INTRODUCTION TO DATA SCIENCE

This lecture is
based on course by E. Fox and C. Guestrin, Univ of Washington

Visual product recommender

I want to buy new shoes, but...




Too many options online...




Visual product recommender

Text search doesn't help...

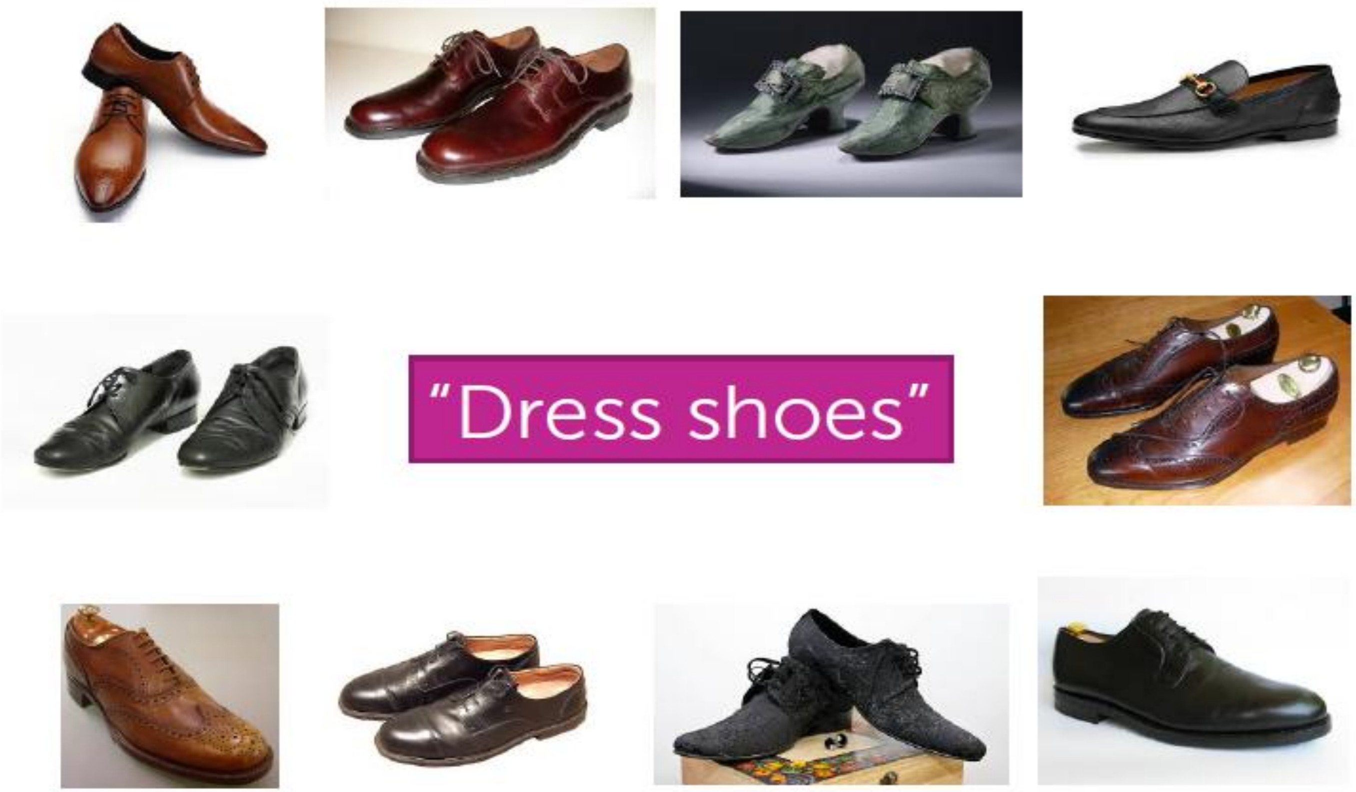

Features are key to machine learning

Classifiers

Goal: revisit classifiers, but using more complex, non-linear features

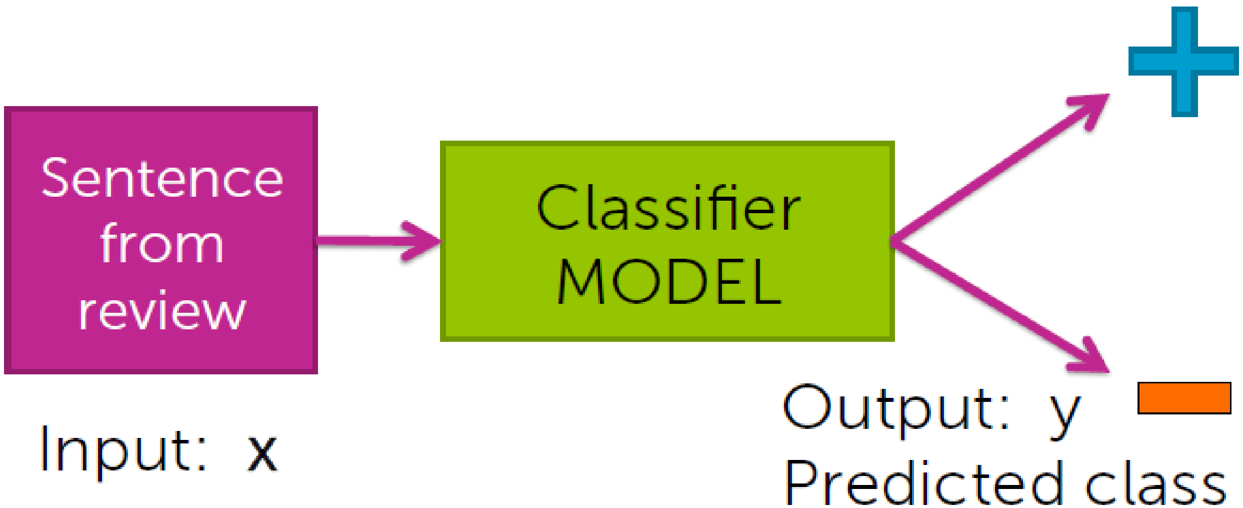

Classifiers

Image classification

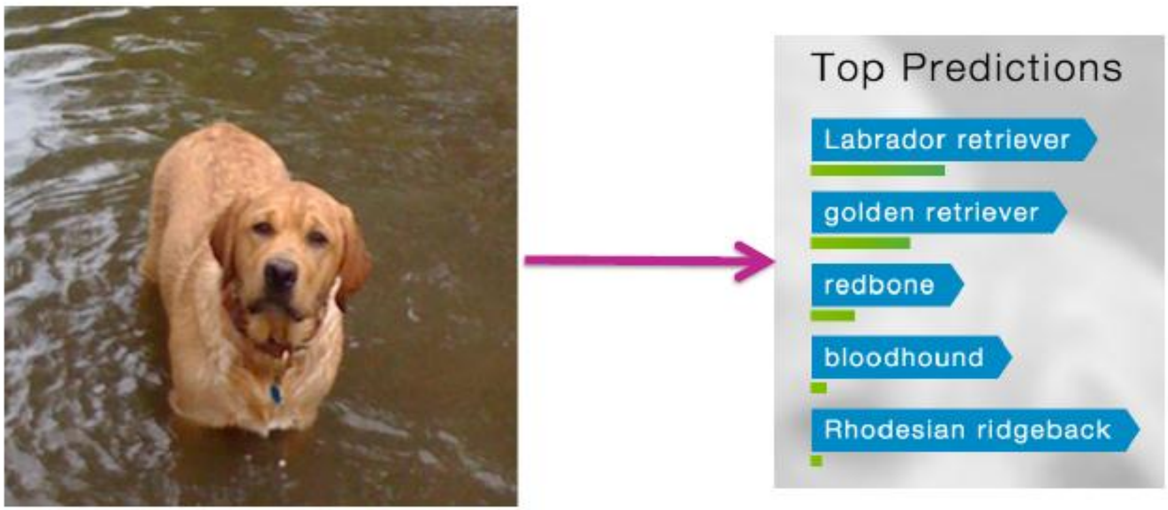

