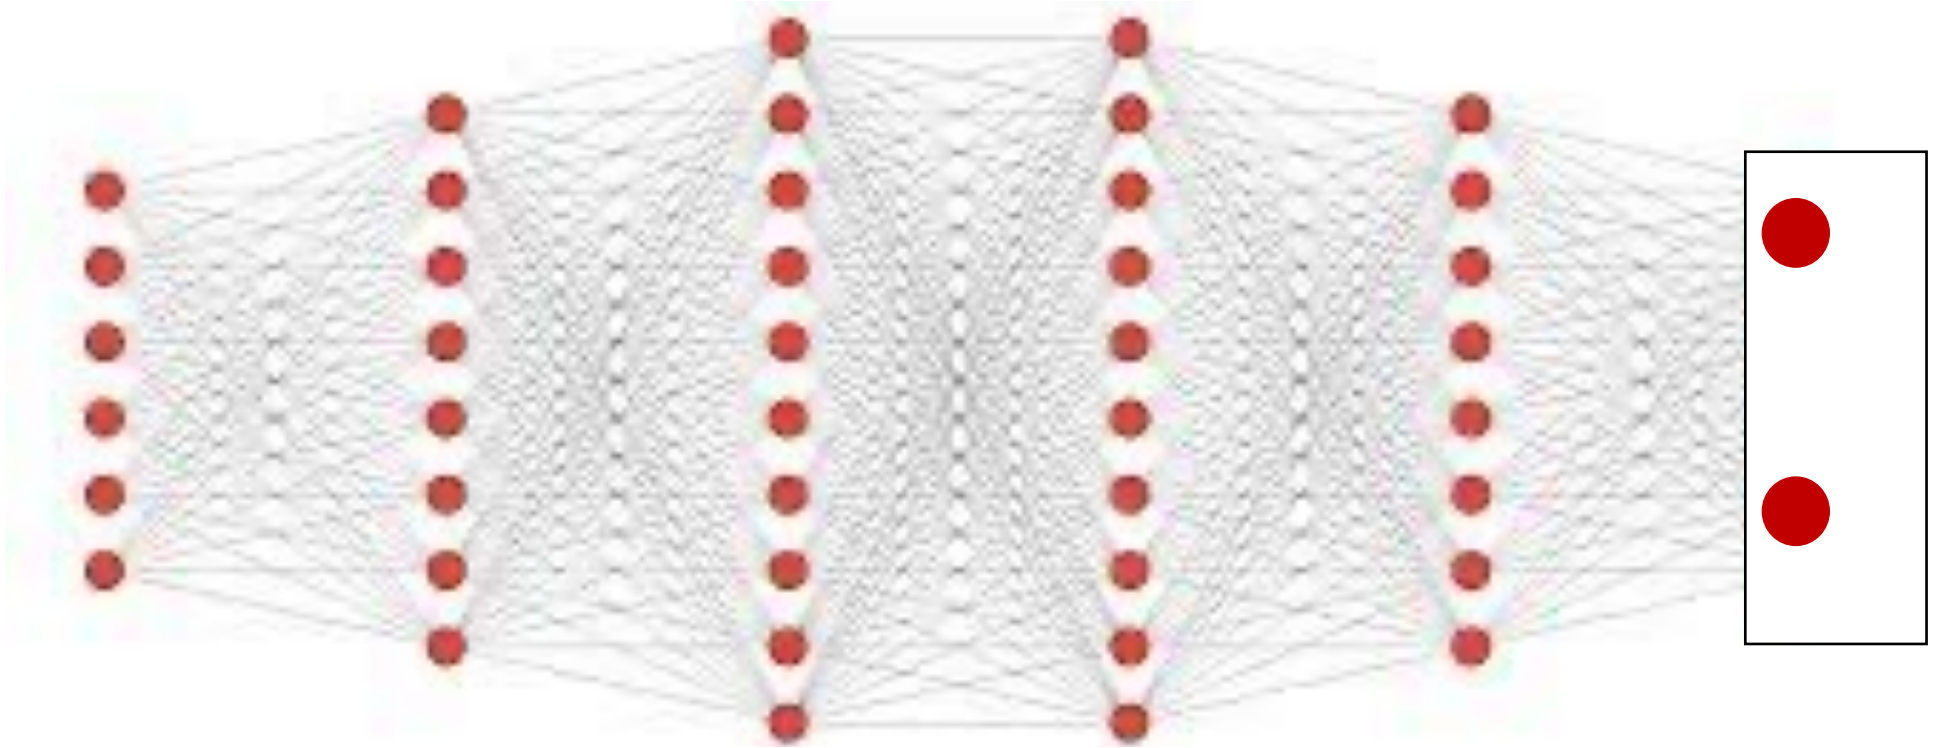

Improving Model Accuracy

- Common culprits:
  - Data
  - **Model**
  - Training Process

Models

- Layers -> generally more will give higher accuracy
- Using a well-established model will generally give higher accuracy – both because structure has been tested and because weights from pre-training are often available
- Details of how model structure (i.e., different layer types, order) affects accuracy is beyond scope of this course

