

Deep Learning

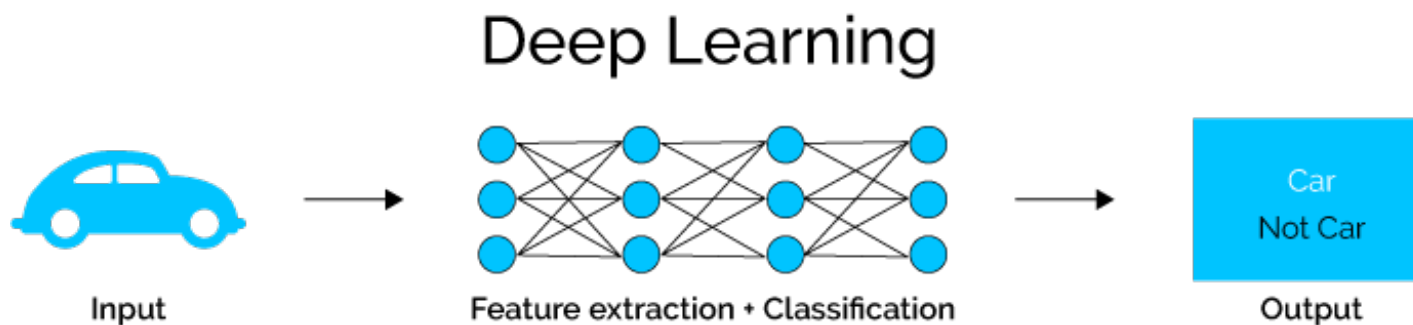

Nima Sedaghat
University of Freiburg

Outline

- Classical Machine Learning vs. Deep Learning
- The Rapid Evolution of Deep Learning
- Deep Learning and Types of Supervision

Classical Machine Learning vs. Deep Learning


From Hand-Crafted Features to Representation Learning




Why didn't we do it before?

We needed to go deeper: more parameters, more sophisticated features

"Non-deep" feedforward
neural network





Deep neural network



It would take forever to tell a cat from a dog!

What made it possible?



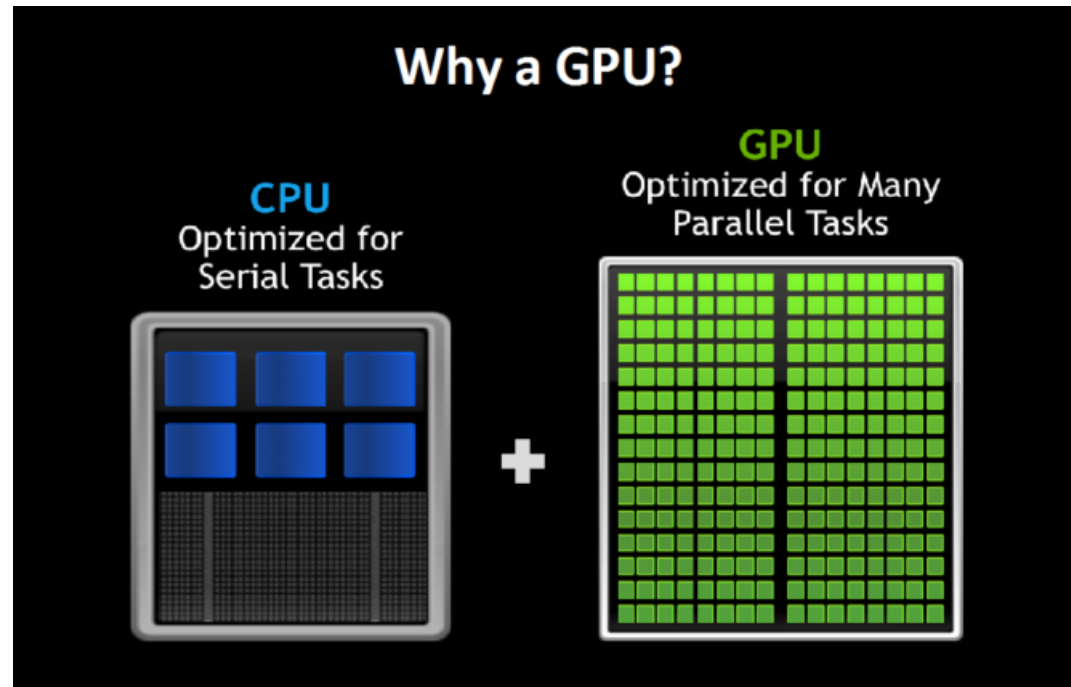
Convolutional Layers

GPU: Graphical Processing Unit




<https://www.wikihow.com/Connect-an-External-TV-Tuner-Card-to-a-Desktop>

GPU: Graphical Processing Unit




Convolutional Layers



<https://www.analyticsindiamag.com/convolutional-neural-network-image-classification-overview>


Why Convolutional Layers help?

“Shared Weights”





Seongwon Hwang, “Convolutional Neural Network (CNN) presentation from theory to code in Theano”


Deep ConvNets



Hierarchical Representation Learning



On the Importance of Data



On the Importance of Data


Small Dataset → over-fitting



It's a cat

The Rapid Evolution of Deep Learning

Classification



<https://becominghuman.ai/building-an-image-classifier-using-deep-learning-in-python-totally-from-a-beginners-perspective-be8dbaf22dd8>

Classification

- Question: examples from other domains?
(E.g. Astronomy)



<https://github.com/EdwardJKim/dl4astro/>

Regression




Sedaghat, Brox, "Unsupervised Generation of a Viewpoint Annotated Car Dataset from Videos"


Regression

- Question: What examples can you think of in Sky/Earth?

