CS 5043: Advanced Machine Learning: Deep Learning

Andrew H. Fagg

Symbiotic Computing Laboratory

What is Machine Learning?

What is Machine Learning?

- Fundamentally: using data to automatically construct a model
- The model must be predictive!
  - I.E.: to be useful, it must produce meaningful output given novel situations.

Brain-Machine Interfaces

Command
prosthetic arm

Estimate of
intended
movement

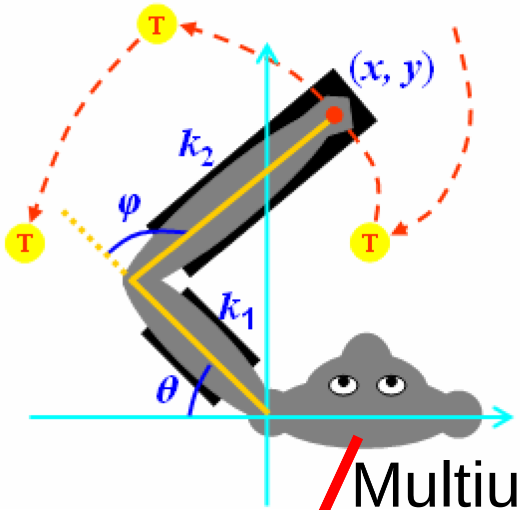

Multiunit
recording

Predictive
model

In collaboration with Nicholas G. Hatsopoulos and Lee E. Miller

Decoding Arm State

Want to predict arm motion at time t given recent history of spiking behavior

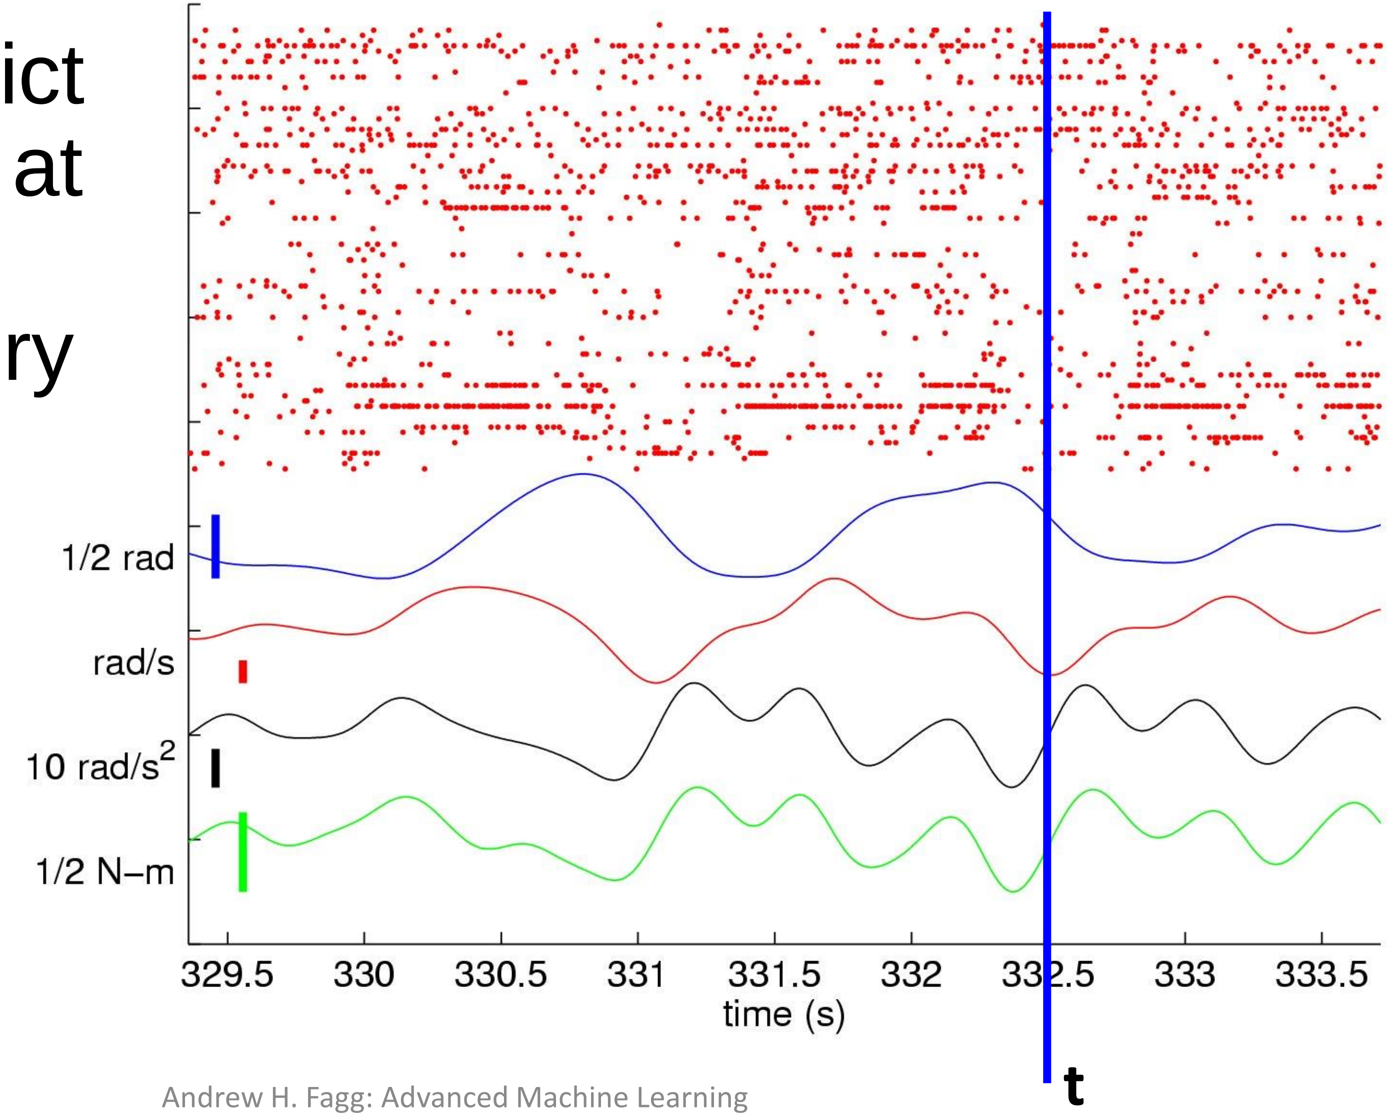

Decoding Arm State

- 50ms bins: 20 descriptors of neural activation for each cell

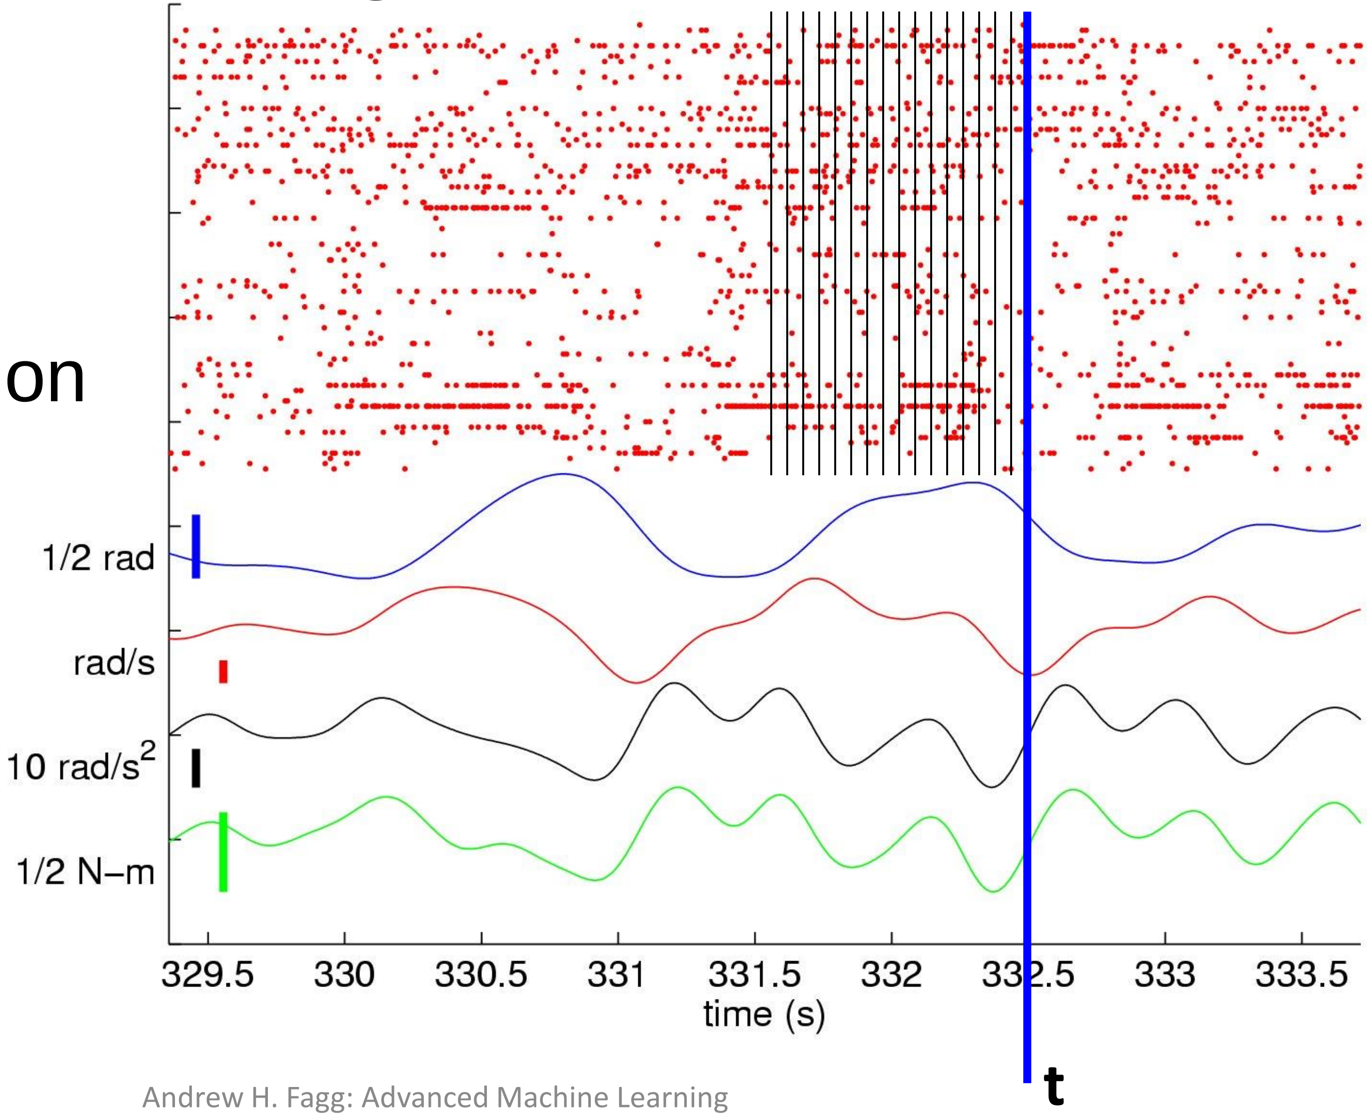

Wiener Filter

Each feature (F_i) is a count of spikes by a neuron for a 50 ms bin

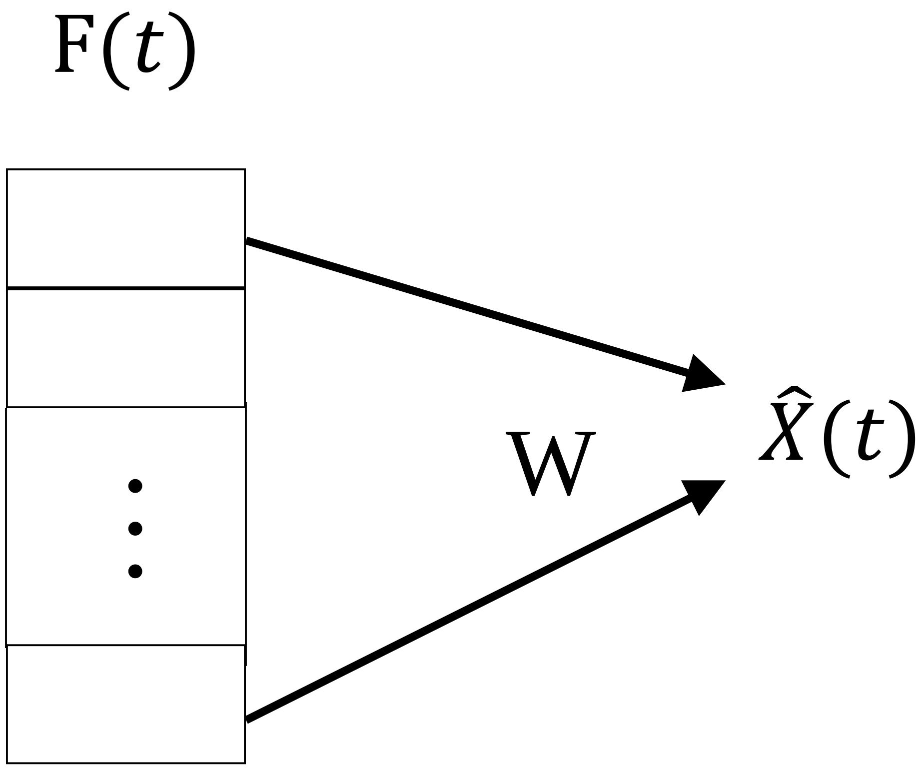

$$\hat{X} = g_W\big(F(t)\big) = W^T F(t)$$

Column vector encoding spike counts for N cells at T taps up to time t

Classes of Models

Defined by the data type of the output. Very broadly:

- Continuous output: regression-type models
- Categorical output: classifier models

Classes of Machine Learning Problems

Classes of Machine Learning Problems

Supervised learning

- Training set contains only input / output (labels) pairs
- Outputs could be continuous, probabilistic or categorical

Classes of Machine Learning Problems

Semi-Supervised learning