Input: x
Image pixels

Output: y
Predicted object

Neural networks:
learning *very*
non-linear features

Classifiers

Linear classifiers

$$\text{Score}(x) = w_0 + w_1 x_1 + w_2 x_2 + \ldots + w_d x_d$$

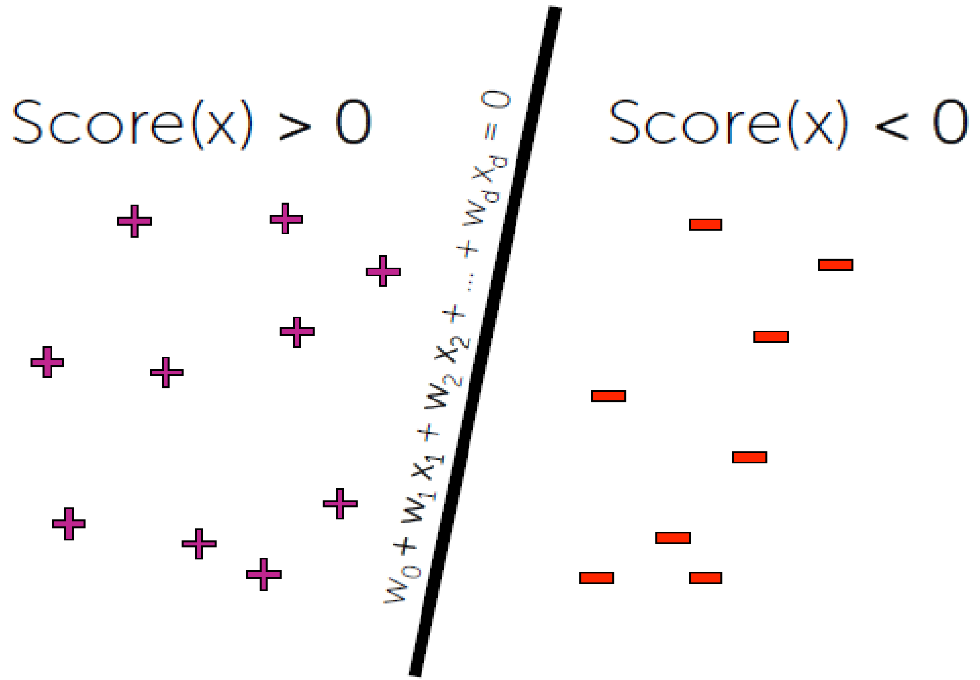

Classifiers

Graph representation of classifier: useful for defining neural networks

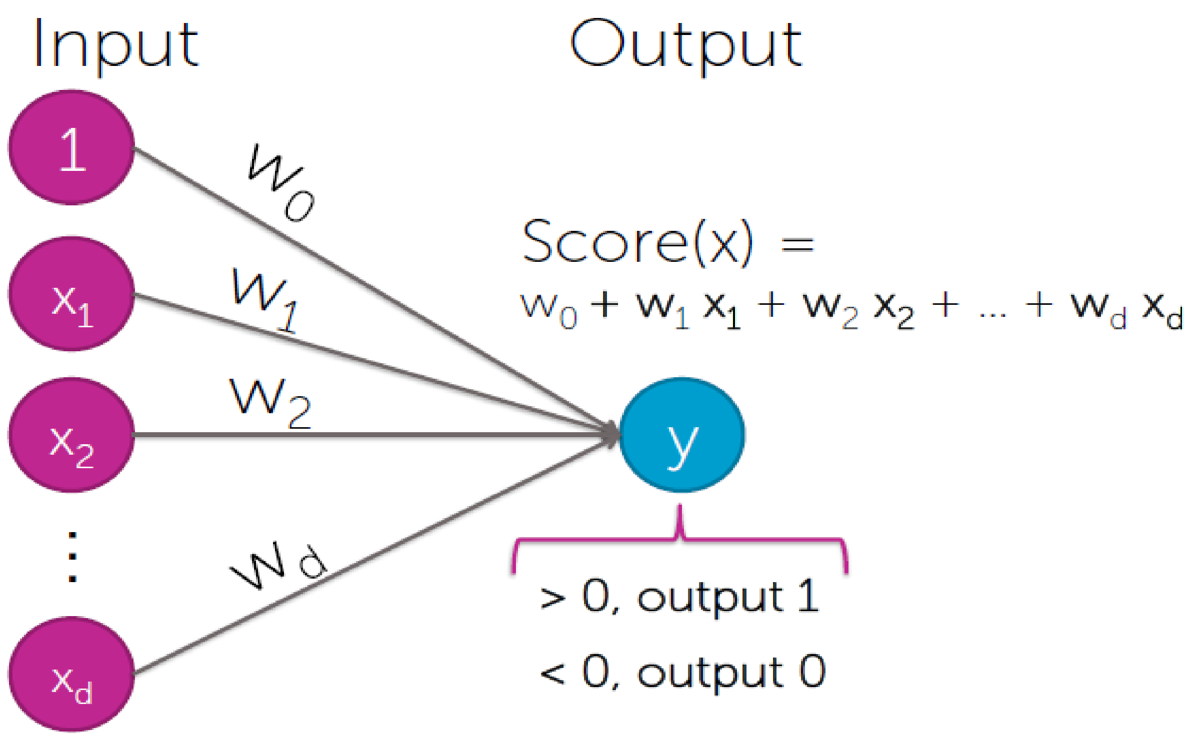

Classifiers

What can a linear classifier represent?

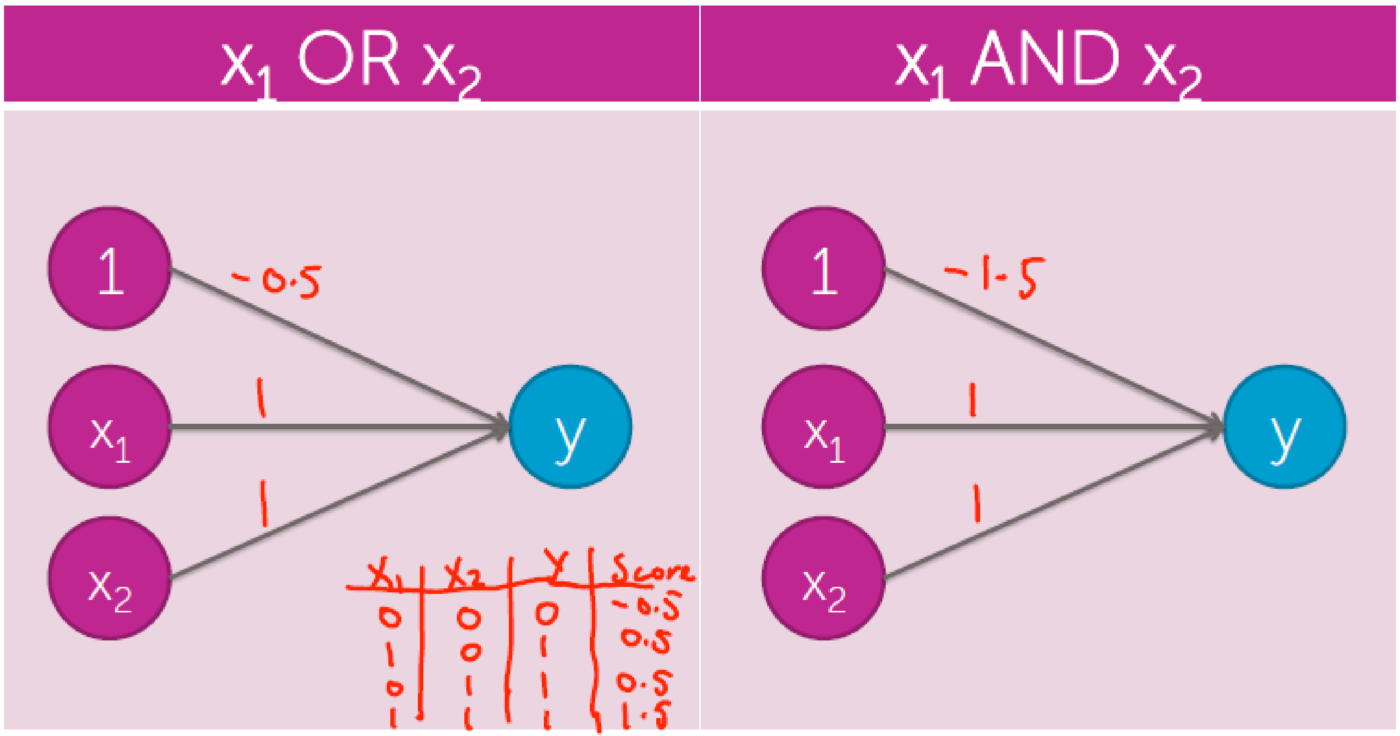

Classifiers

What can't a simple linear classifier represent?