Improving Model Accuracy

- Common culprits:
  - Data
  - Model
  - **Training Process**

Training Process

- Learning rate
- Optimizer
- Common mistakes: eval() vs. train(), didn’t zero the gradients, didn’t pass the correct model’s parameters to the optimizer

Object Detection

- Need to identify where objects are AND what they are

Faster RCNN

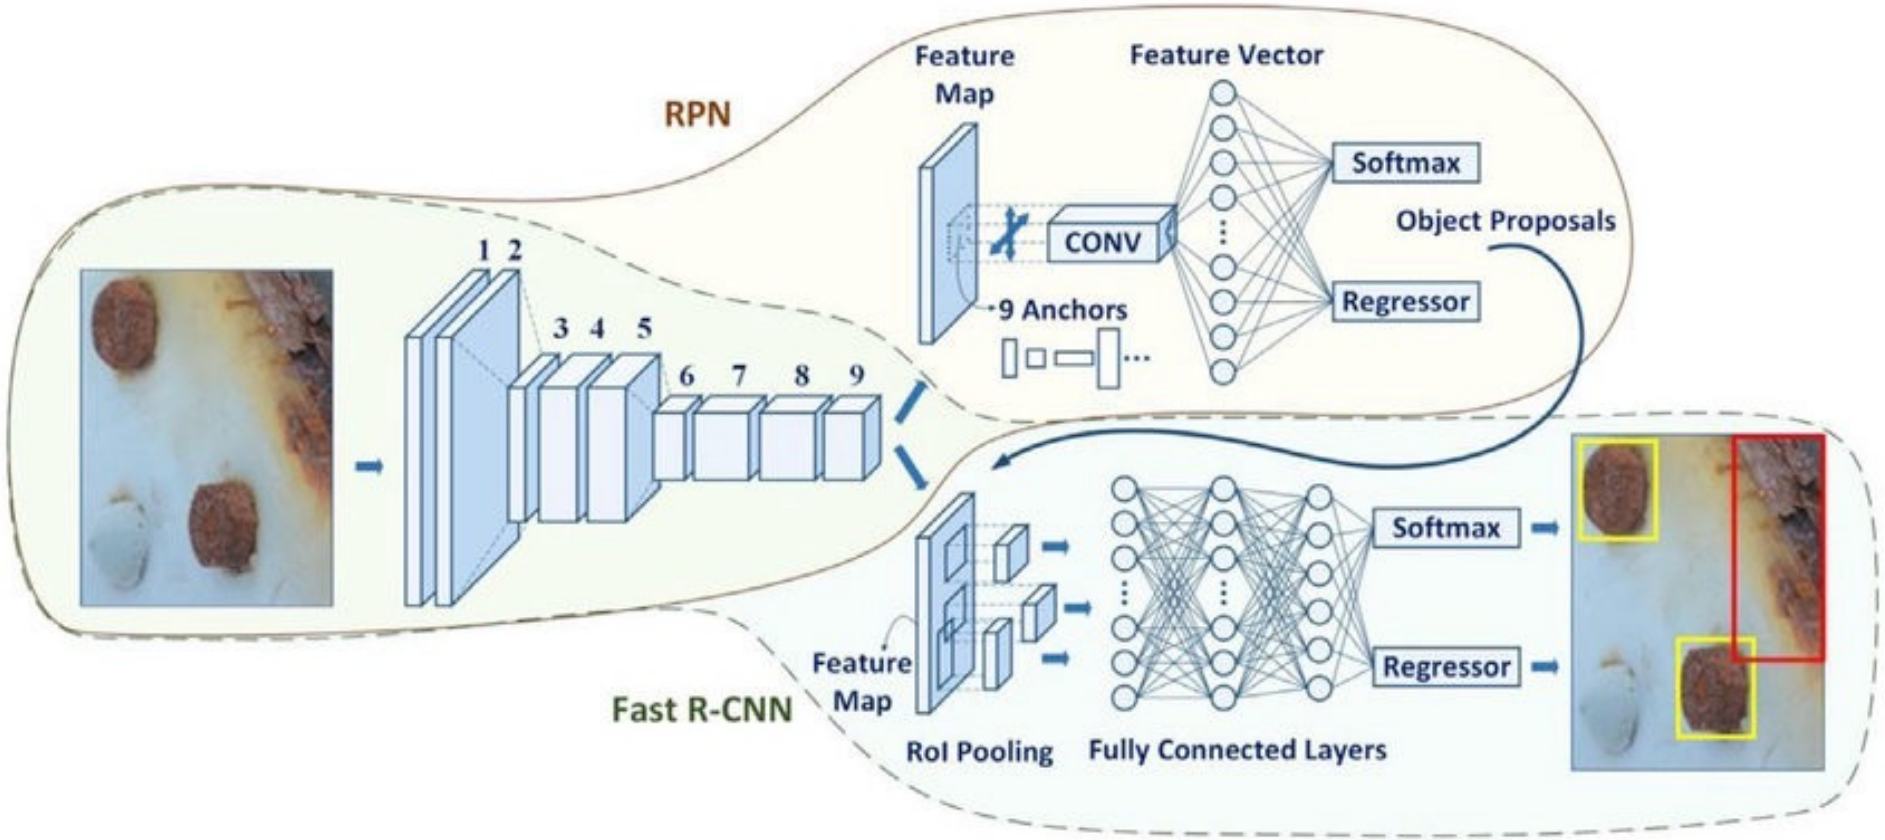
RPN
Feature Map
Feature Vector
CONV
Softmax
Regressor
Object Proposals
9 Anchors
1 2
3 4 5
6 7 8 9
Fast R-CNN
Feature Map
RoI Pooling
Fully Connected Layers
Softmax
Regressor