- Part of the training set contains input / output pairs
- The rest of the training set contains only inputs

- Using all of the data can yield a better model than if we only used the labeled data

Classes of Machine Learning Problems

Unsupervised learning

- The training set contains only inputs
- Fundamental question: what is the structure of these inputs?
  - Acommon case: algorithm assigns categorical labels to each sample
  - But we can also ask continuous questions. For example: are there linear or nonlinear manifolds that the data live on?

Classes of Machine Learning Problems

Reinforcement learning:

- Different than direct prediction or classification: RL is about taking sequences of actions in some environment
- At each step:
  - In response to an input, the model (agent) produces some action
  - The feedback signal is an evaluation of the results of this and previous actions

Classes of Machine Learning Problems

Reinforcement learning:

- A common case: a single evaluation can be a function of the sequence of outputs that is generated
  - How much time did it take to solve a problem?
  - How much energy did you use while solving the problem?
- Learning problem: for a given input, what is the output that maximizes the expected reinforcement over time?

Neural Networks…

Biological Neural Networks

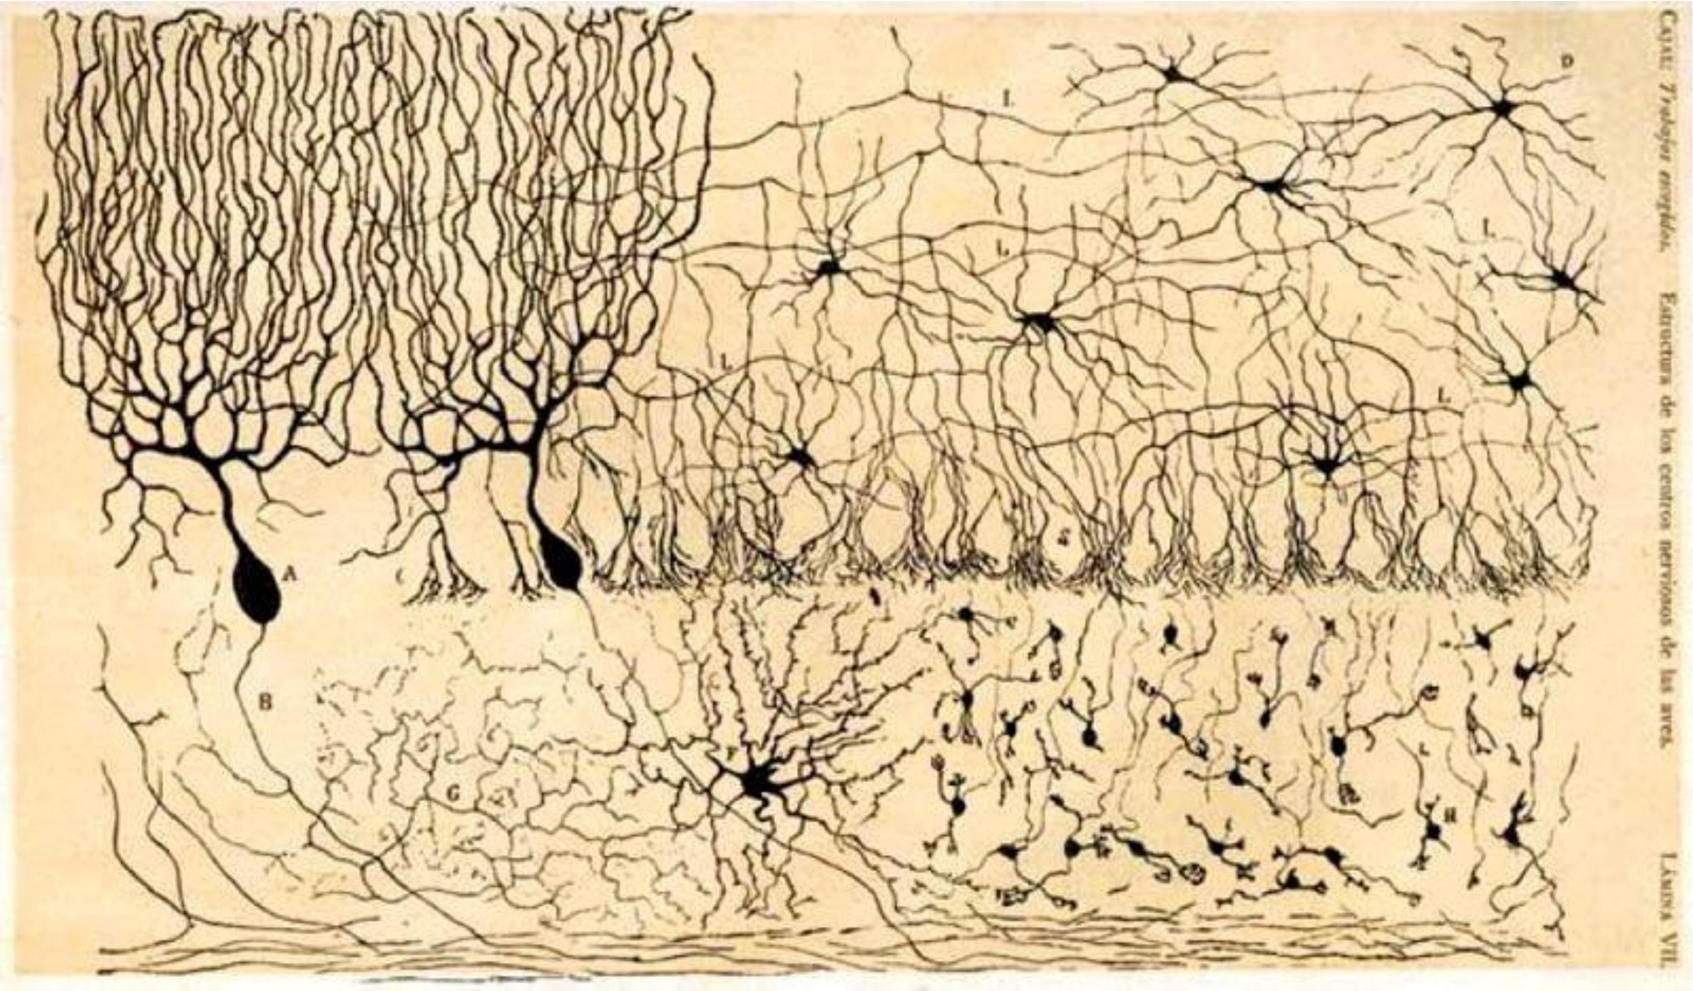

Ramón y Cajal

https://www.tabakalera.eu/en/aesthetic-knowledge-and-sensibility-drawings-santiago-ramon-y-cajal