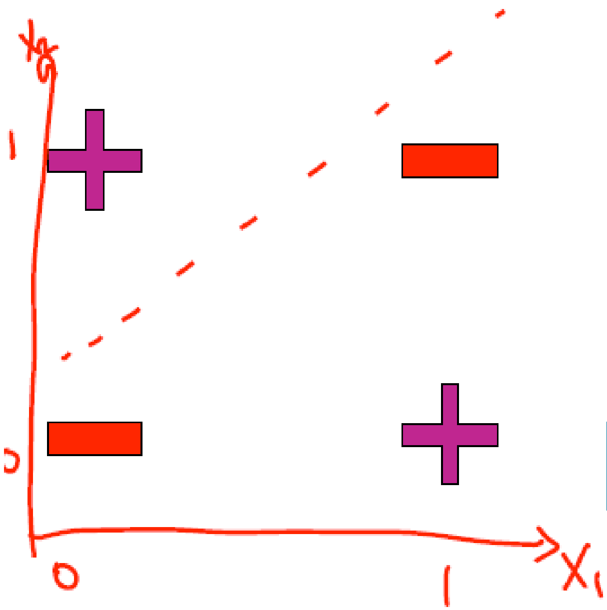

XOR
the counterexample to everything

Need non-linear features

Classifiers

Solving the XOR problem: Adding a layer

$$\text{XOR} = \underbrace{x_1 \text{ AND NOT } x_2}_{z_1} \text{ OR } \underbrace{\text{NOT } x_1 \text{ AND } x_2}_{z_2}$$

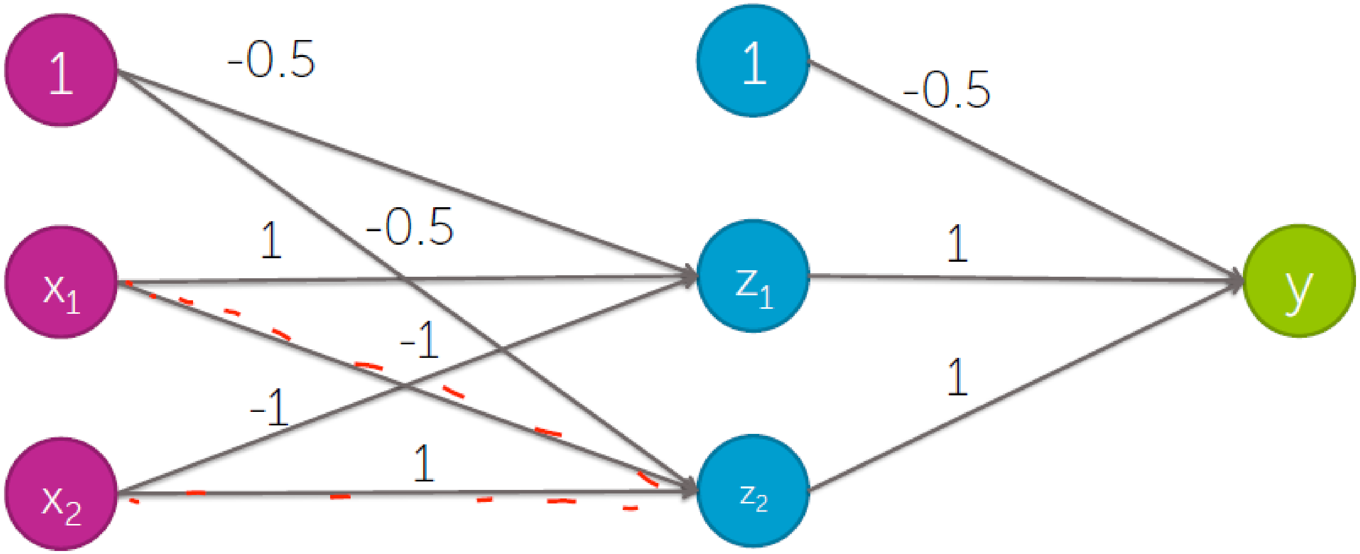

Thresholded to 0 or 1

A neural network

- Layers and layers and layers of linear models and non-linear transformations

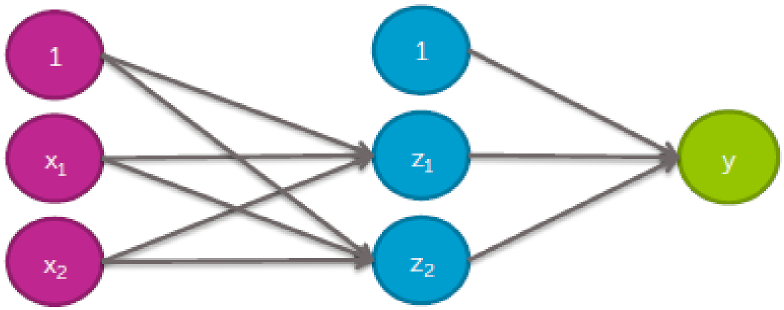

- Around for about 50 years
  - Fell in "disfavor" in 90s
- In last few years, big resurgence
  - Impressive accuracy on several benchmark problems
  - Powered by huge datasets, GPUs, & modeling/learning alg improvements

Application of deep learning to computer vision

Deep Learning

Image features

- Features = local detectors
  - Combined to make prediction
  - (in reality, features are more low-level)

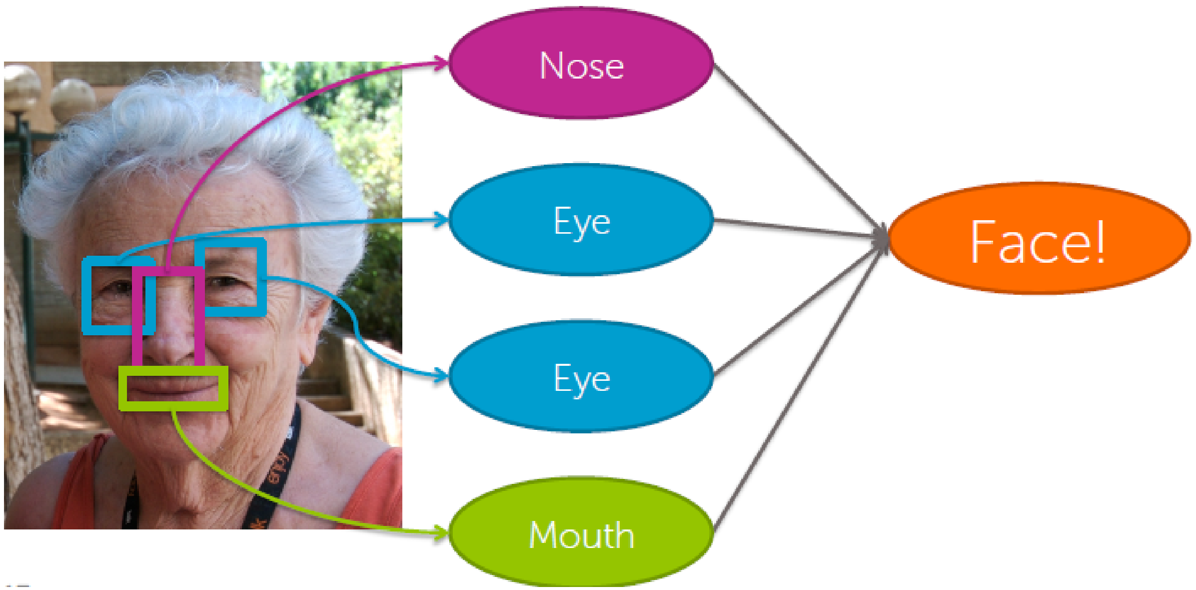

Deep Learning

Typical local detectors look for locally "interesting points" in image

- *Image features*: collections of locally interesting points
  - – Combined to build classifiers

