

# CS342 - NEURAL NETWORKS

Philipp Krähenbühl

# OVERVIEW

[Philipp Krähenbühl](#)

office hours T 11-12  
GDC 4.824

**TA Dian Chen**

office hours M 11-12  
GDC 1.302

**TA Ankur Garg**

office hours M 4-5  
GDC 1.302

Try canvas first!

[www.philkr.net/  
cs342/](http://www.philkr.net/cs342/)

# OVERVIEW

T



Th



Fr



# IN CLASS

- Learn background material
  - “Theory”
  - Math behind deep learning



# SECTION

- Coding
  - PyTorch
  - Start homework




# IN CLASS QUIZZES

- In average one per week (~13 in total)
- Can work in groups
- Can fail 2, at least 5 required
- Ways to fail
  - Do not show up
  - Hand in an empty quiz
  - Write a purposefully wrong answer



# ASSIGNMENTS

- One homework per week
  - Out Thursday night (or Friday morning)
  - due next Th 11:59pm
  - Submit a pytorch module
  - Automatically graded
    - partial grader with assignment



# ASSIGNMENTS





# FINAL PROJECT

- Learn to race in SuperTuxKart
  - Last 3 weeks
  - open ended strategy
- Competition in class



# PREREQUISITES

- 311 or 311H - Discrete math for computer science (or equivalent)
- 343 or 363D - Artificial Intelligence or Statistical Learning and Data Mining
- Proficiency in Python

# GOALS

- Implement and train neural networks in pytorch
- Have a basic understanding of the inner workings of neural networks
- Know several types of neural networks, including convolutional and recurrent neural networks

# GRADES

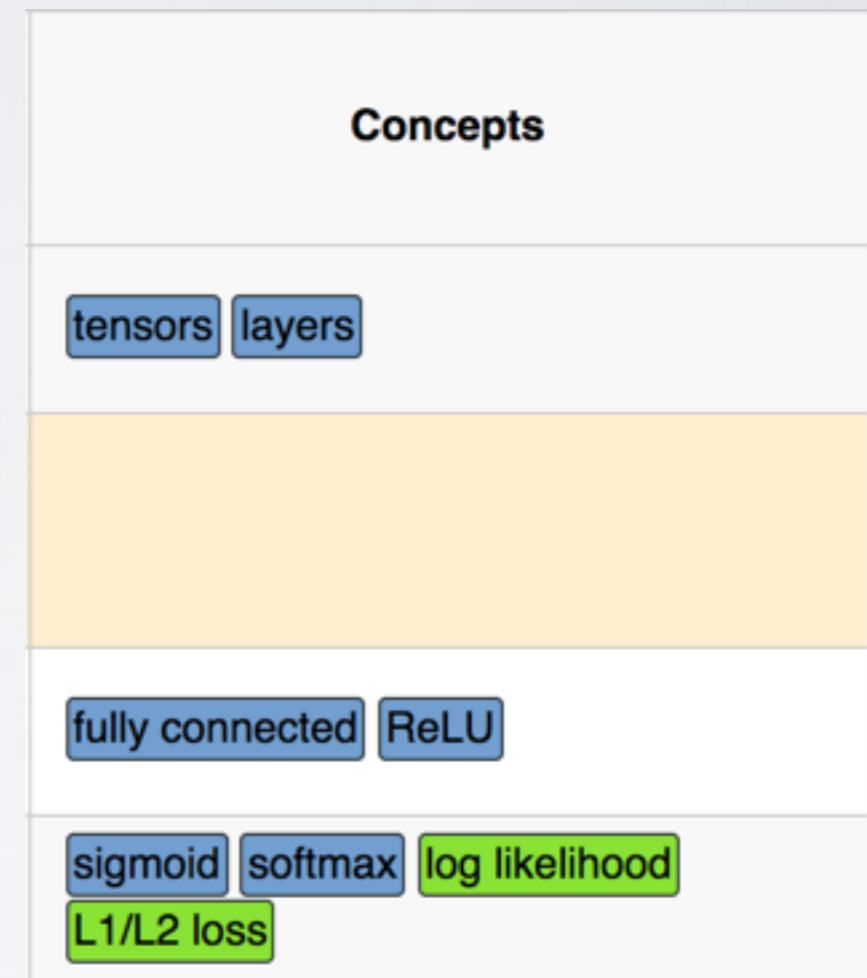
- 10% Quizzes
- 55% homework
- 35% final project
- 1 day late: -25%
- 2 days late: -50%
- 3+ days late: -100%

# BOOK

- No official book
- For background reading:
  - Deep learning, Goodfellow, Bengio and Courville 2016
  - [www.deeplearningbook.org](http://www.deeplearningbook.org)


# WIKI / DOC

- The webpage as a little writeup for all important concepts you'll learn in class