Biological to Artificial

Biological Neuron:

- Cell body acts like a leaky capacitor: summation + smoothing
- Inputs (generally) to dendrites: short-duration injections of + or – current
- Cell body sums inputs from all dendrites
- Axon contains active processes: nonlinear response to cell body voltage

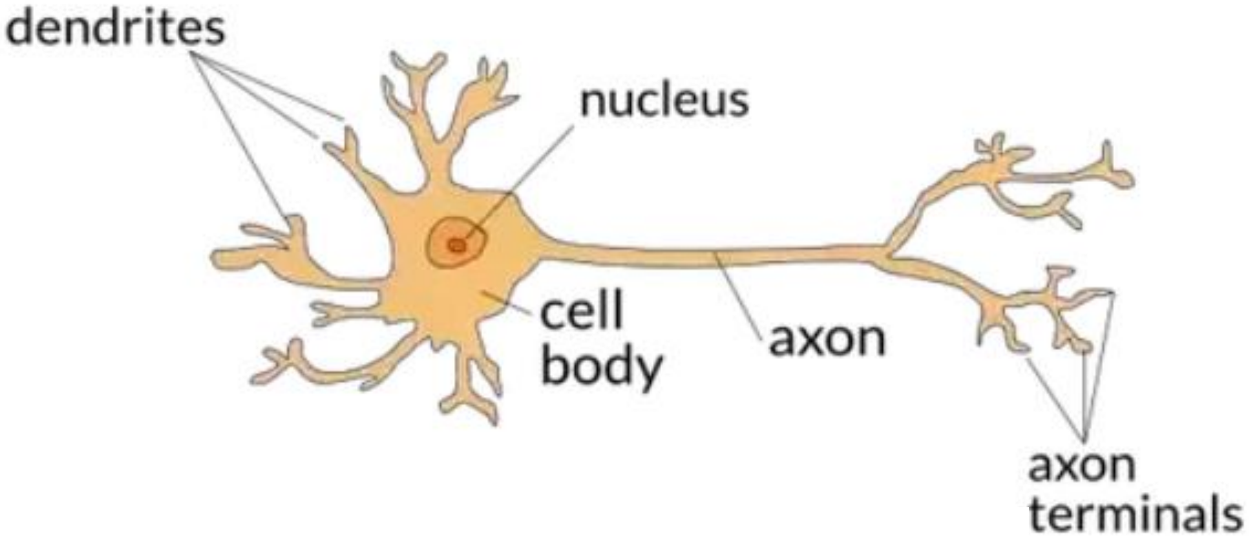

https://towardsdatascience.com/the-differences-between-artificial-and-biological-neural-networks-a8b46db828b7

Biological to Artificial

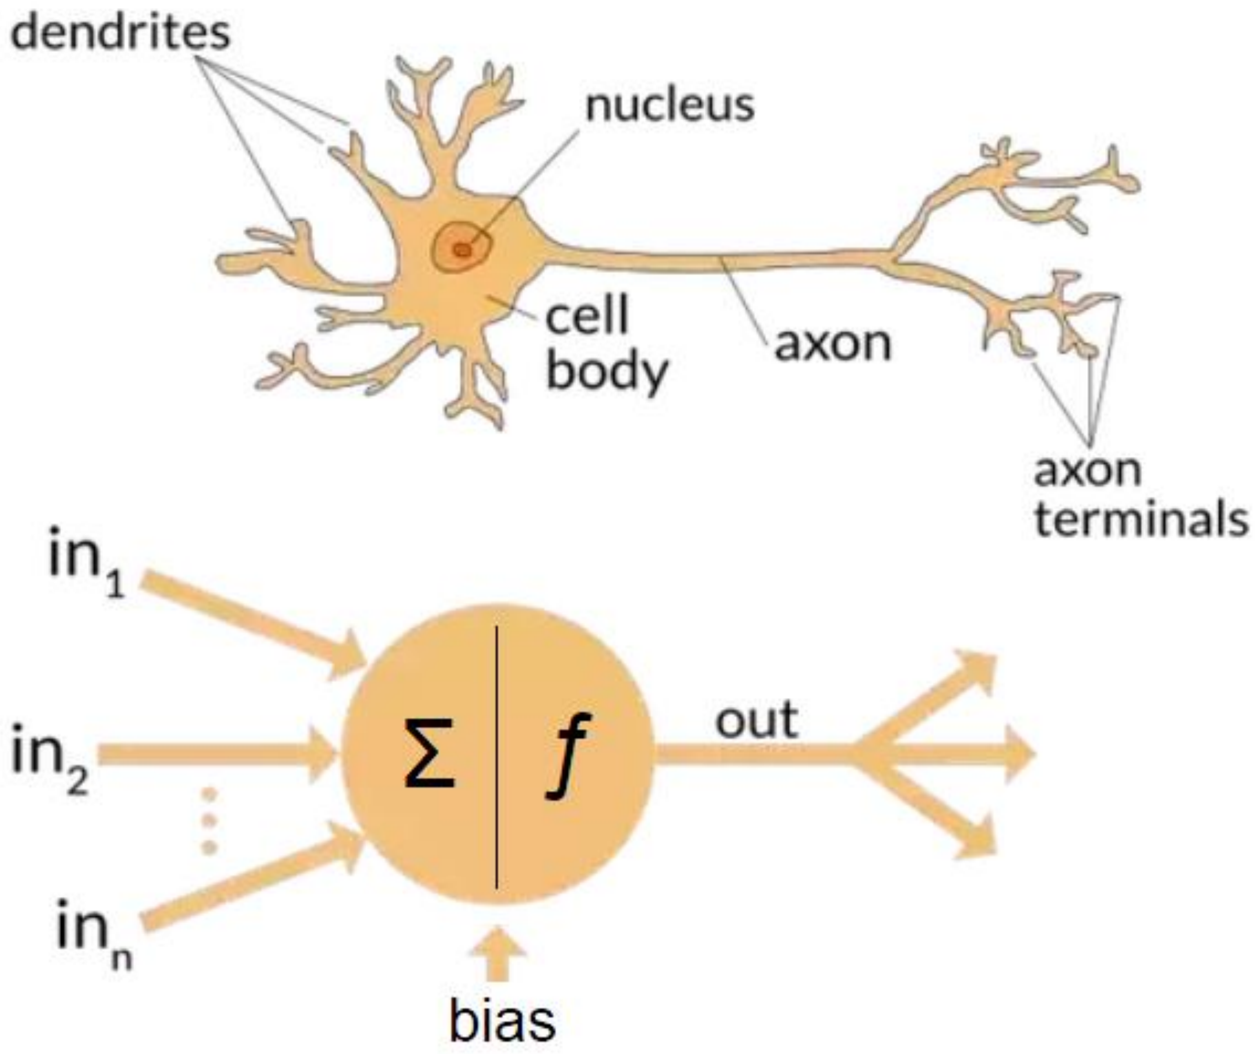

Abstraction:

- Each input is weighted
- Weighted inputs are summed
- Output is a nonlinear function of the sum

For many models, we ignore the issue of time

https://towardsdatascience.com/the-differences-between-artificial-and-biological-neural-networks-a8b46db828b7

What can a neuron compute?

Layered Networks: Local Circuits

- Layers within a small slice of cortex
- Different types of neurons receive inputs from different sources & produce outputs to different parts of the circuit

Lynch (2015)