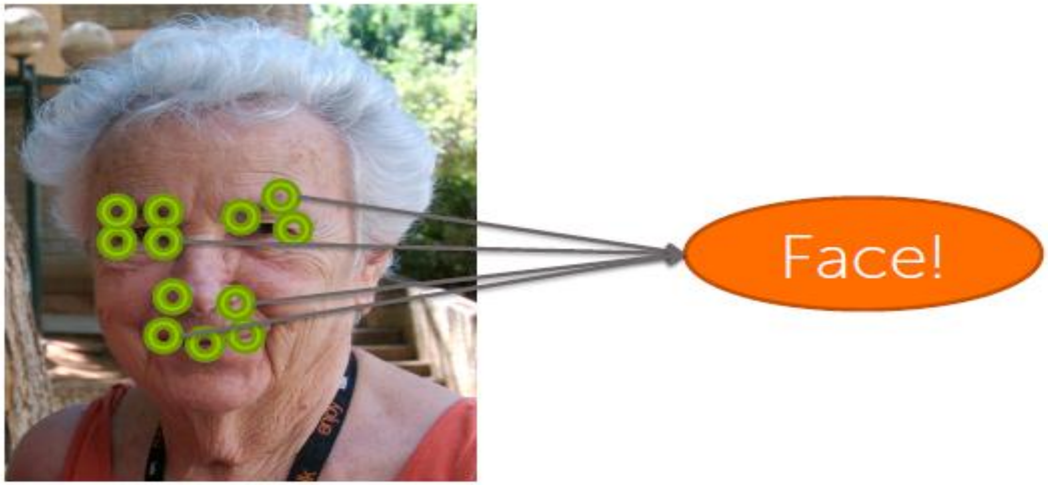

Deep Learning

Many hand created features exist for finding interest points…

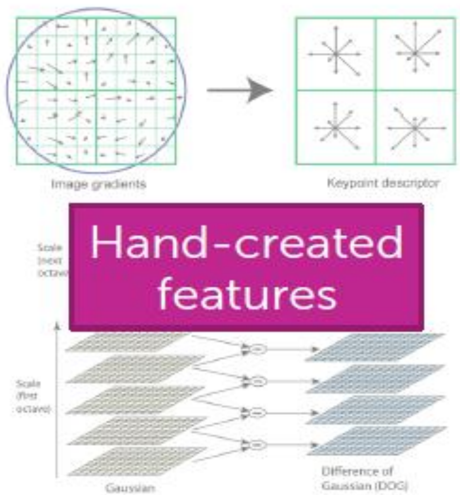

SIFT [Lowe '99]

- *Spin Images* [Johnson & Herbert '99]
- *Textons* [Malik et al. '99]
- *RIFT* [Lazebnik '04]
- *GLOH* [Mikolajczyk & Schmid '05]
- *HoG* [Dalal & Triggs '05]
- …