ONCE UPON A TIME

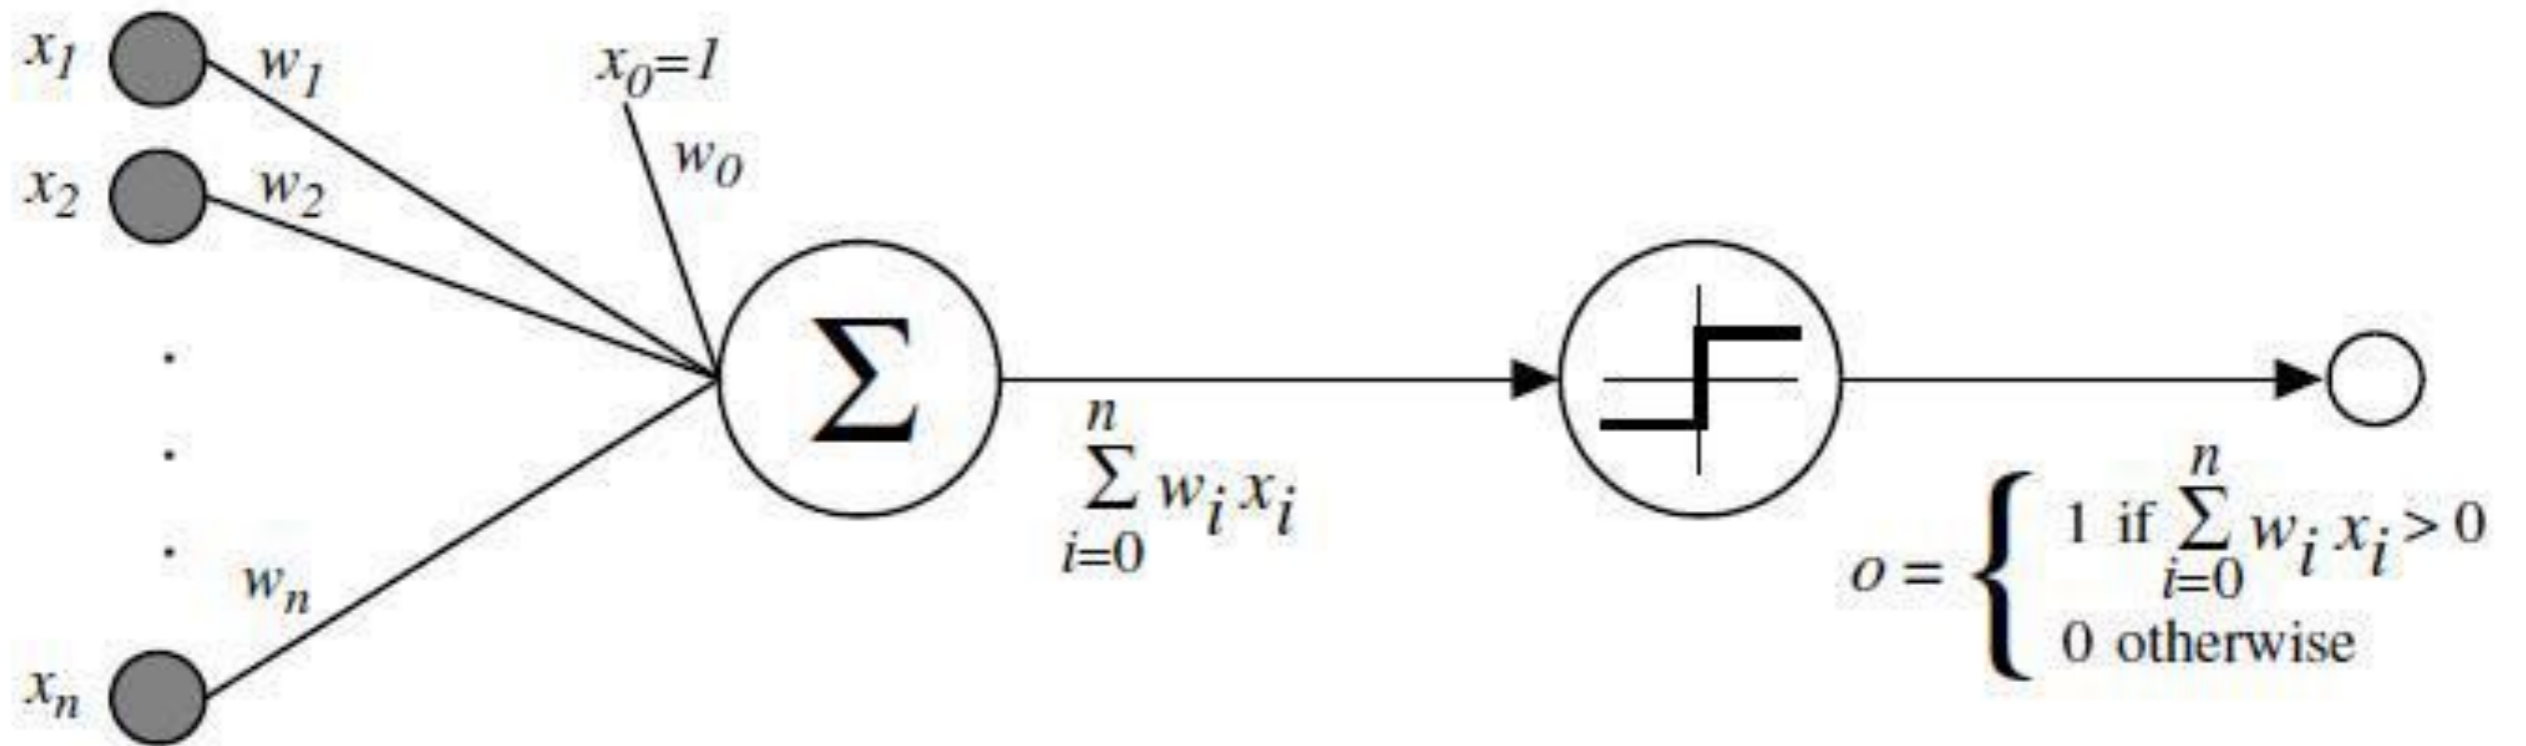
1950s






# PERCEPTRON

Frank Rosenblatt [1957]




# PERCEPTRON

$$f(x) = \begin{cases} 1 & \text{if } w^*x > 0 \\ 0 & \text{otherwise} \end{cases}$$




# HISTORY OF DEEP LEARNING


Hubel & Wiesel [early 1950s]