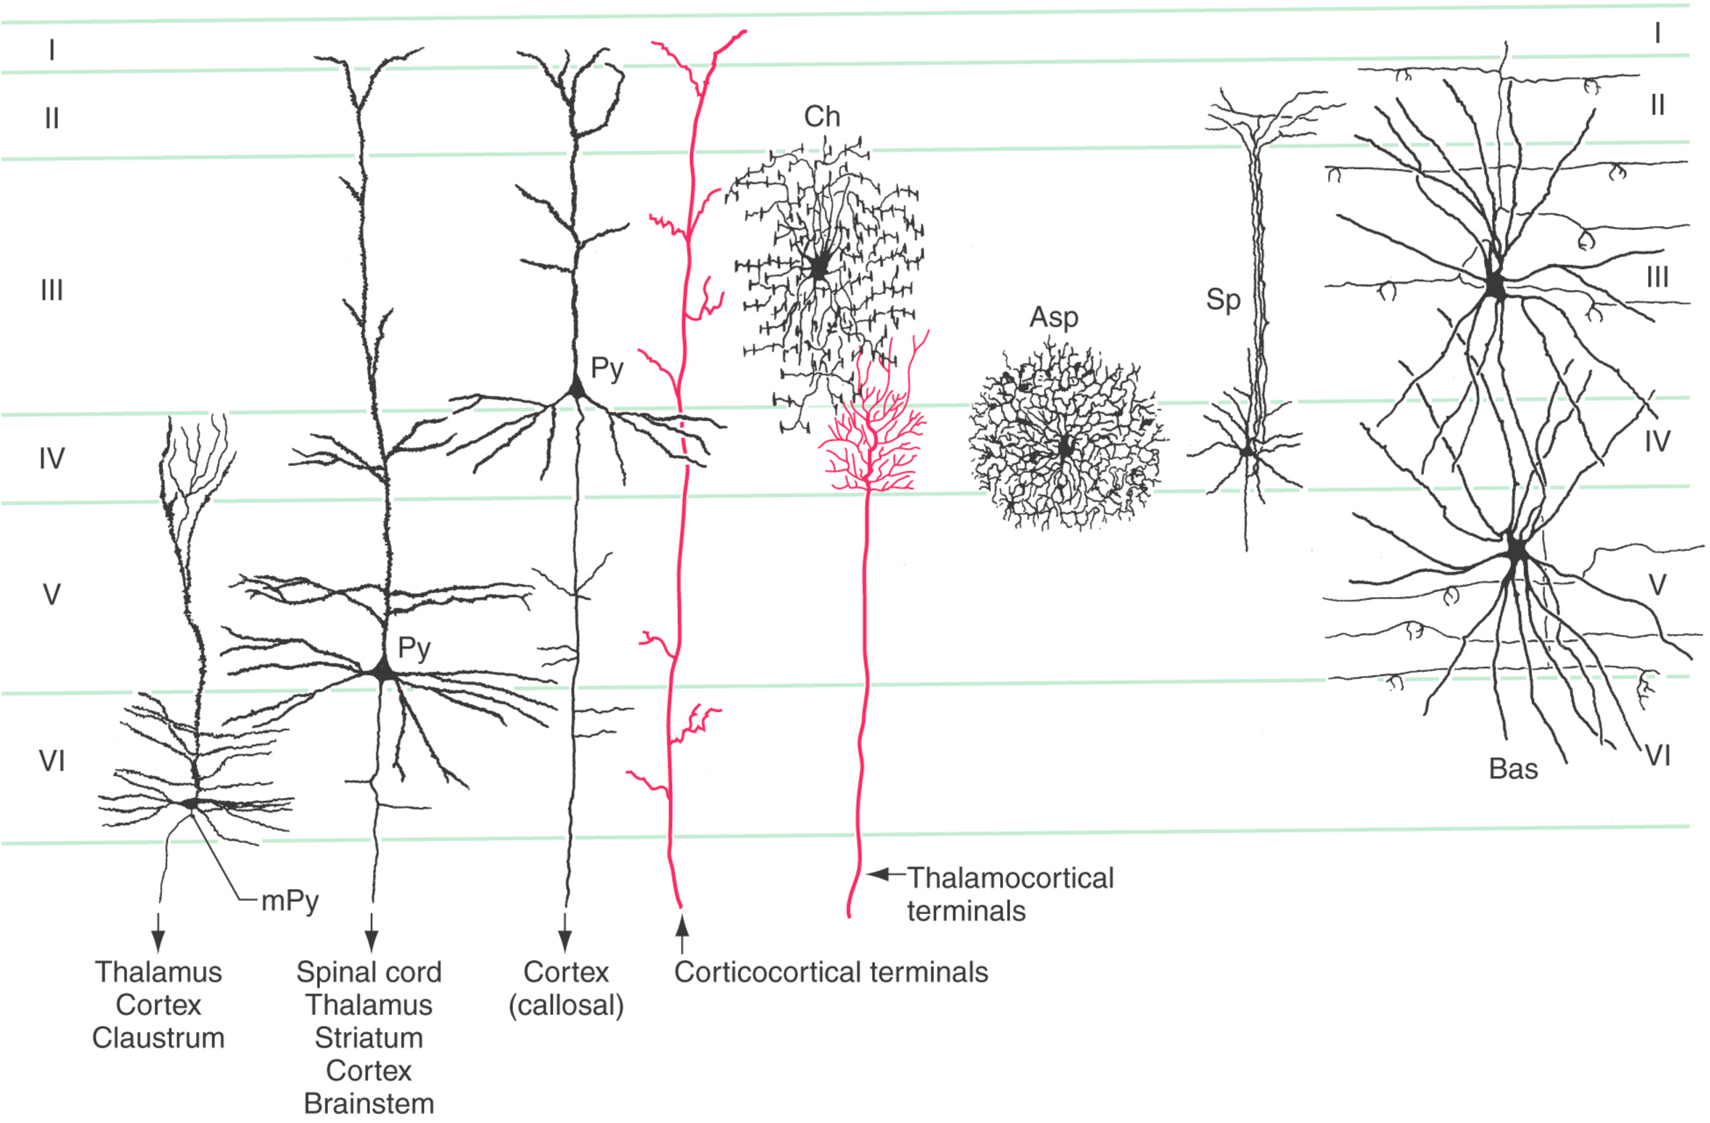

Layered Networks: Large-Scale Circuits

- Primary visual pathways
- V1 (primary visual cortex): compute simple visual features
- Downstream regions recognize more complex patterns

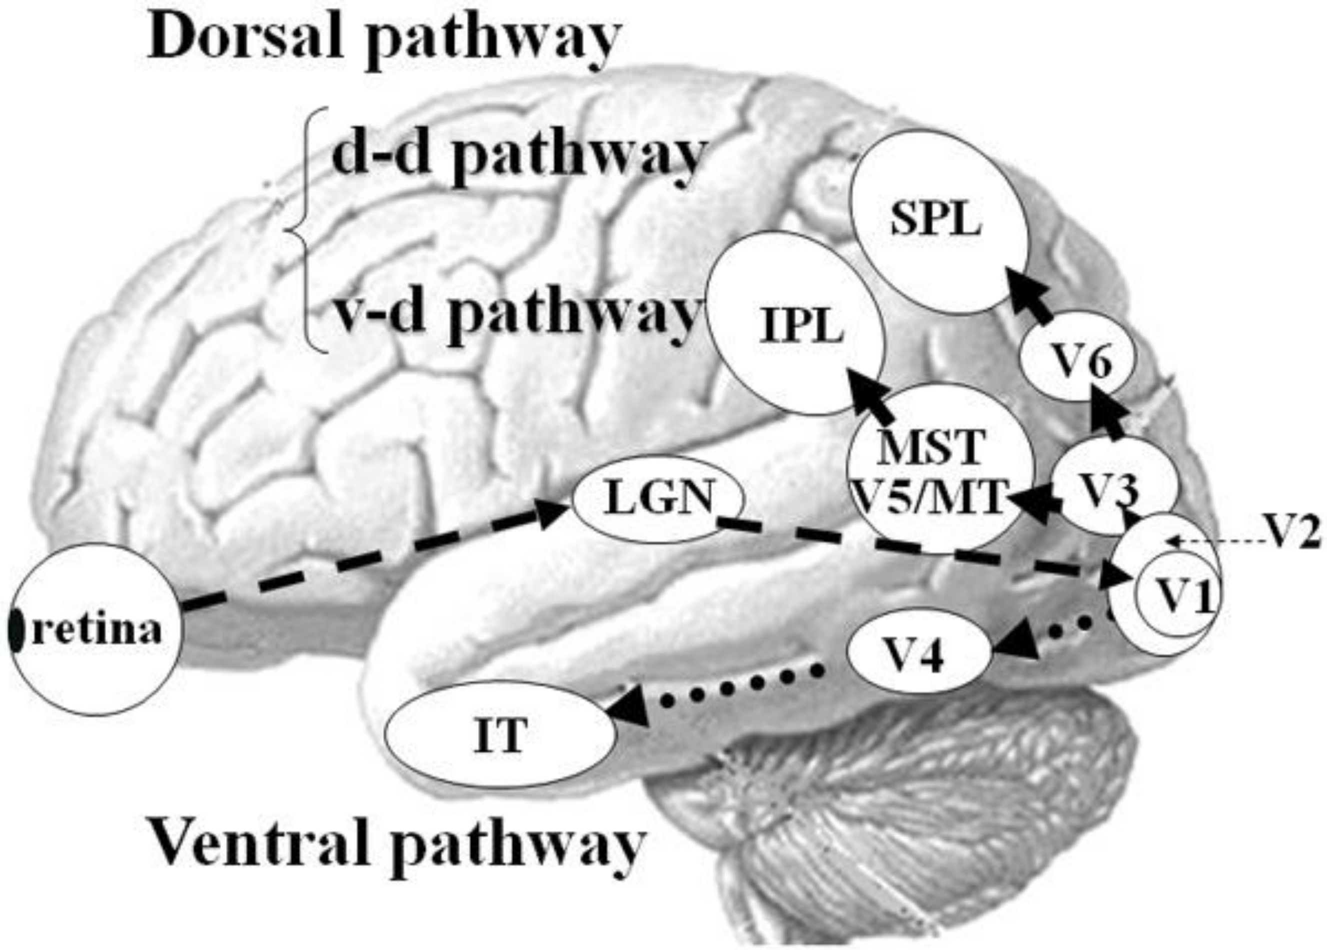

Yamasaki & Tobimatsu (2011)

Feed-Forward Neural Networks

- Forward propagation: input units to hidden to output
  - Form of function approximation
- Backpropagation of error: measure output error & compute changes to parameters to reduce that error
  - Form of regression