SSD (Single Shot Detector)

- Alternatively, these could be done together
- Divide image to cells. At each cell, 4 bounding boxes/object pairs are predicted. The object could be None

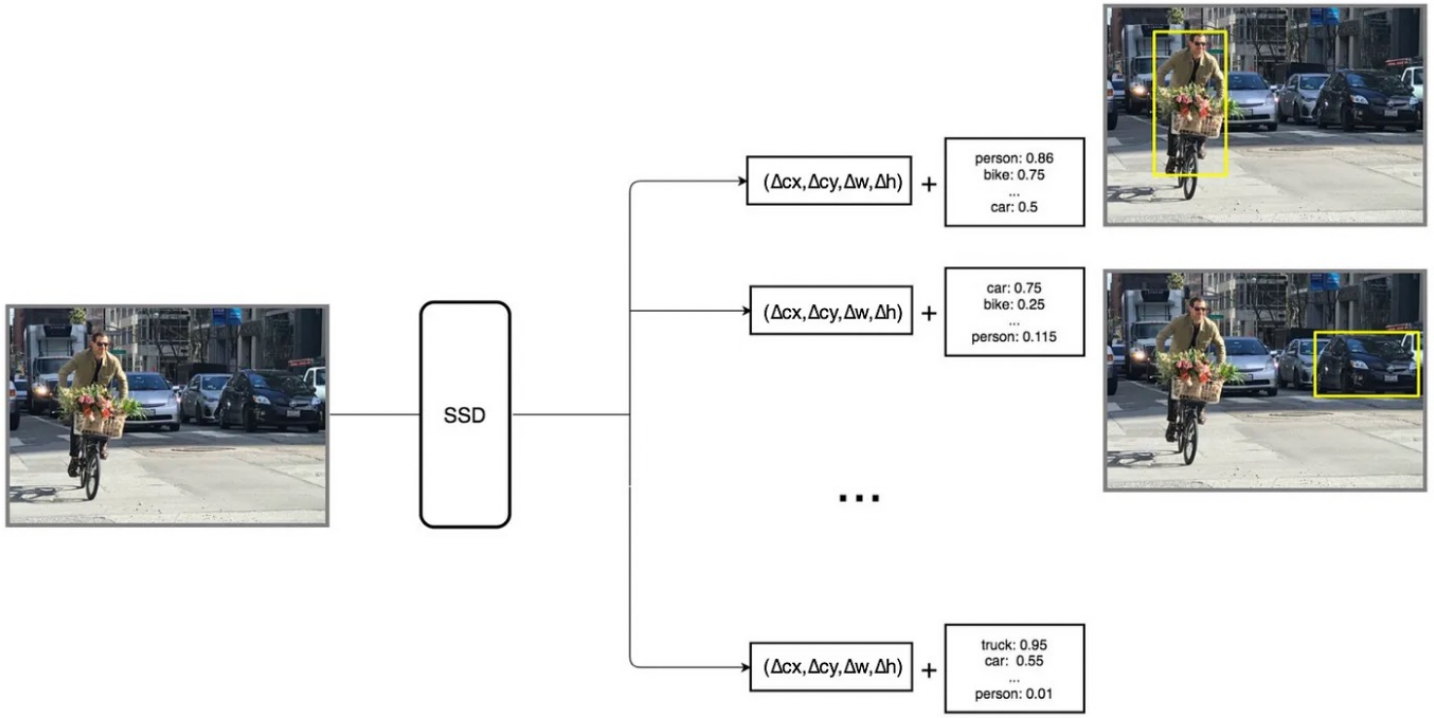

Cells of Different Sizes

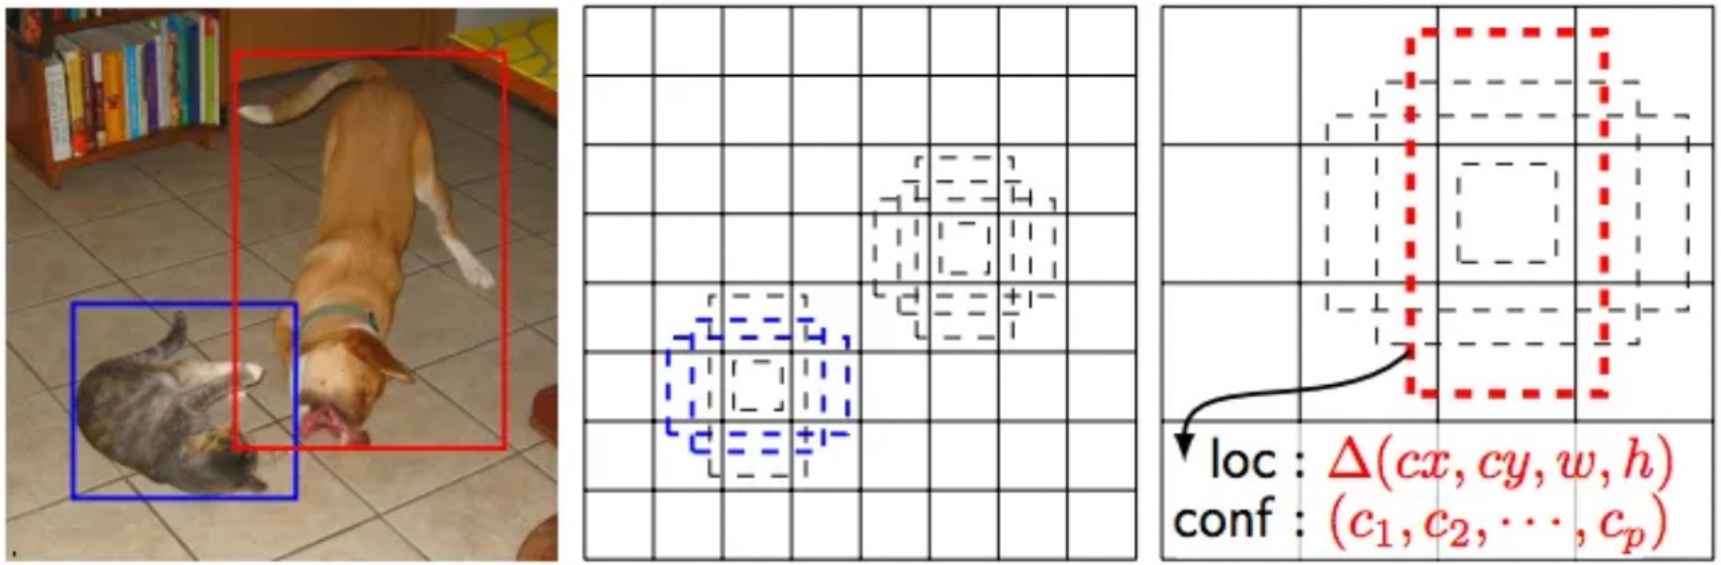

(a) Image with GT boxes (b) 8×8 feature map (c) 4×4 feature map