… but very painful to design

Standard image classification approach

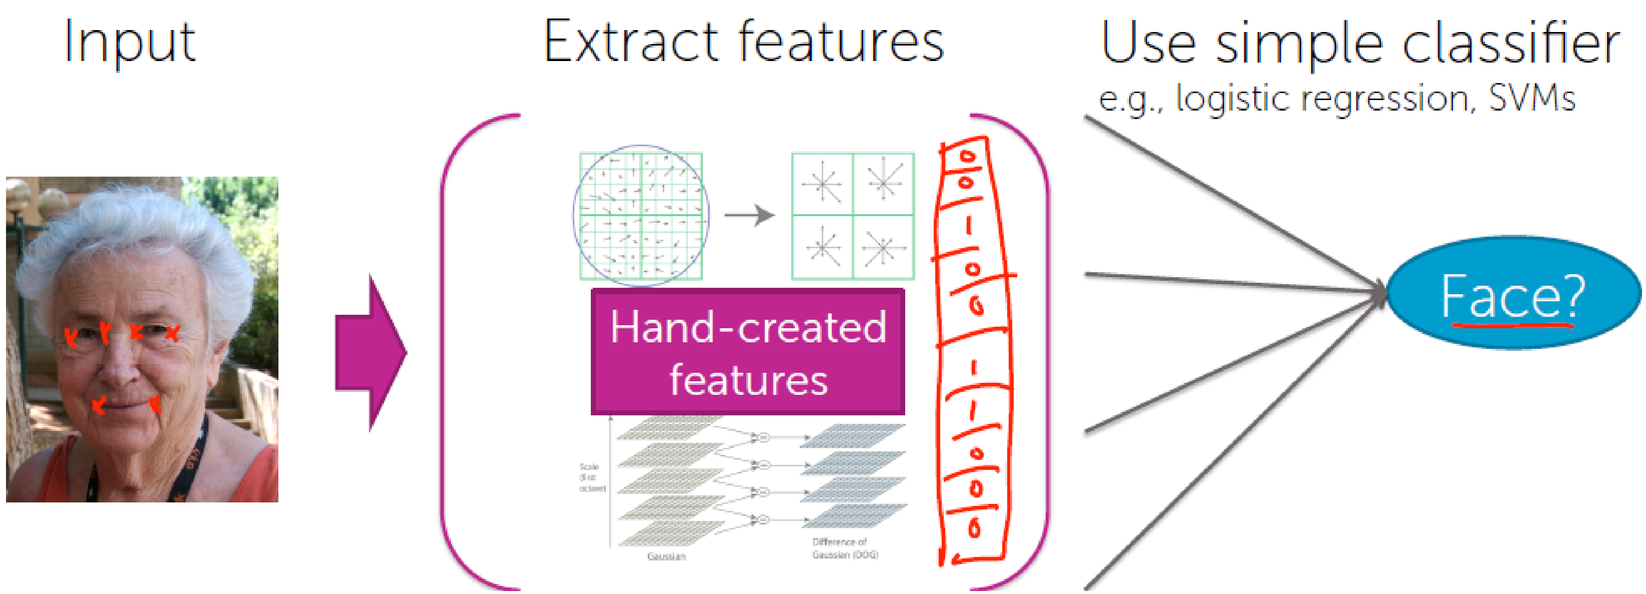

Deep Learning

Deep learning:
implicitly learns features

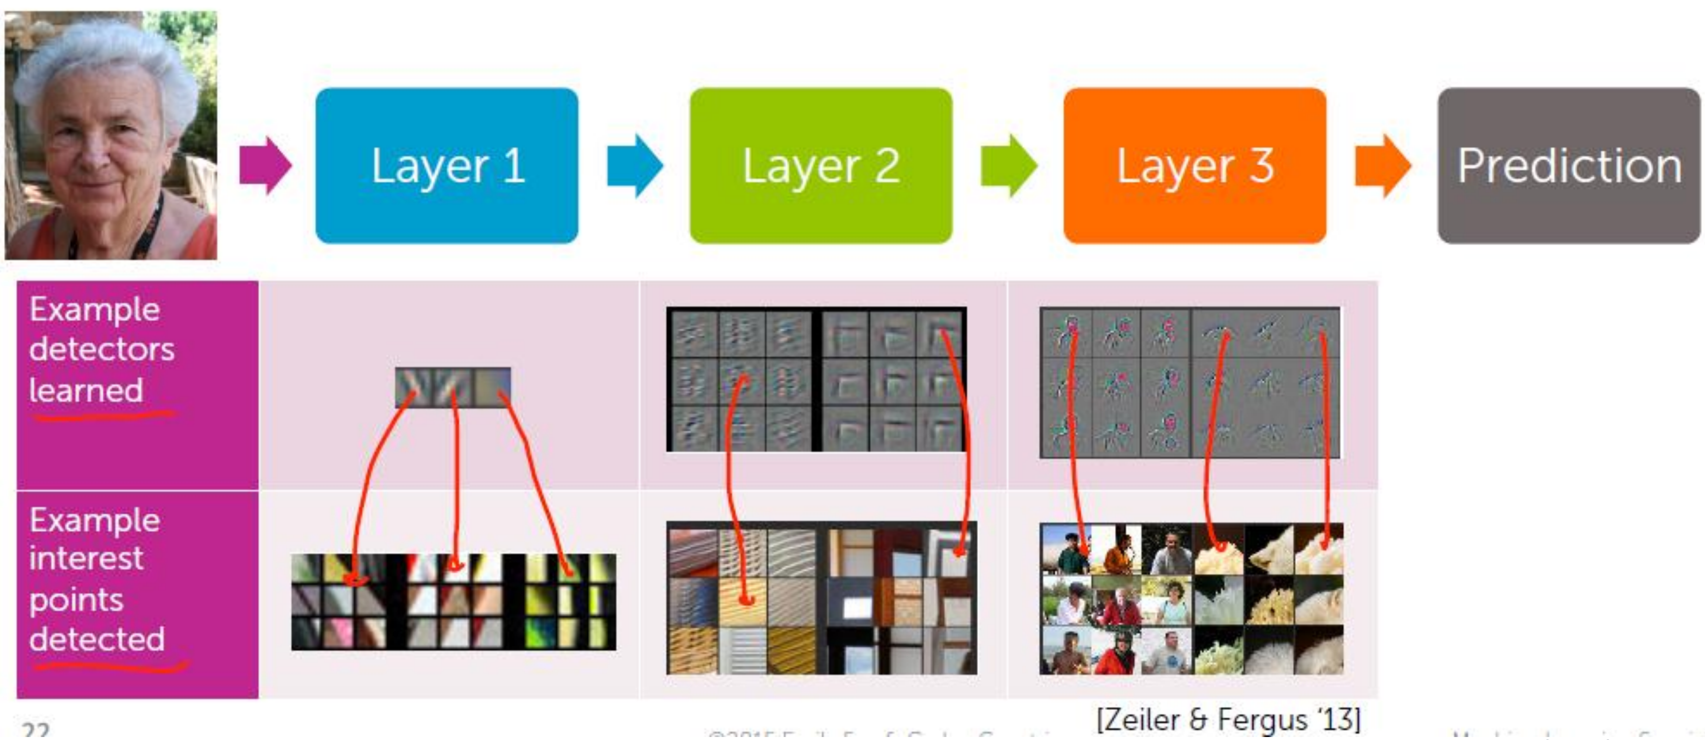

Deep Learning performance

Sample results using deep neural networks

- German traffic sign recognition benchmark
  - 99.5% accuracy (IDSIA team)

- House number recognition
  - 97.8% accuracy per character [Goodfellow et al. '13]

Deep Learning performance

ImageNet 2012 competition:
1.2M training images, 1000 categories

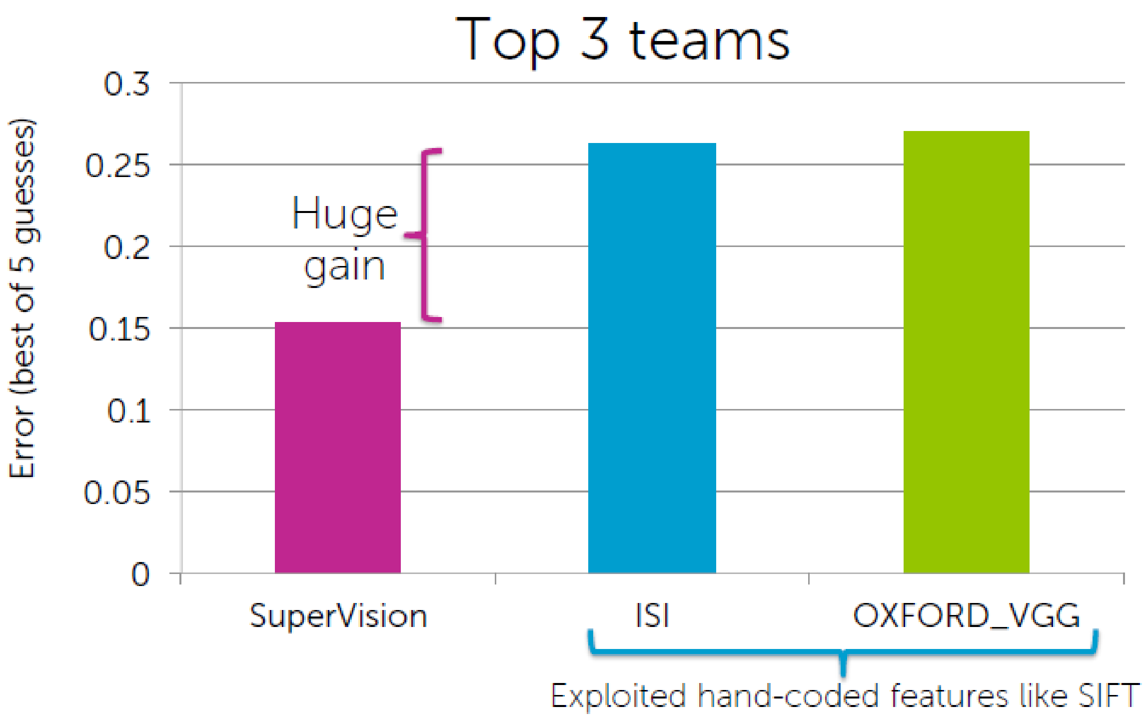

Deep Learning performance

ImageNet 2012 competition:
1.2M training images, 1000 categories

Winning entry: SuperVision
8 layers, 60M parameters [Krizhevsky et al. '12]

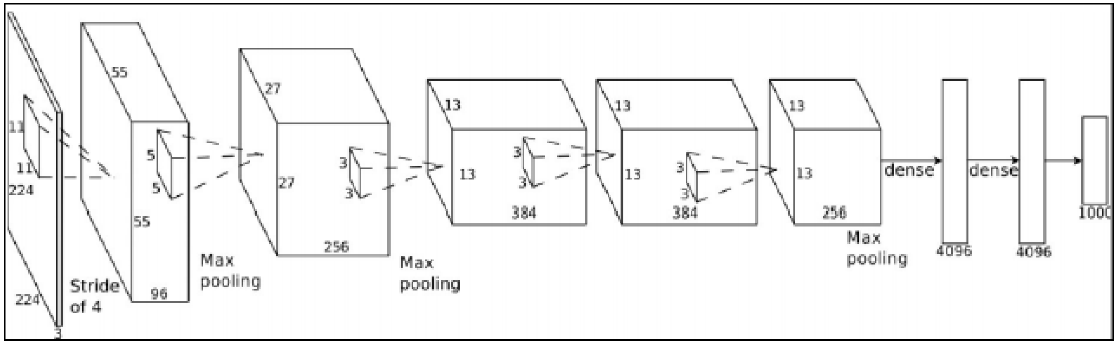

Achieving these amazing results required:
- New learning algorithms
- GPU implementation