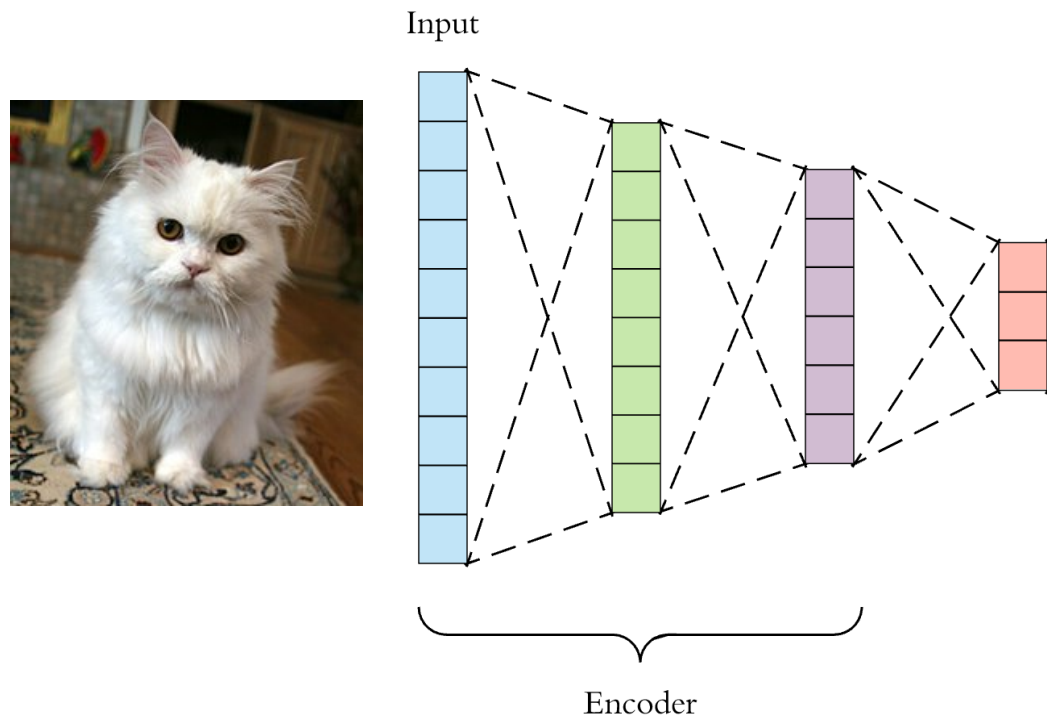
Encoder-Decoder Networks




E.g. Auto-Encoder



Why an Auto-Encoder?




E.g. De-noising




Question: How would you train this?


E.g. Segmentation



E.g. Optical Flow (Motion) Estimation




Generative Adversarial Network (GAN)



<https://towardsdatascience.com/understanding-generative-adversarial-networks-4dafc963f2ef>

GAN Use-case example



Antreas Antoniou et al., "Data Augmentation Generative Adversarial Networks"

Deep Learning and Types of Supervision

Deep Learning and Types of Supervision

- Fully Supervised Learning
- Unsupervised/Self-Supervised Learning
- Reinforcement Learning
- (Imitation Learning)

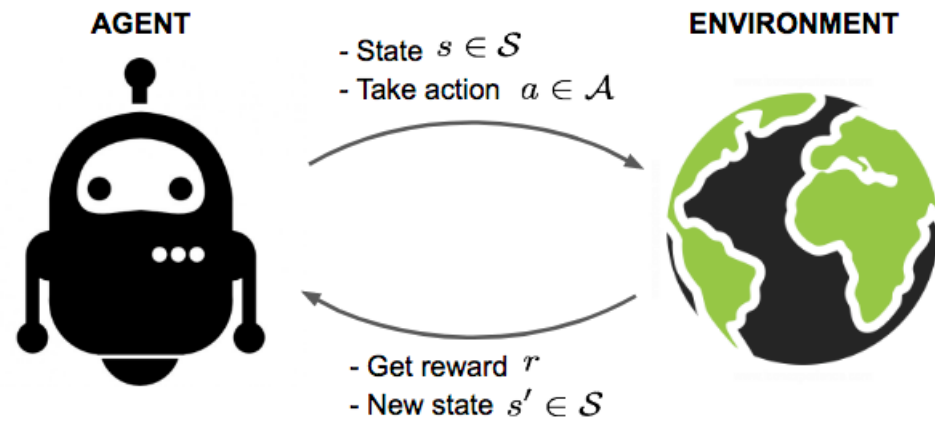
Fully Supervised Learning

- During Training: Input AND Output
Labels
Annotations



Unsupervised/Self-supervised Learning

- No Labels
- No “unsupervised task” such as clustering
- Use self-supervision instead
 - Auto-encoder
 - Denoising
 - Jigsaw-puzzle
 - ...


Reinforcement Learning




Reinforcement Learning



Covariate Shift and Domain Adaptation




Covariate Shift and Domain Adaptation



Hybrid Learning of Optical Flow and Next Frame Prediction to Boost Optical Flow in the Wild.
Nima Sedaghat, Zolfaghari M., Brox T.


Probabilistic Networks





$$p(w|x, y) = \frac{p(x, y|w)p(w)}{\int p(y|x, w)p(w)dw}$$

Application Example: **TransiNet**

Reference Image



Science Image



Ideal output