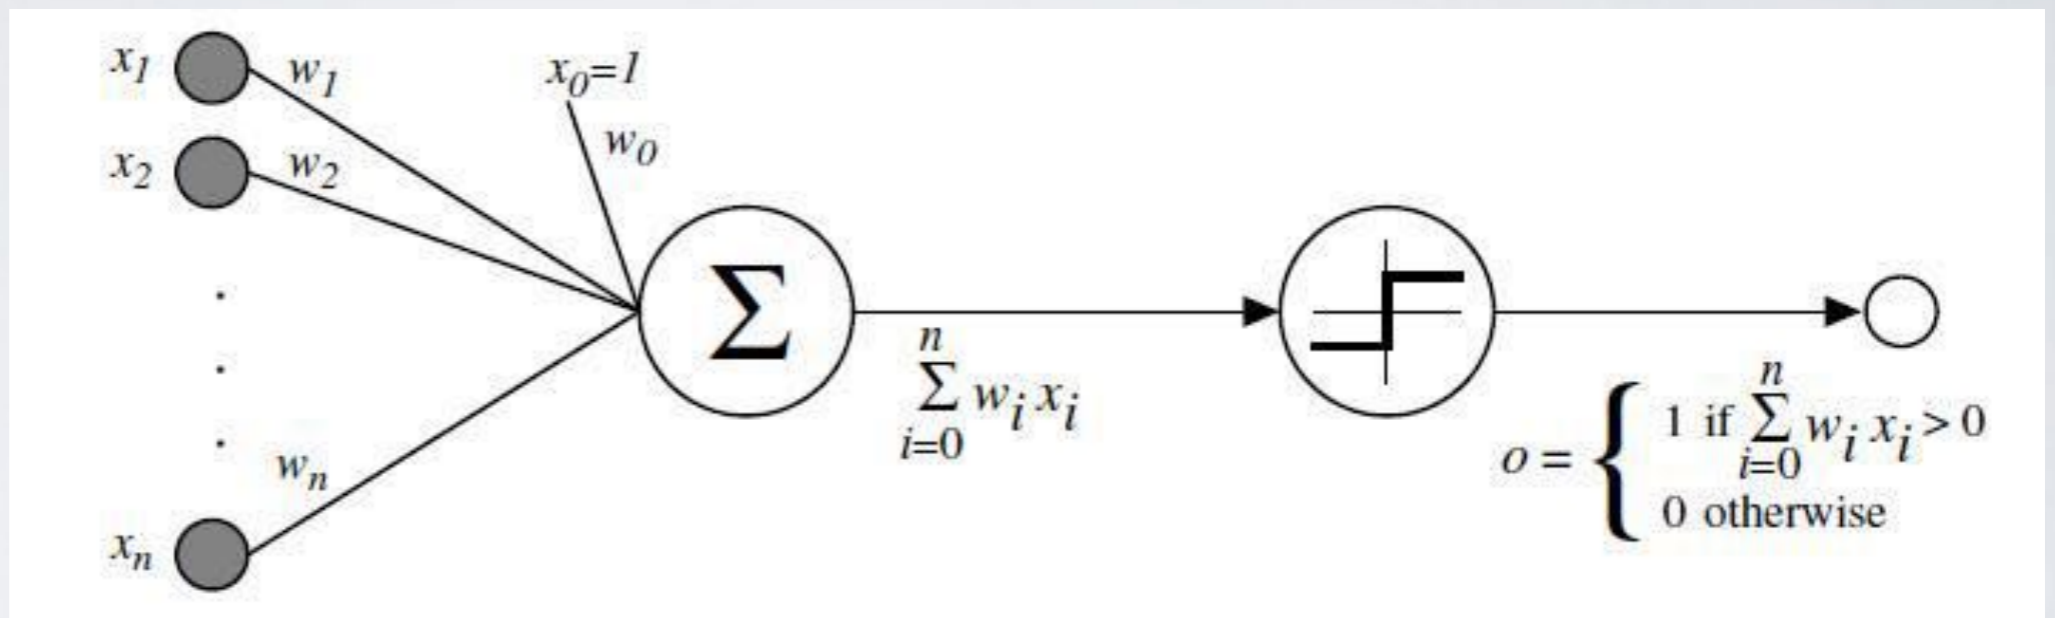




# MULTI-LAYER PERCEPTRON





# MULTI-LAYER PERCEPTRON



# MULTI-LAYER PERCEPTRON



# MARVIN MINSKY [1968]



symbolic AI  
rule based systems



AI winter [1970s]



# BACK-PROPAGATION [1988]



=





# MACHINE LEARNING + OPTIMIZATION

[1990 - 2005]