Deep learning in computer vision

Scene parsing with deep learning

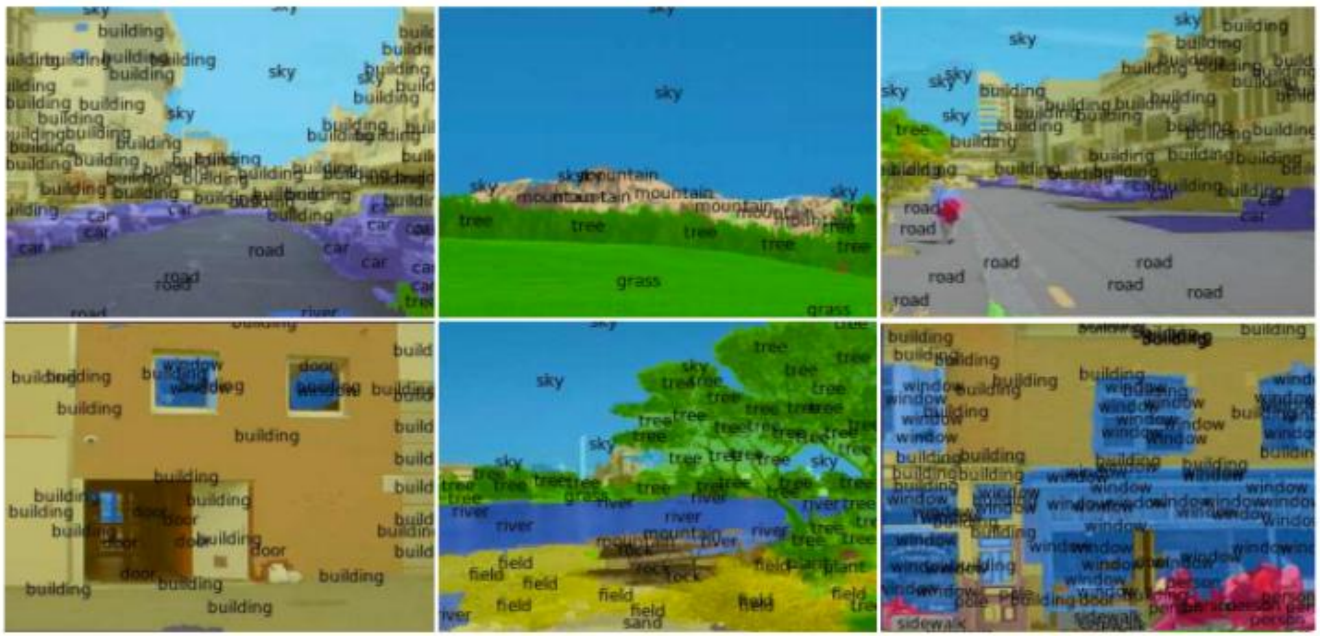

[Farabet et al. '13]

Deep learning in computer vision

Retrieving similar images

Challenges

Deep learning score card

Year 2015

Pros

- Enables learning of features rather than hand tuning
- Impressive performance gains
  - Computer vision
  - Speech recognition
  - Some text analysis
- Potential for more impact

Deep learning workflow

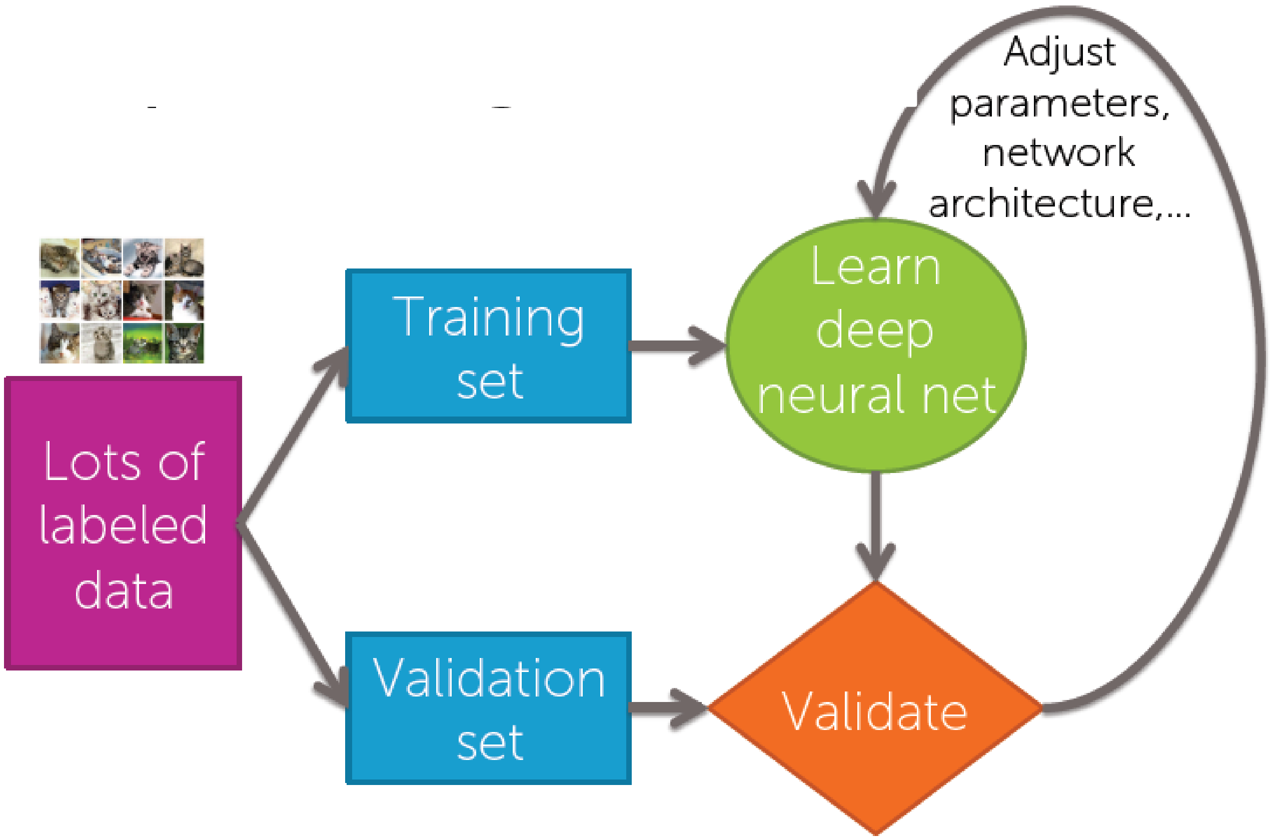

Deep learning workflow

Many tricks needed to work well...

Different types of layers, connections,... needed for high accuracy

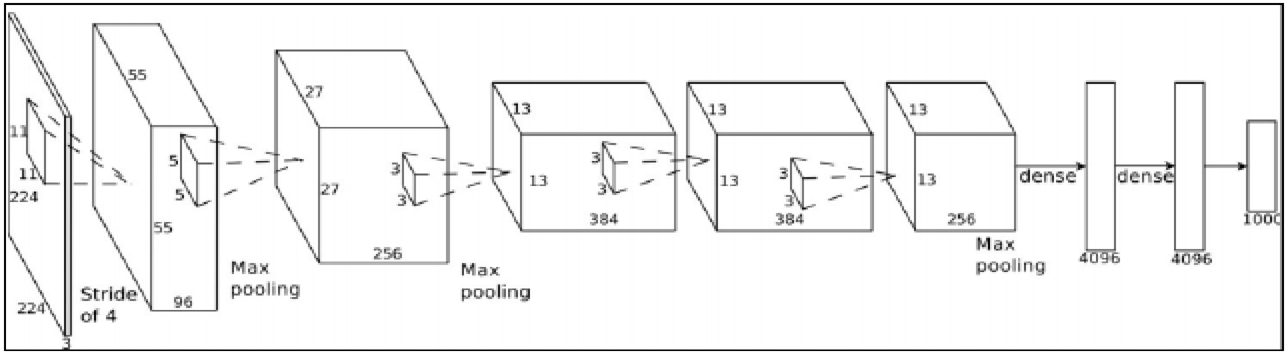

[Krizhevsky et al. '12]

Challenges

Deep learning score card

Year 2015

Pros

- Enables learning of features rather than hand tuning
- Impressive performance gains
  - Computer vision
  - Speech recognition
  - Some text analysis
- Potential for more impact

Cons

- Requires a lot of data for high accuracy
- Computationally really expensive
- Extremely hard to tune
  - Choice of architecture
  - Parameter types
  - Hyperparameters
  - Learning algorithm
  - ...