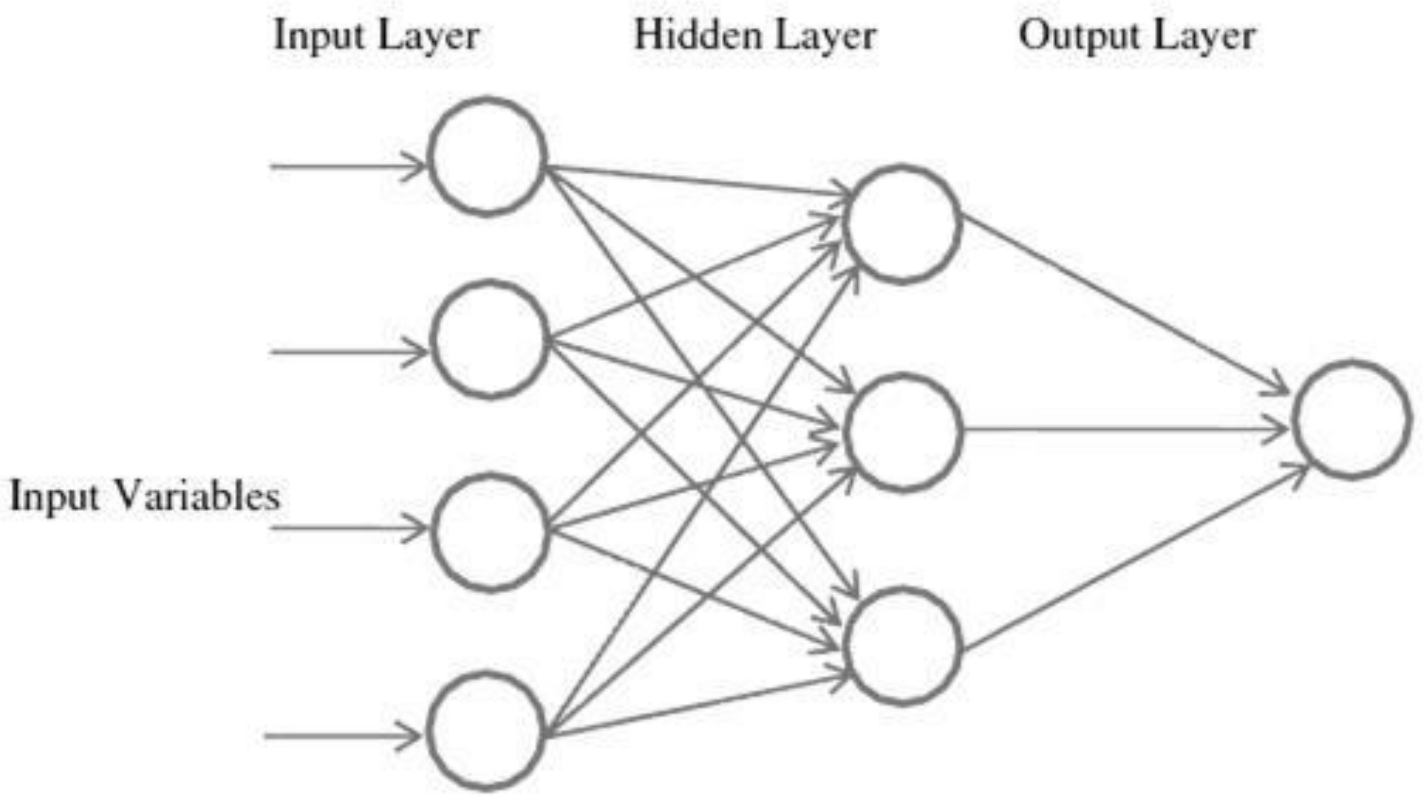

Roushangar & Homayounfar (2015)

Feed-Forward Neural Networks

- Can represent more interesting functions than a single neuron
- But: selecting parameters becomes more complex

Feed-Forward Neural Networks

- Two hidden layers can approximate any continuous function with arbitrary accuracy
- Deeper networks can capture complicated functions more easily
  - But much harder to find solutions

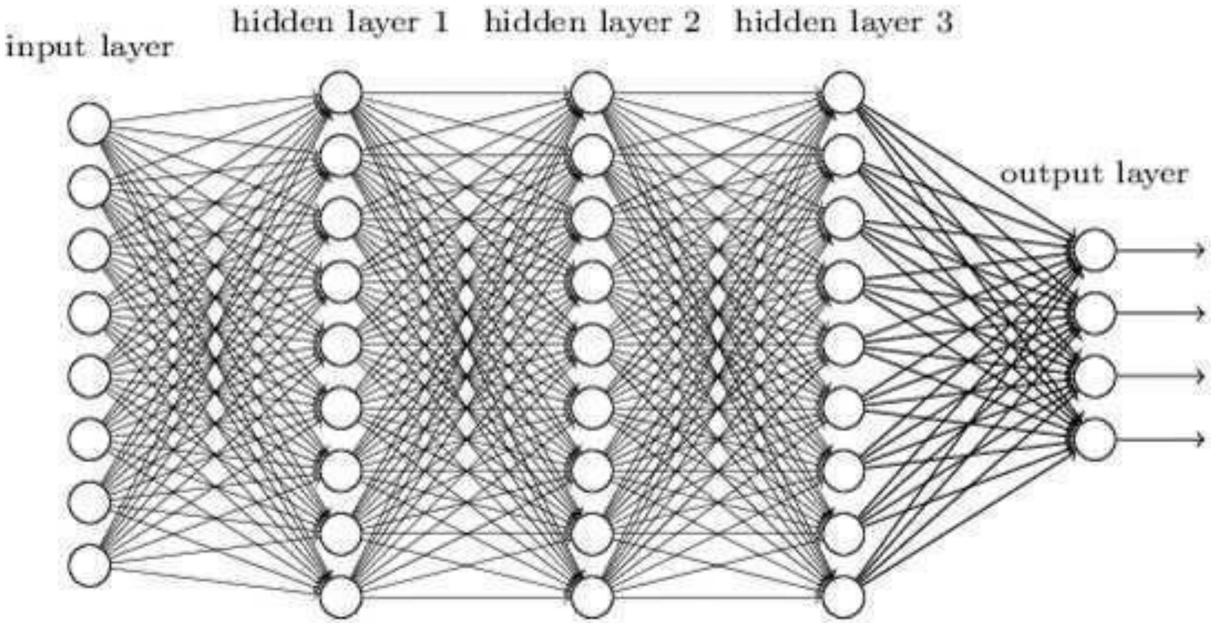

https://www.researchgate.net/figure/Deep-Neural-Network_fig5_316451216

Deep Learning

- Many layers (sometimes 100s)
- Successful because of the confluence of many factors

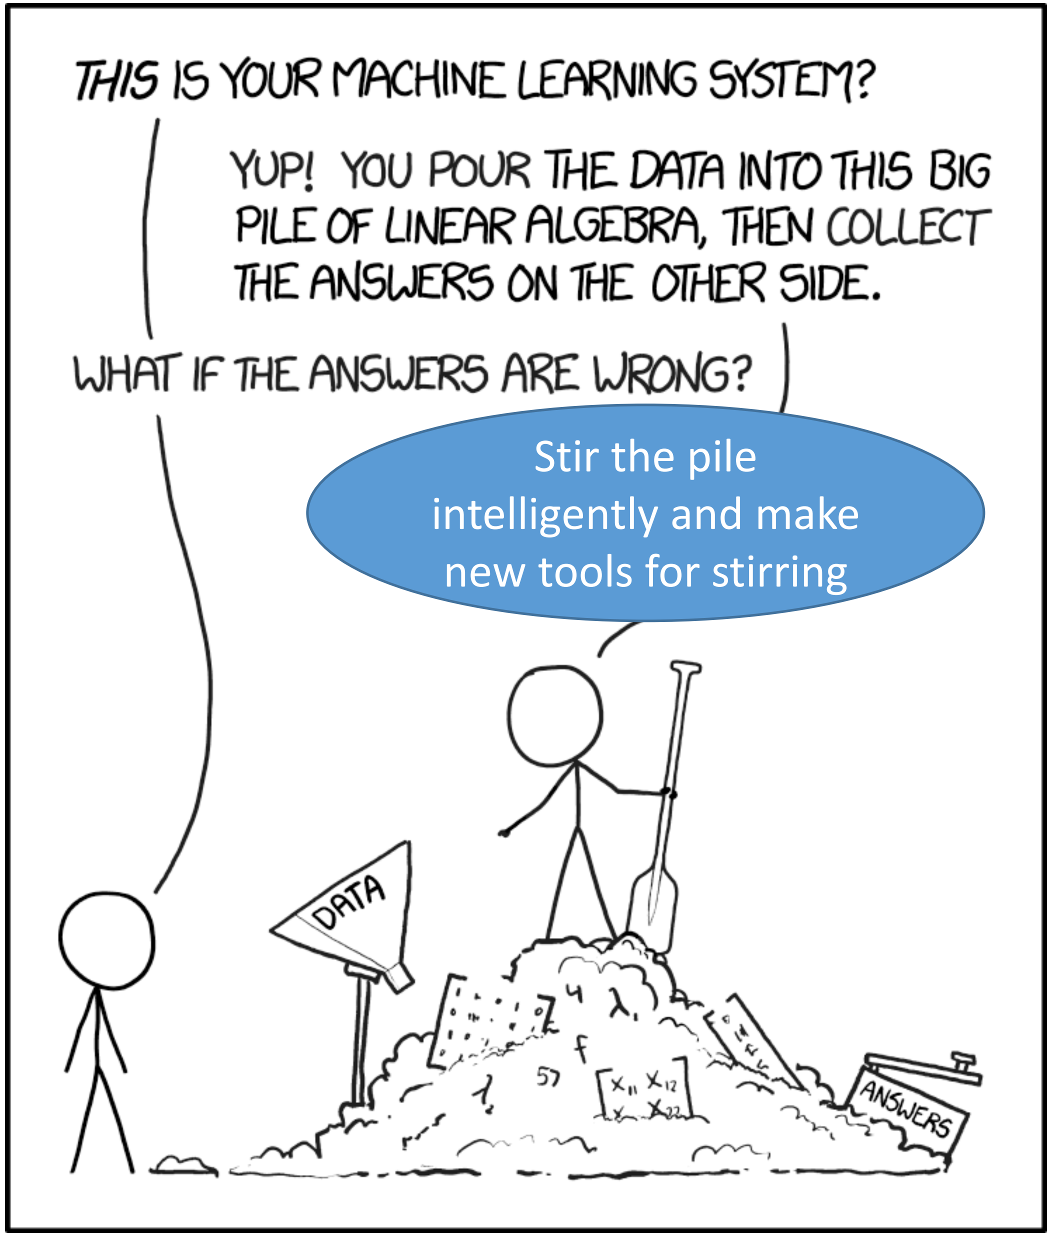

https://xkcd.com/1838/

Deep Learning

- Better and more specialized computation (especially GPUs and TPUs)
- Large data sets
- Large & efficient memory structures
- Many algorithmic “tweaks”
  - Easy to compute non-linearities
  - Fast estimation of gradients in the network (e.g., autodiff and mini batching)
  - Incremental learning of subnetworks (e.g., autoencoders)
- Software tools, including: python, TensorFlow and Keras