Random Forests

Support Vector Machines


Boosting

Hand engineering

Computer Vision: Geometry

# THE DEEP LEARNING CONSPIRACY

Minsky and  
Papert 1969



# THE DEEP LEARNING CONSPIRACY

## REVIEW

---

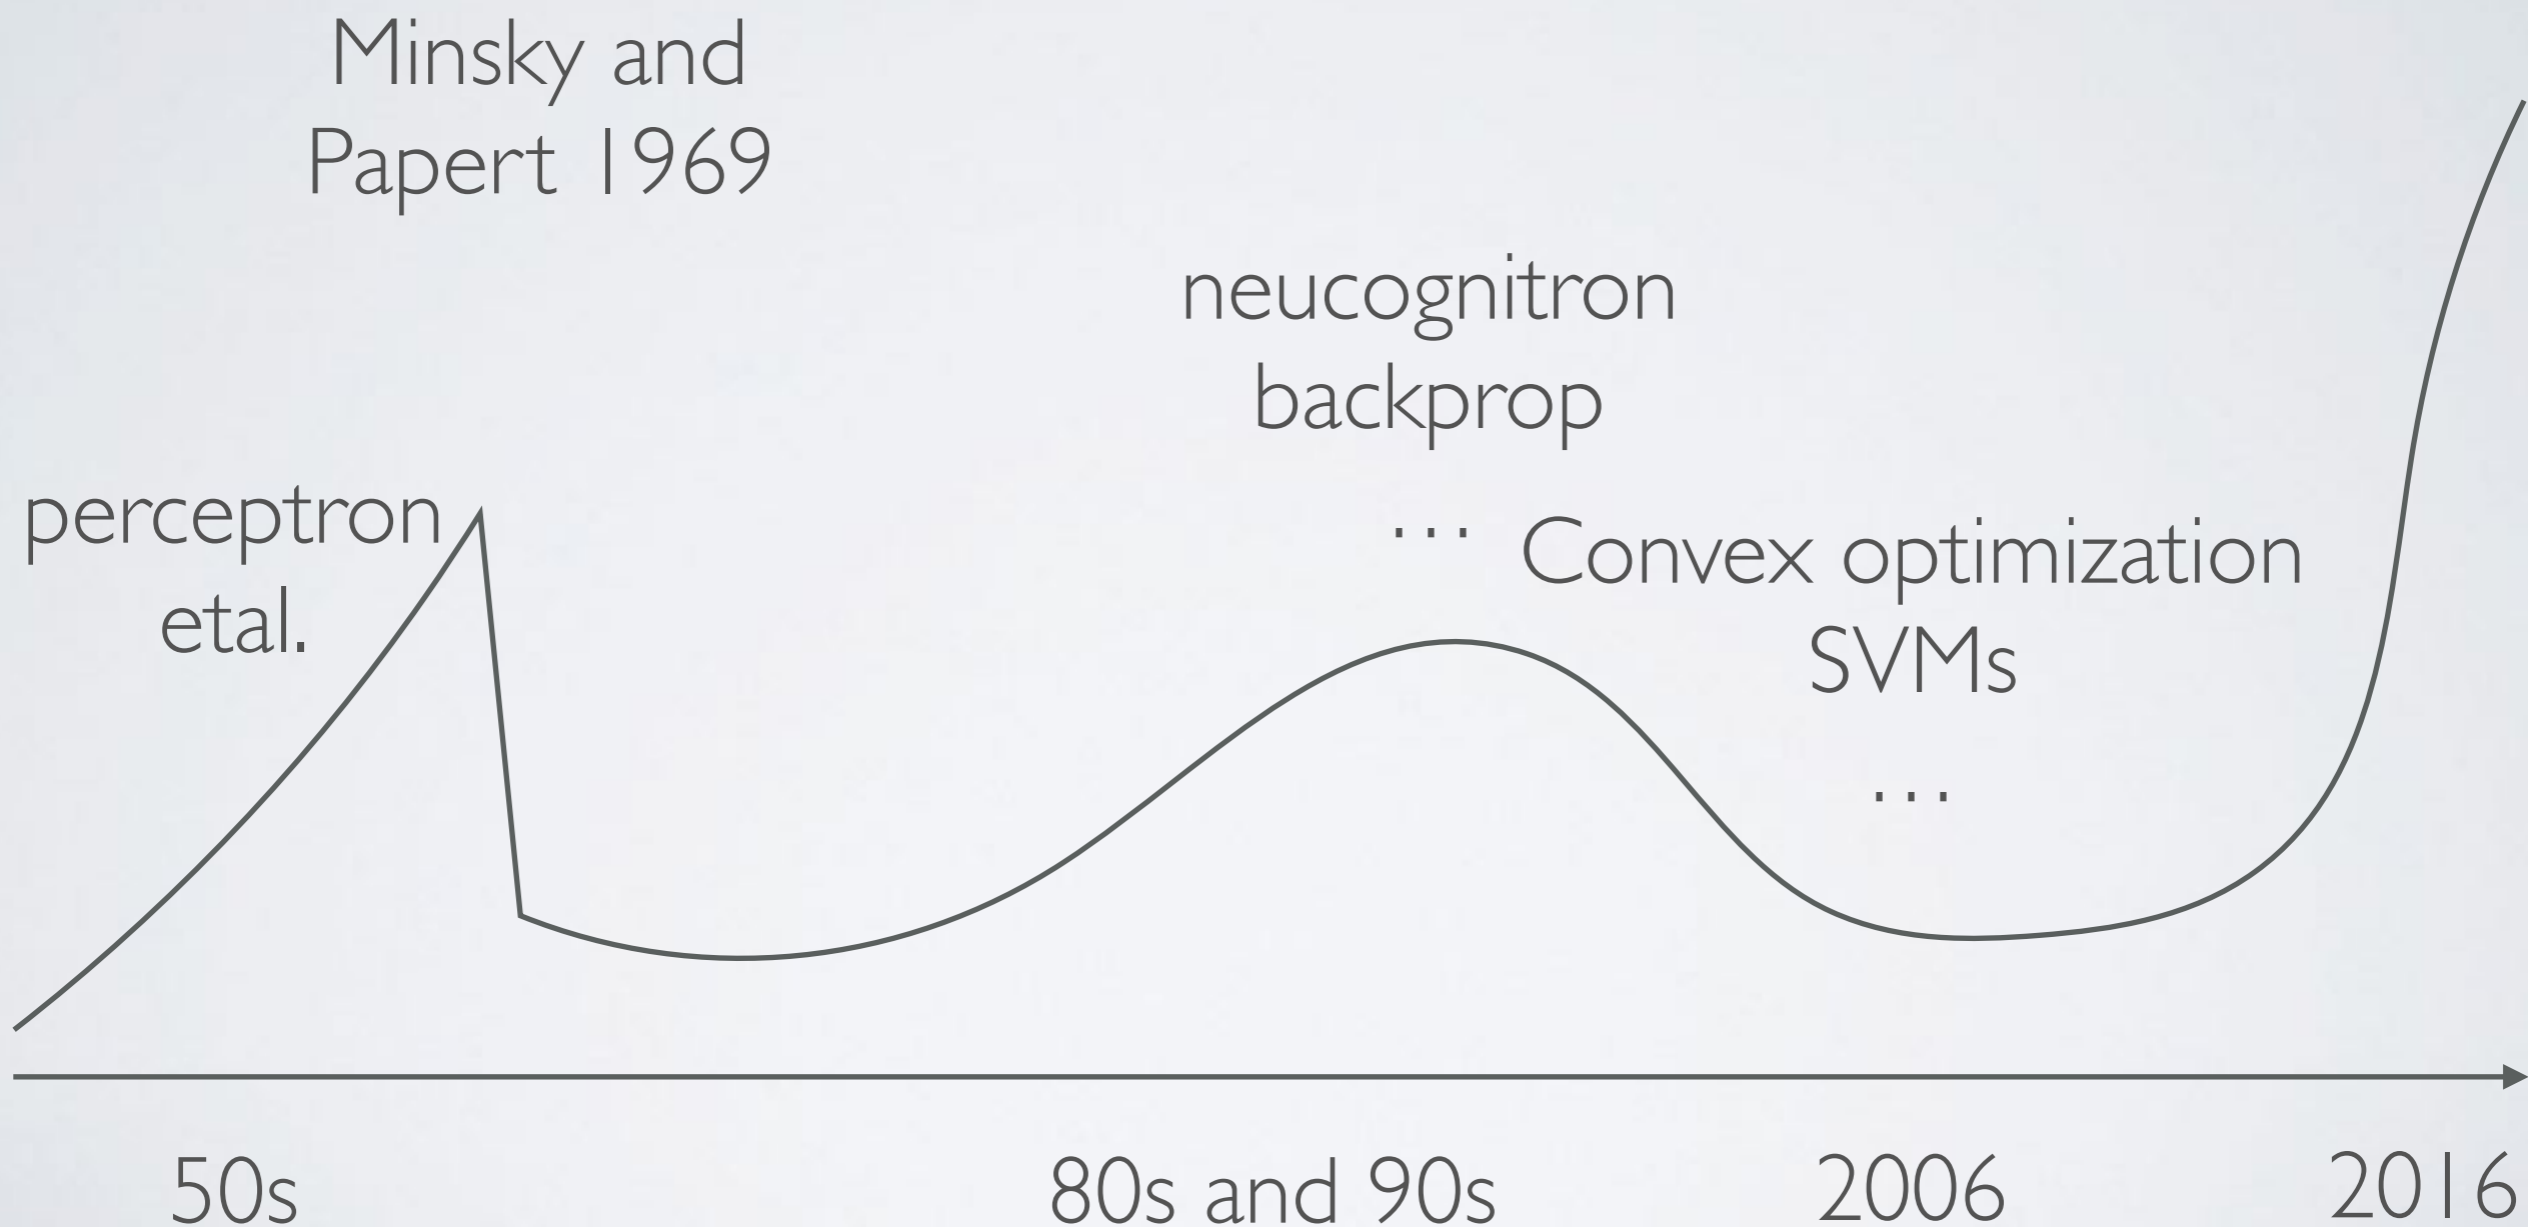
---

doi:10.1038/nature14539


## Deep learning

Yann LeCun<sup>1,2</sup>, Yoshua Bengio<sup>3</sup> & Geoffrey Hinton<sup>4,5</sup>