Computational cost+ so many choices
=
incredibly hard to tune

Transfer learning

Transfer learning: *Use data from one task to help learn on another*

Old idea, explored for deep learning by Donahue et al. '14 & others

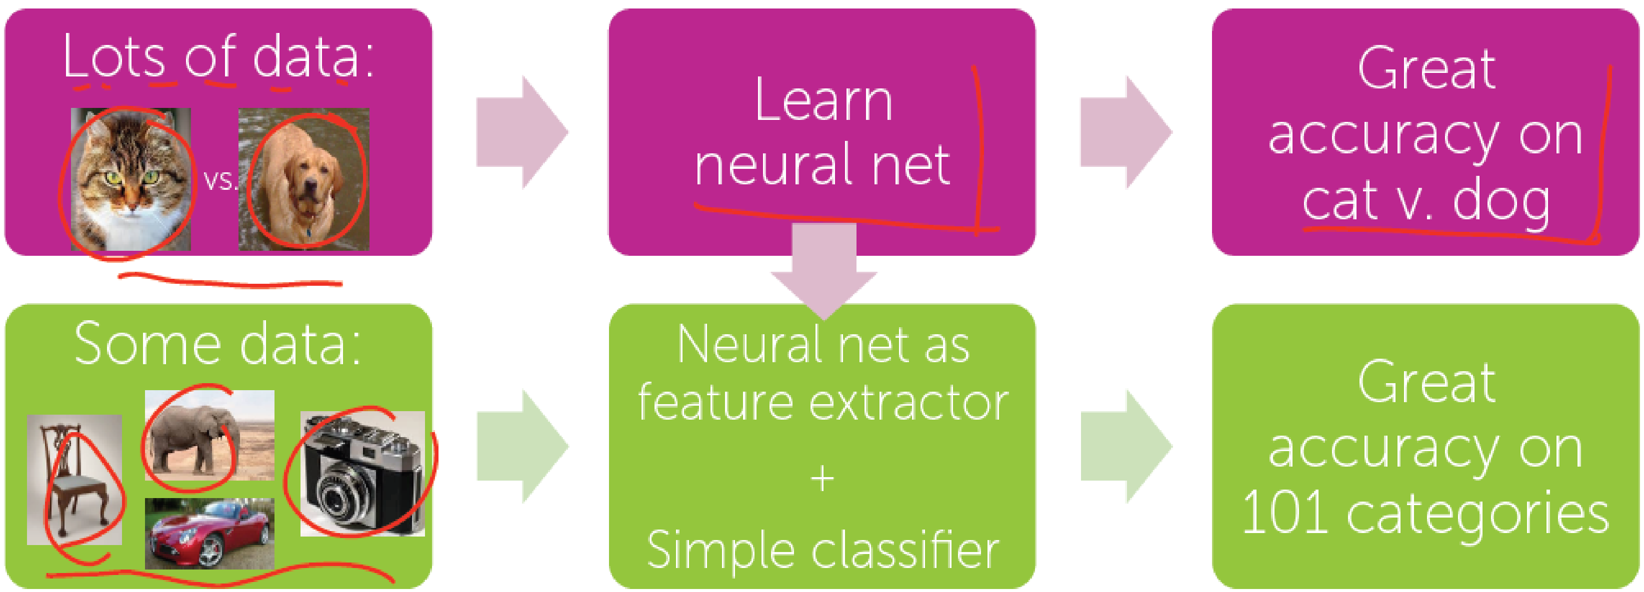

Transfer learning

What's learned in a neural net

vs.