Imagenet Data Set

http://www.image-net.org/

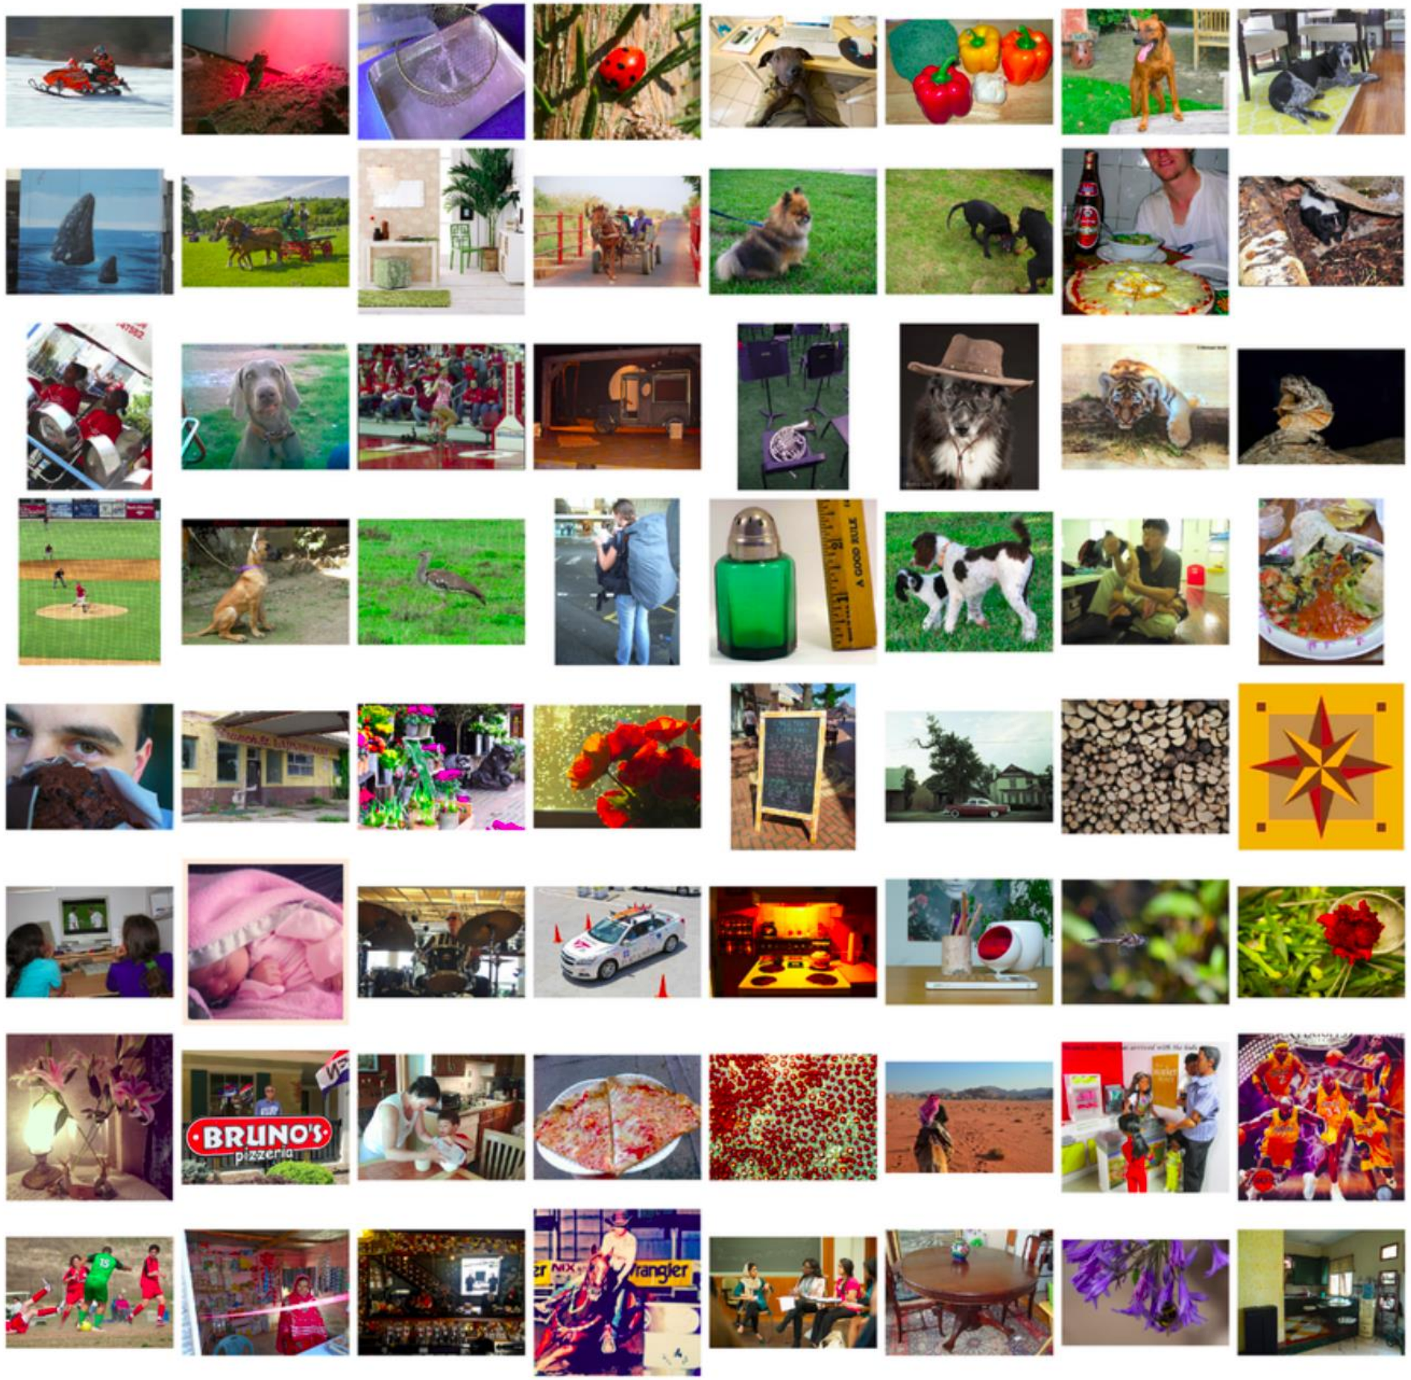

Convolutional Neural Networks

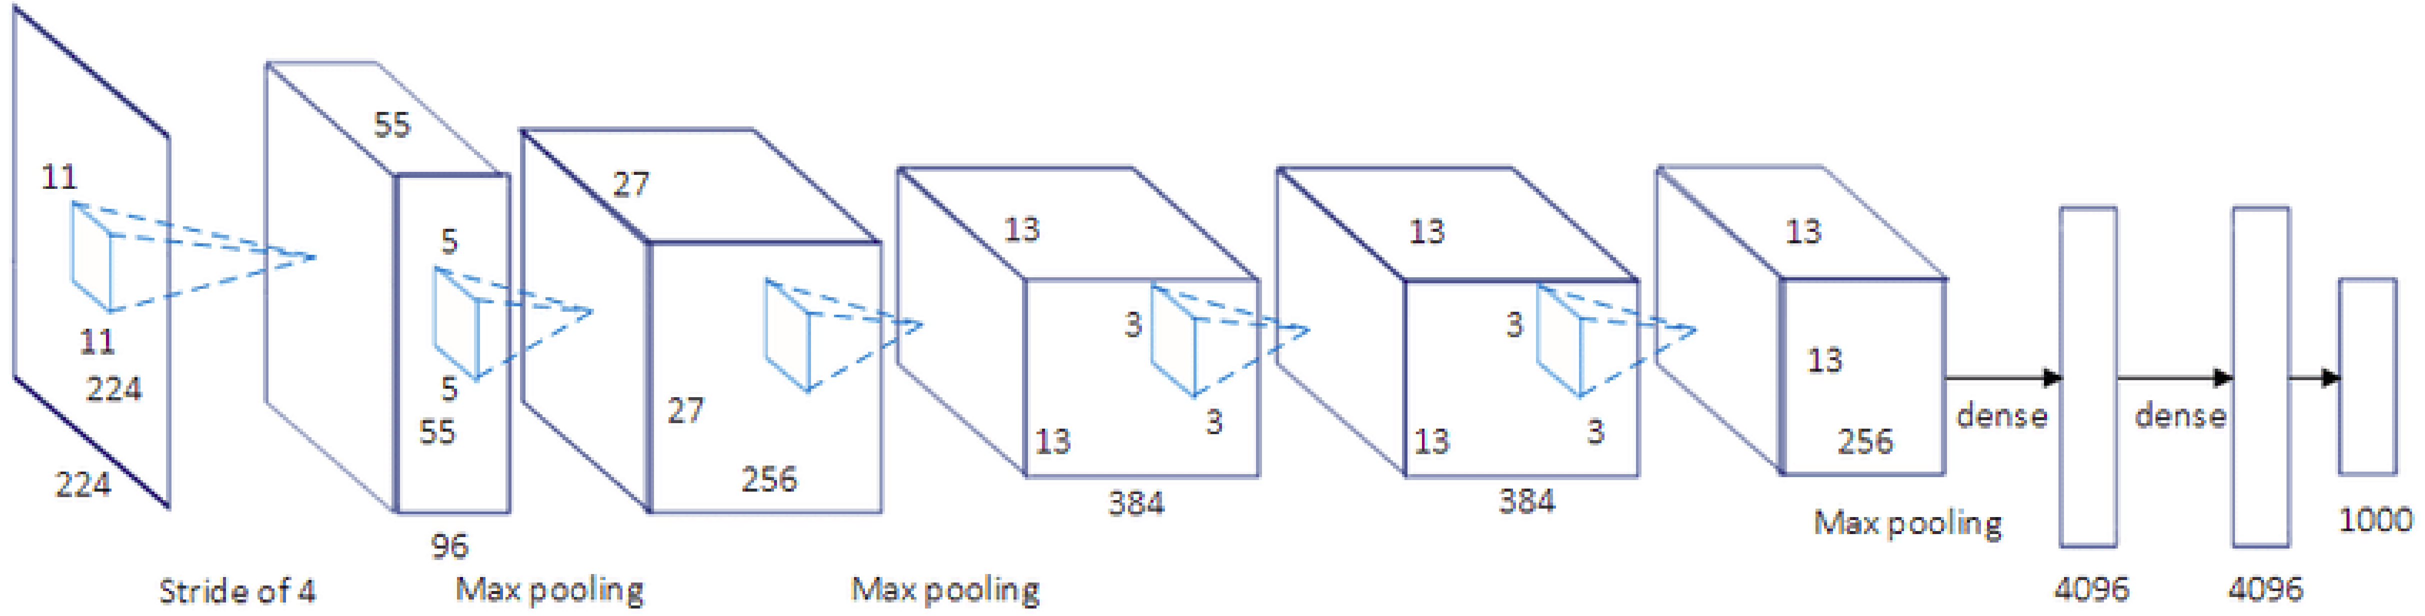

https://www.researchgate.net/figure/AlexNet-Convolutional-Neural-Network-architecture-Figure-reproduced-from-14_fig1_316450908