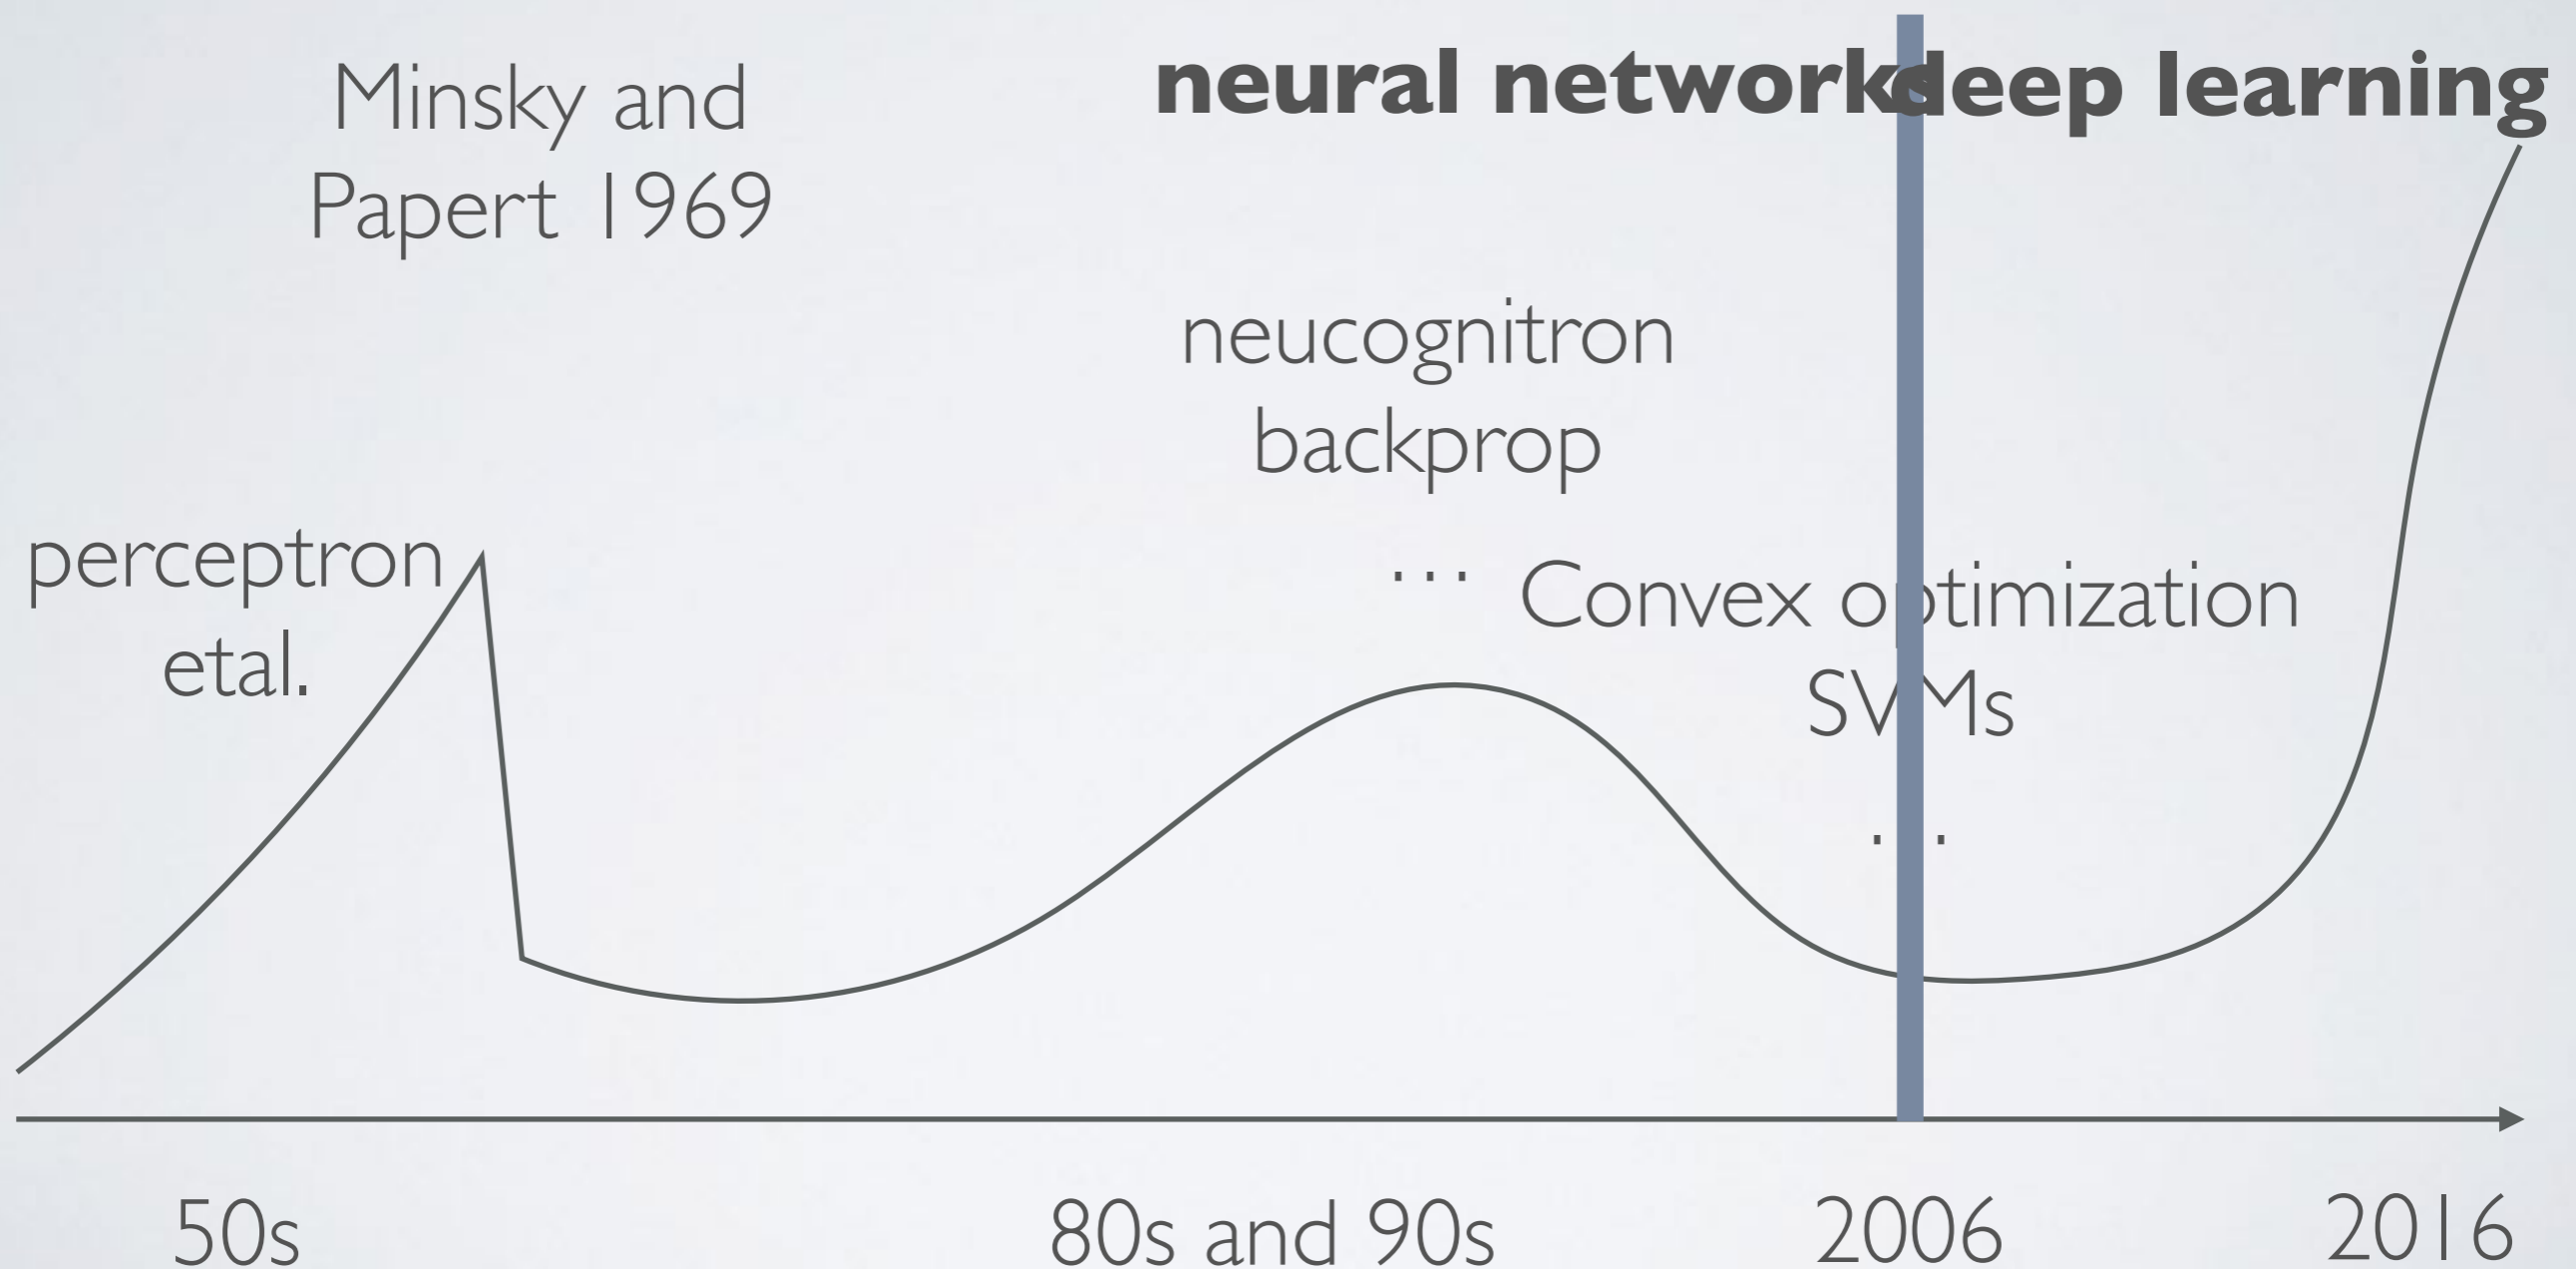
# THE DEEP LEARNING CONSPIRACY



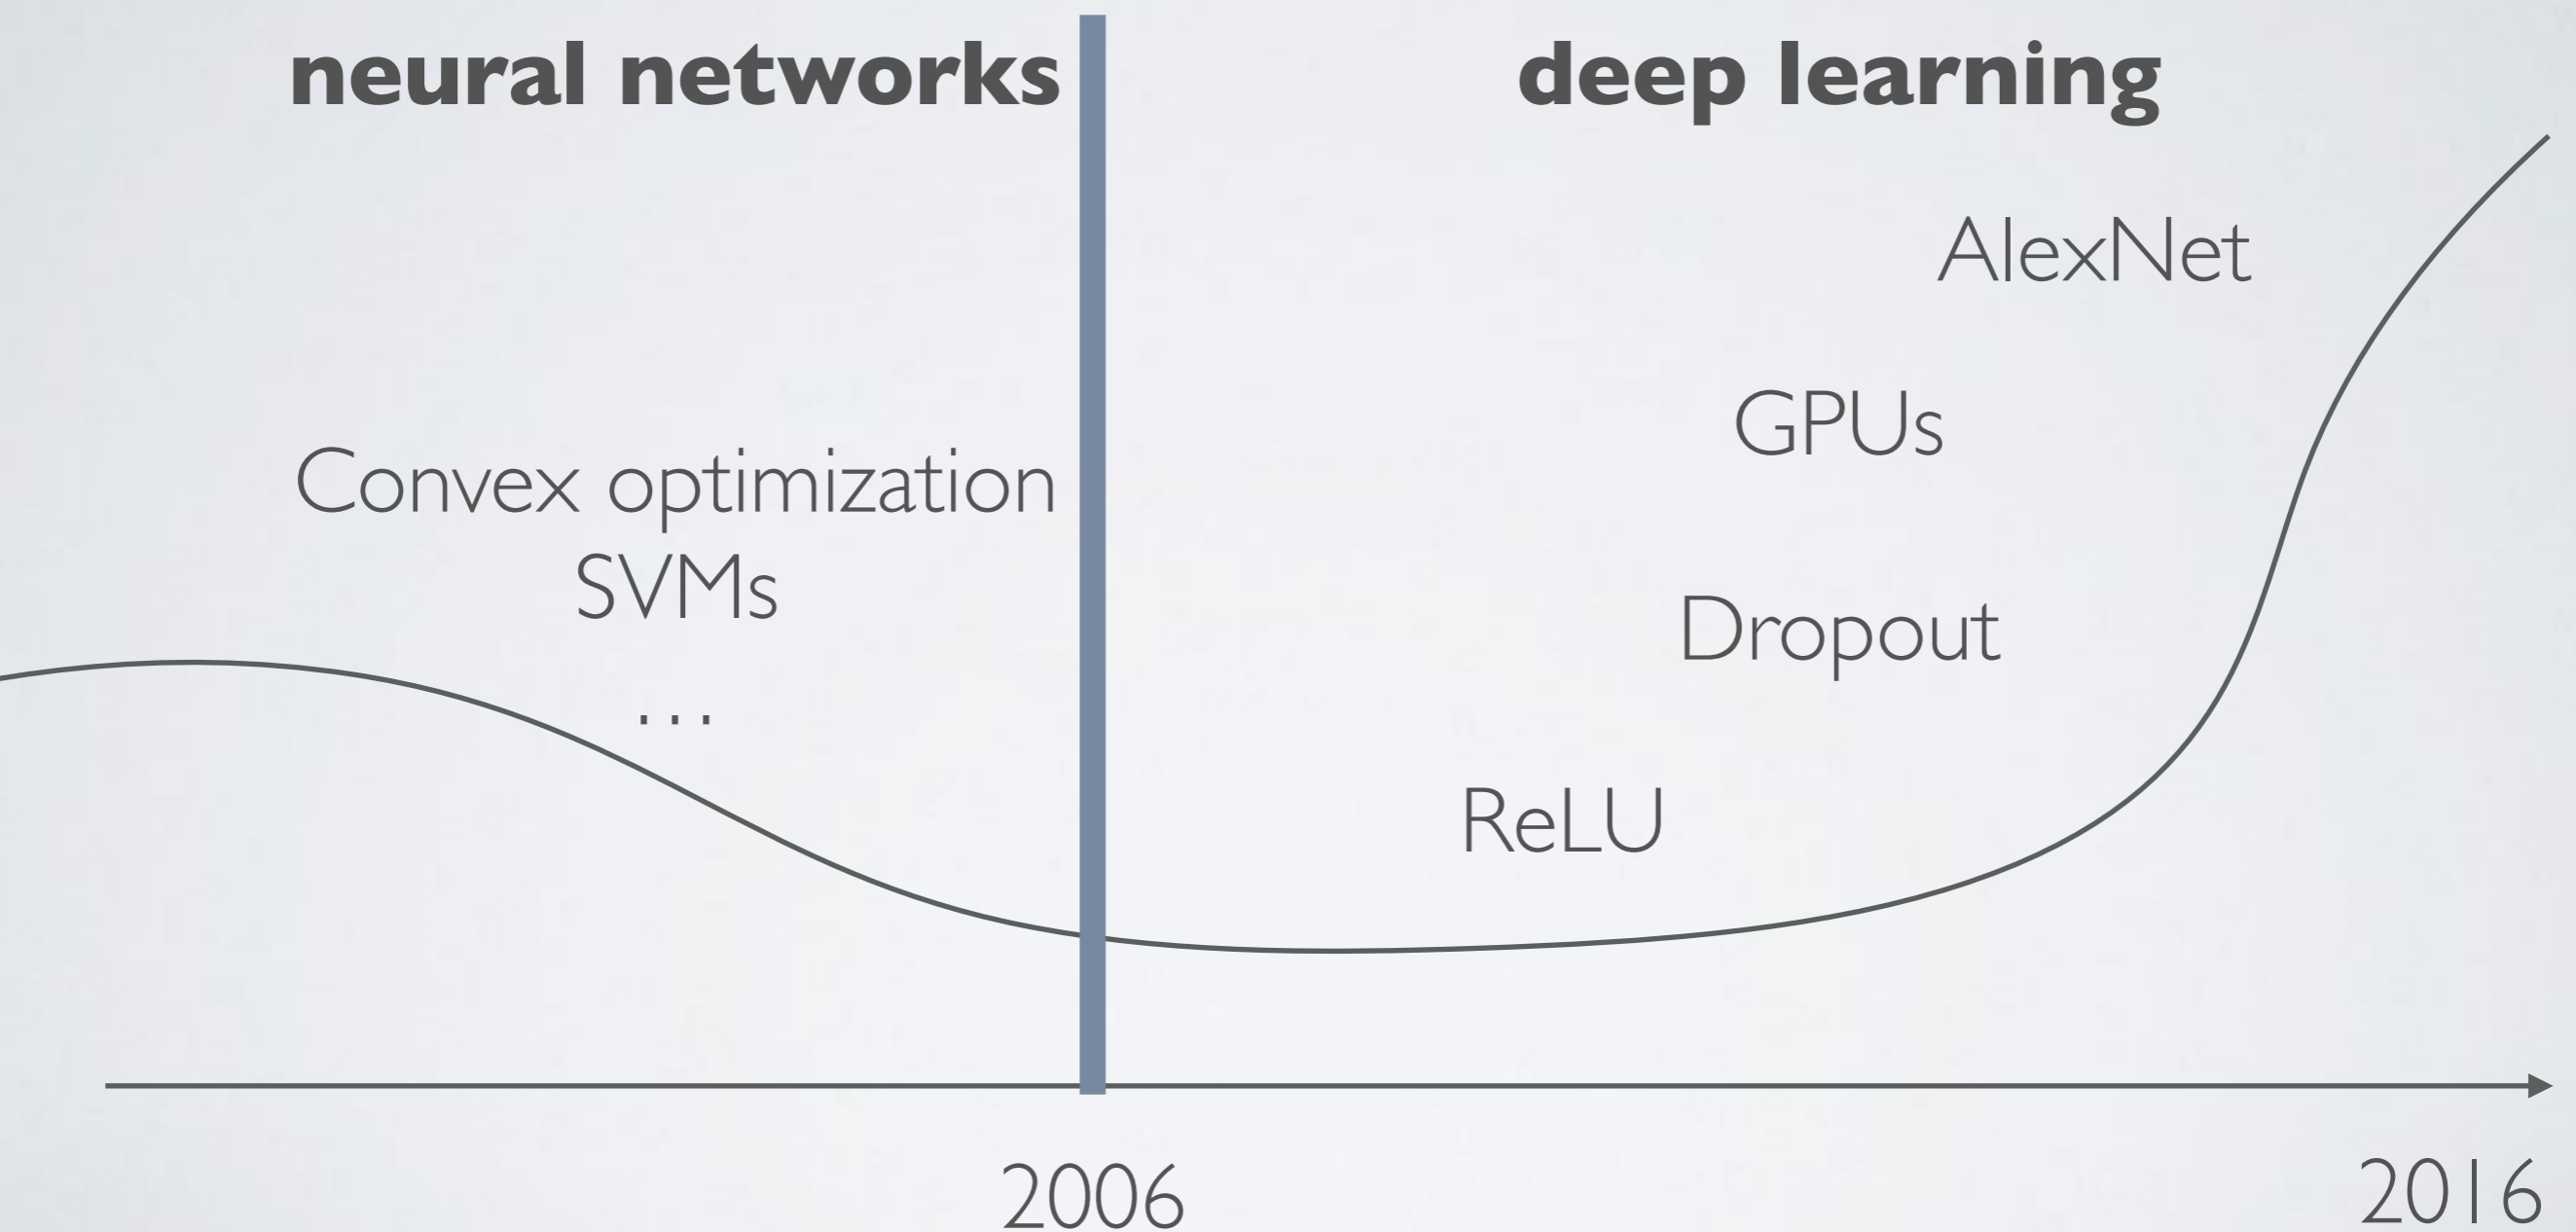
# THE DEEP LEARNING CONSPIRACY



# THE DEEP LEARNING CONSPIRACY



# THE DEEP LEARNING CONSPIRACY



# THE DEEP LEARNING CONSPIRACY

30 days of = x days of






# THE DEEP LEARNING CONSPIRACY

30 c

ars of



# THE DEEP LEARNING CONSPIRACY

30 days of = x years of



# THE DEEP LEARNING CONSPIRACY

30 days of = 26 years of



# THE N-WORD


- Neural
  - Deep Learning
  - try to keep Neuroscience out of this class
  - try to motivate through optimization and ML
    - instead of biology



HISTORY LESSON OVER

# WHAT IS A DEEP NETWORK?

A “differentiable” function  
composed out of multiple layers  
of computation



# TENSORS

- A tensor is a  $d$ -dimensional array
  - A 1-d tensor is a vector
  - A 2-d tensor is a matrix
  - ...
- Tensors are inputs and outputs of layers, as well as their parameters




# LAYER

- Basic unit of computation
- Simple function
  - With parameters





# WHY DO DEEP NETWORKS WORK SO WELL IN PRACTICE?




prairie chicken



# THINGS TO KNOW BEFORE YOU TAKE THE CLASS




training set



test set


# THINGS TO KNOW BEFORE YOU TAKE THE CLASS

linear classifier




# THINGS TO KNOW BEFORE YOU TAKE THE CLASS

logistic regression



# THINGS TO KNOW BEFORE YOU TAKE THE CLASS

training logistic regression



# PREPARATION FOR SECTION

- Windows 10
  - install bash (ubuntu within windows)
- Install python3, pip
- Install pytorch, torchvision
- Bring your laptop to section!