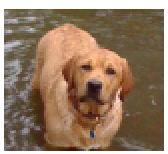

Neural net trained for Task 1: cat vs. dog

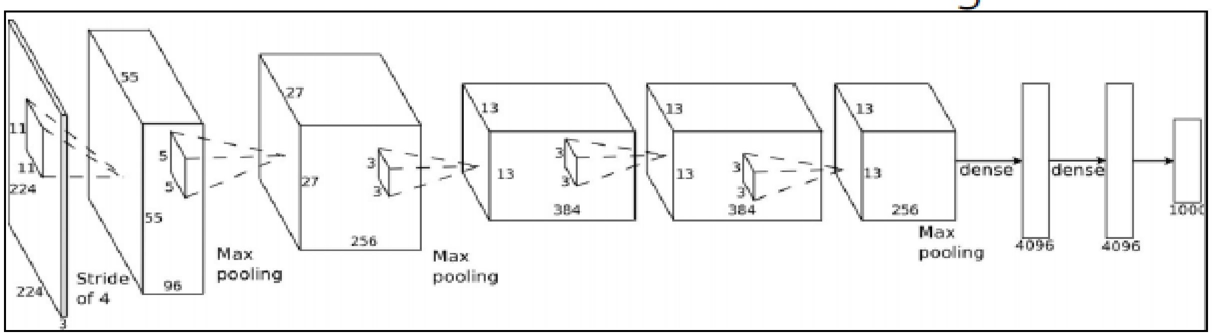

More generic
Can be used as feature extractor

Very specific
to Task 1
Should be ignored
for other tasks

Transfer learning

Transfer learning in more detail…

For Task 2, predicting 101 categories,
learn only end part of neural net

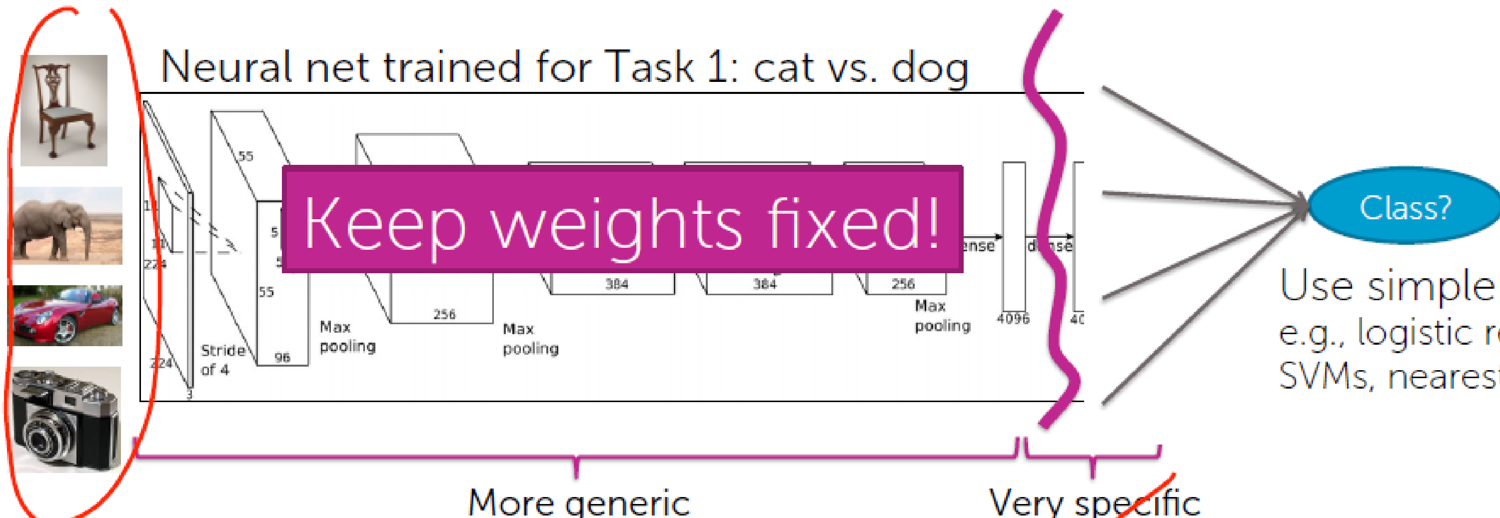

Use simple classifier
e.g., logistic regression,
SVMs, nearest neighbor,…

More generic
Can be used as feature extractor

Very specific
to Task 1
Should be ignored
for other tasks

Transfer learning

Careful where you cut:
latter layers may be too task specific

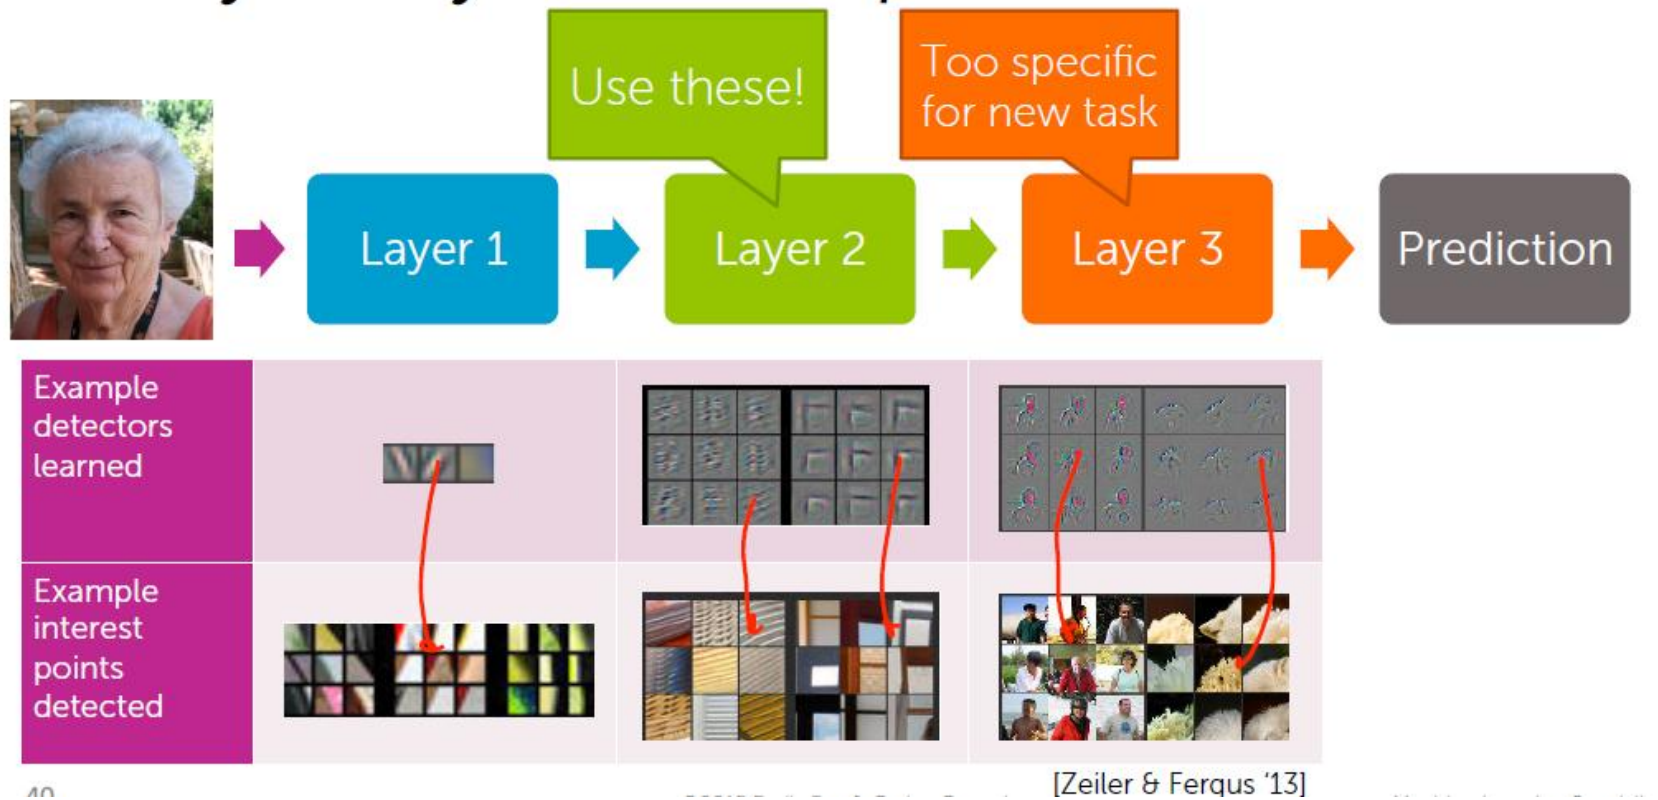

Transfer learning

Transfer learning with deep features workflow

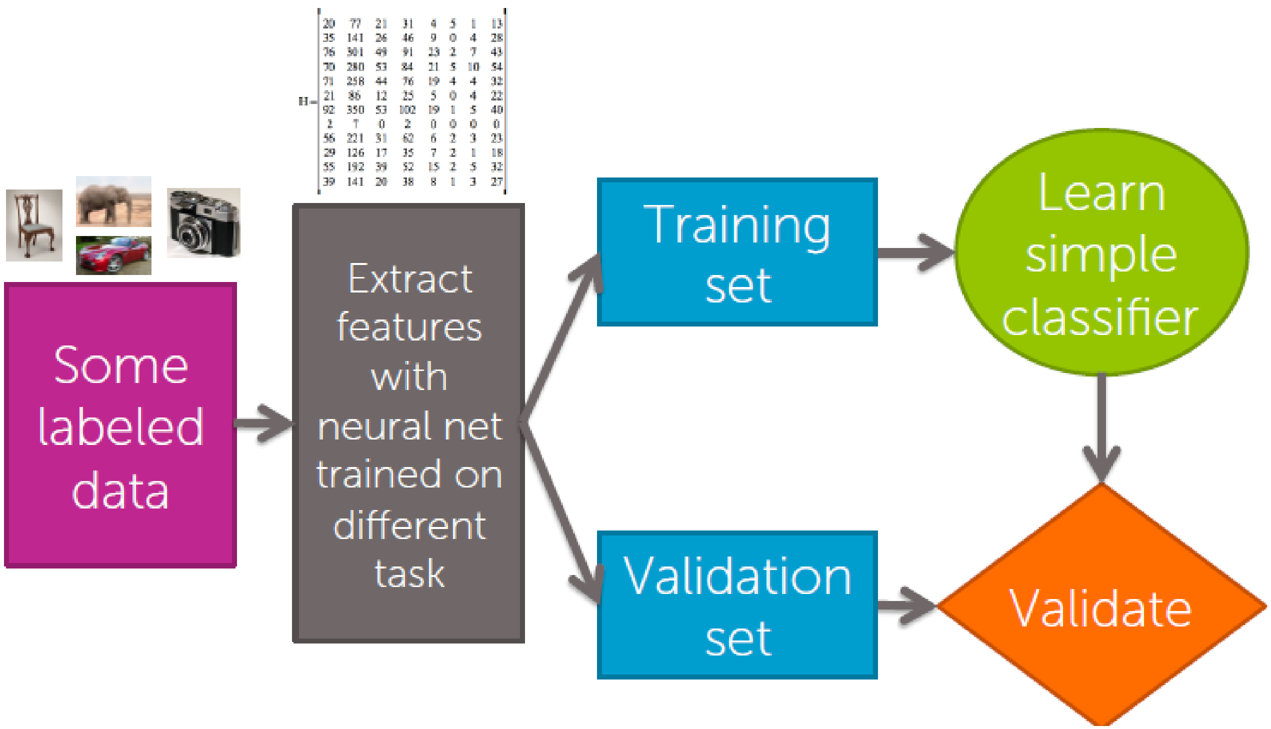

What you can do now …

- Describe multi-layer neural network models
- Interpret the role of features as local detectors in computer vision
- Relate neural networks to hand-crafted image features
- Describe some settings where deep learning achieves significant performance boosts
- State the pros & cons of deep learning model
- Apply the notion of transfer learning
- Use neural network models trained in one domain as features for building a model in another domain
- Build an image retrieval tool using deep features