Inception Net (Google)

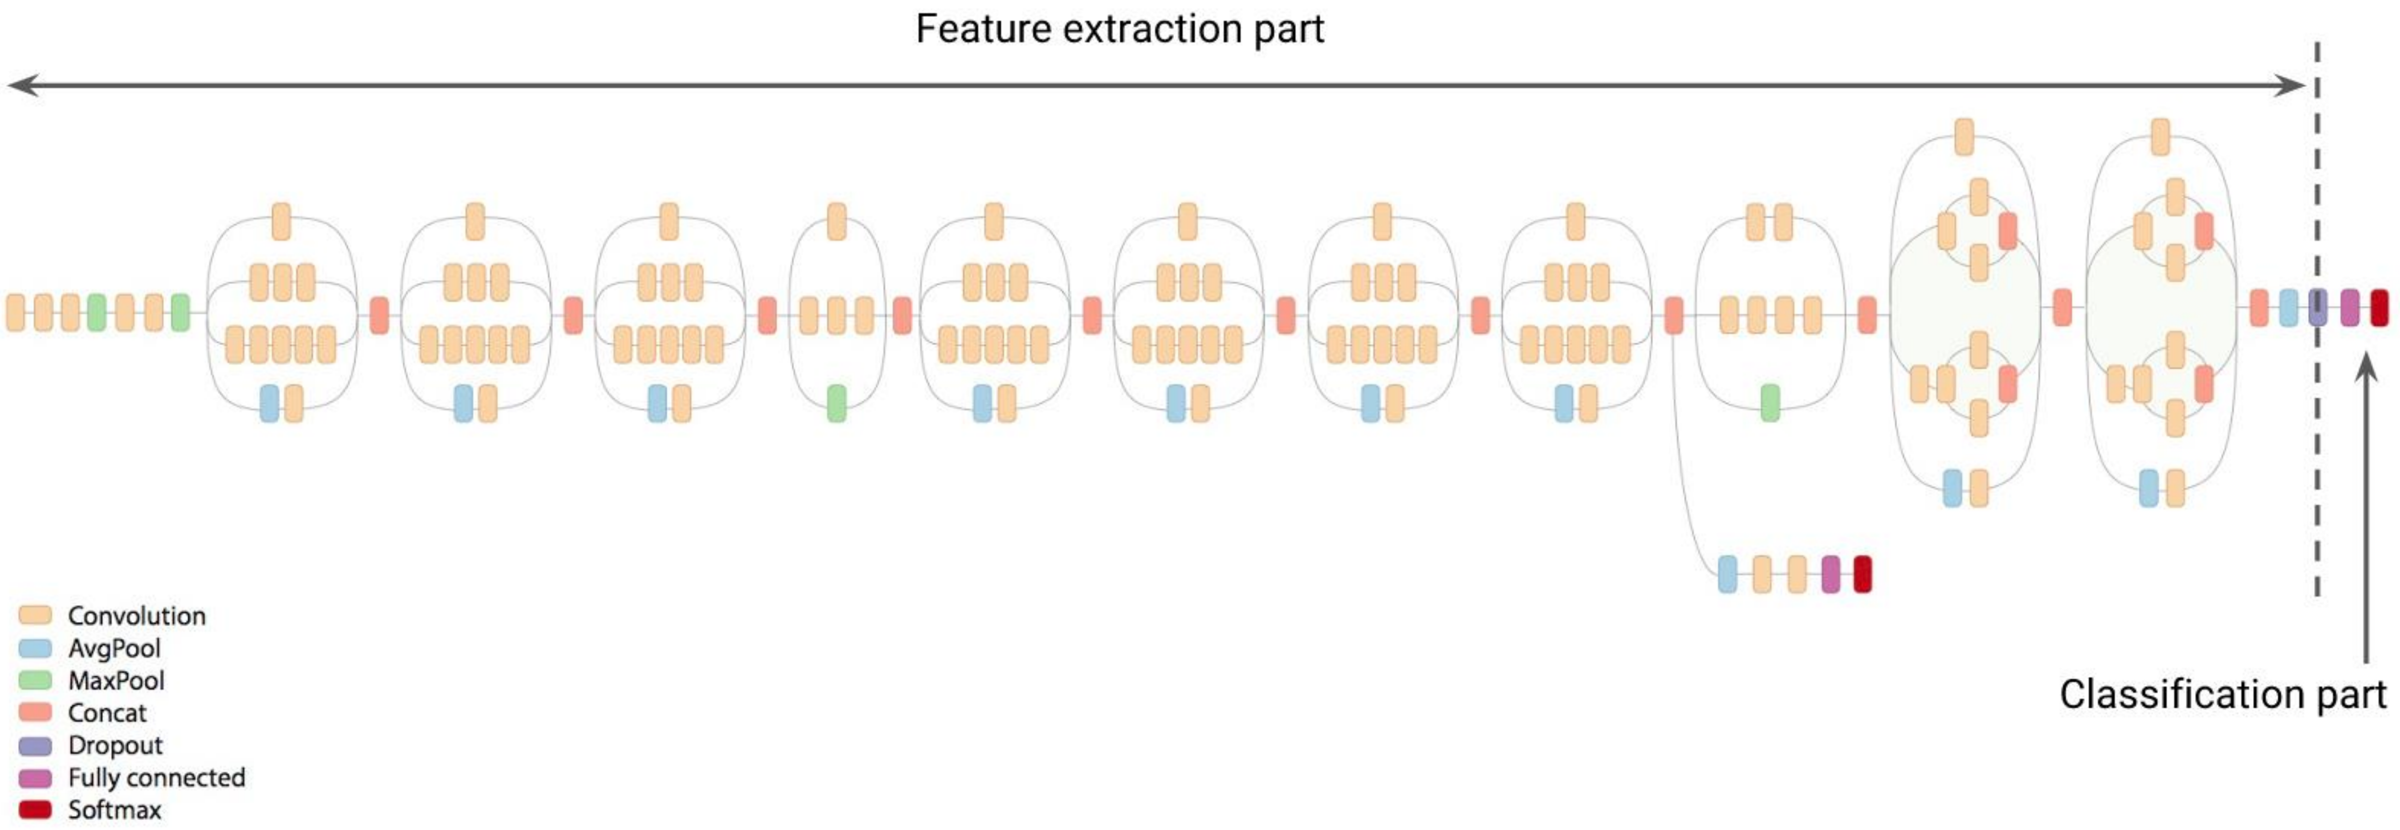

googleblog.com

Image "Understanding"

- Finding interesting areas in the image
- Classifying subimages
- Recognizing specific instances of a class in a subimage

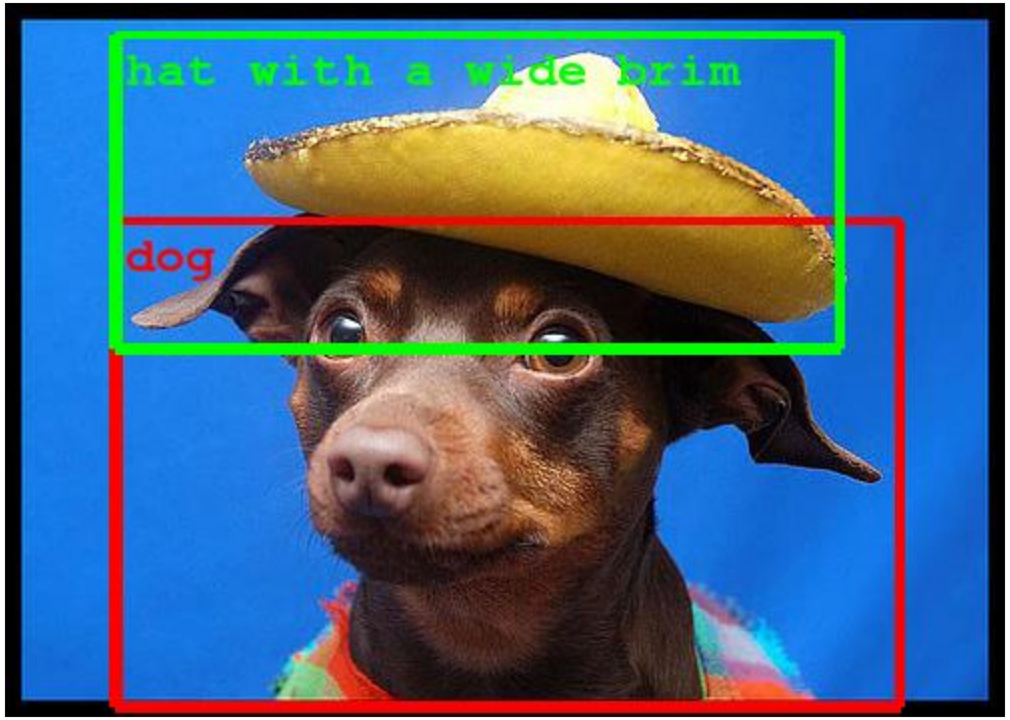

Google AI Vision Kit ($99)

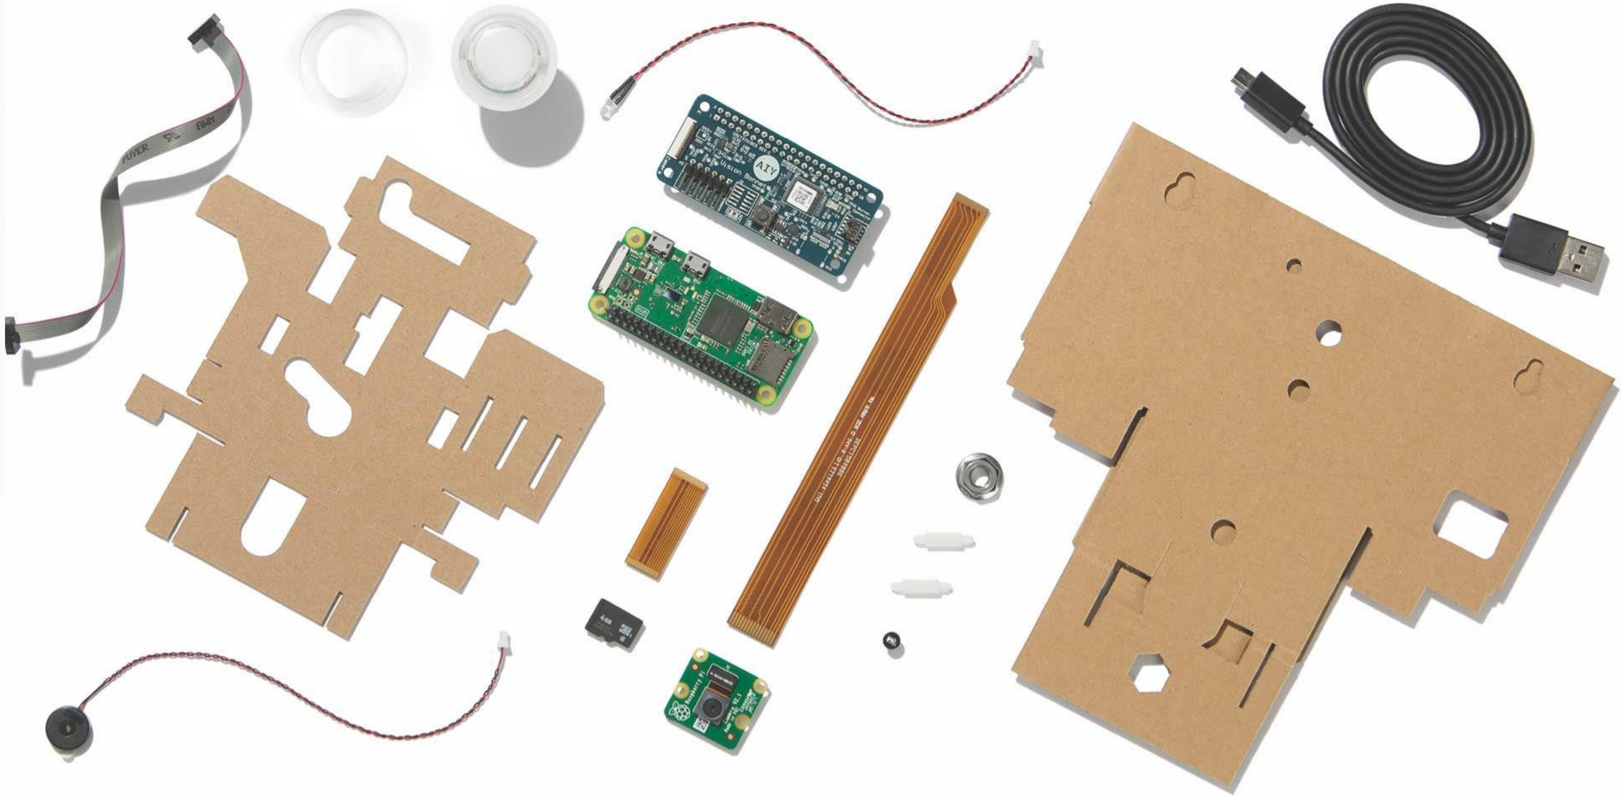

Handling Time: Recurrent Neural Networks

Elman Network (adaptation of a Jordan Network):

- Sequence of inputs
- Sequence of outputs
- Network can use context as a memory of what has happened in the past

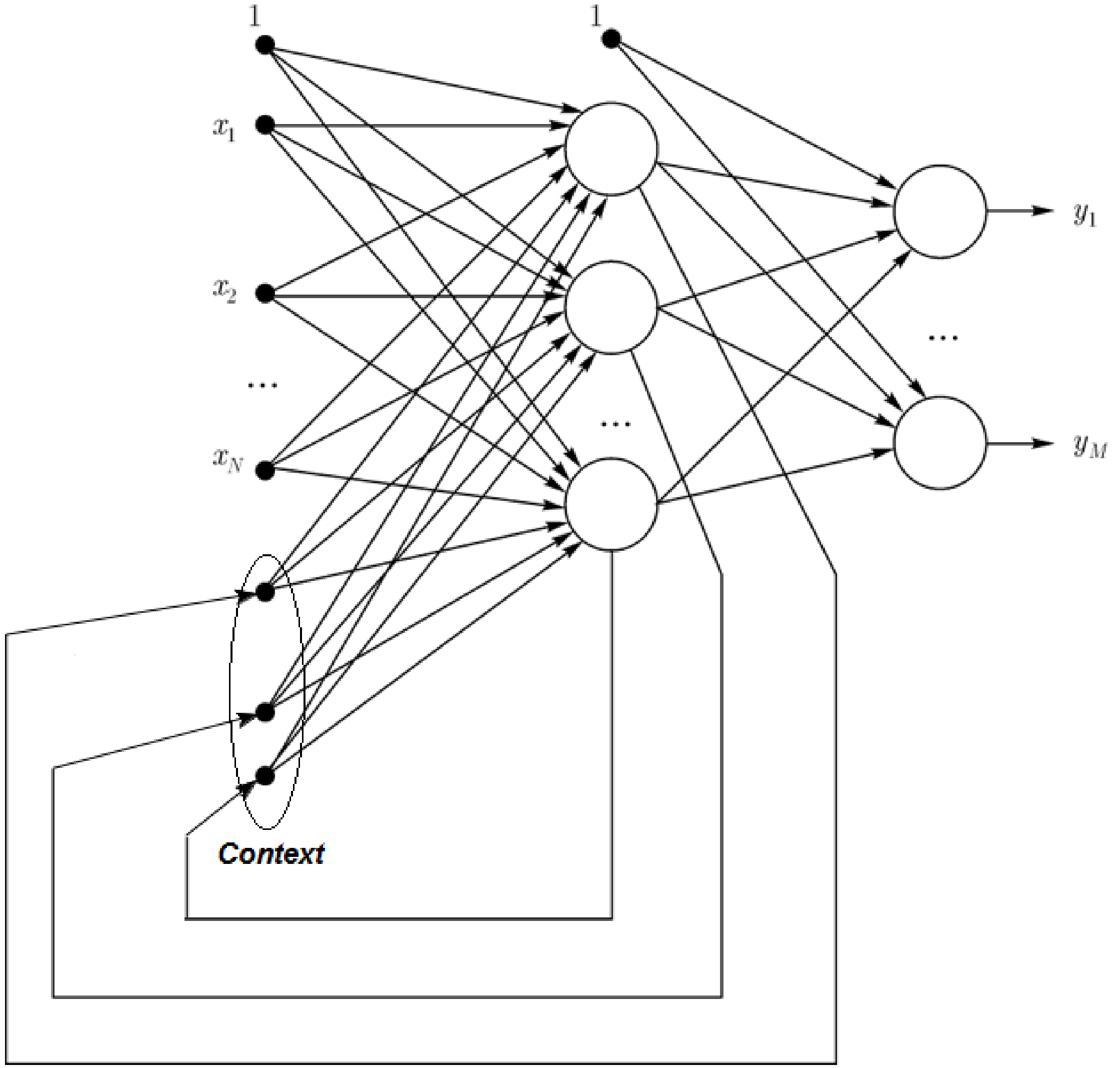

https://www.mql5.com/en/articles/1103

Core50 Data Set

Labeled objects as they are manipulated

https://vlomonaco.github.io/core50

Video to Audio Translation

- https://www.youtube.com/watch?v=0FW99AQmMc8&feature=youtu.be

Deep Reinforcement Learning

- https://www.youtube.com/watch?v=Ih8EfvOzBOY&index=2&list=PLujxSBD-JXglGL3ERdDOhthD3jTlfudC2

Generative Adversarial Networks

- https://www.youtube.com/watch?v=u7kQ5lNfUfg

Our Topics

- Backpropagation
- Model Evaluation Process: metrics, cross-validation, statistics, addressing the multiple comparisons problem
- Tools: TensorFlow and Keras
- Convolutional Neural Networks
- Recurrent Neural Networks
- Timeseries Processing
- Deep Reinforcement Learning
- Generative Models
- Generative Adversarial Networks

To “Solve” a Model Building Problem We Must Answer:

- What is the nature of the data that we have?
- How much data do we have?
- What is the prediction problem?
- How do we measure performance of a model?
- How to select an appropriate model and learning algorithm?
- How to choose parameters?
- How to convince ourselves (and others) that we have a useful model?

What I am assuming about you…

- Programming skills
- Able to jump into Python, including the "Object-Orientedness" of it
- Know or can learn unix command-line tools

Resources

- Course web page: `http://www.cs.ou.edu/~fagg/classes/aml`
- Francois Chollet (2017) Deep Learning with Python, 1st Edition, ISBN-13: 978-1617294433, Manning
- Web resources: documentation, tutorials, papers (linked from the schedule or announced on Canvas)

Computing Environment

Setting up a ML environment (especially one based on TensorFlow) can be a bear …

- We are providing a pre-configured compute cluster on Amazon Web Services (AWS)
- Key tools: Python, Scikit-learn, TensorFlow (Deep Learning), Keras, Jupyter (Interactive Development Environment)
- Other software: editors (emacs, vi, gedit)
- Will also house our common data sets

Computing Environment

AWS machines are nice, but cost us based on the resources we use

- Right now, we have one AWS machine configured (1 processor; 2 GB of memory; 32 GB of swap).
- This machine will increase in resources as needed
- Additional machines will be added as needed
  - They will share a common user accounting system and file system
- As our computing needs get larger, we will shift over to the OU Supercomputer

Configure your Cluster Account

Details coming soon for your cluster account …

Homework Assignments

- First 1/3 of the semester
- Explore different DL methods and data sets

Projects

- Last 2/3 of the semester
- Topic / data set are your choice, but must be approved
- Several in-class presentations
- Final paper

Grading

- In-class participation: 10%
- Homework: 30%
- Project work: 60%

Proper Academic Conduct

- Homework assignments are to be done on your own
  - No communication of solutions in any form
  - Do not copy code off the net

- Projects:
  - Groups of 2 working on related problems

For Next Time

- For today: chapter 1
- Next time: start of chapter 2

- We will get you started on python and numpy