data&modelingsciences

Customizing and Assessing Deep Learning for Specific Tasks

Amir Tavanaei, Matthew Barker, William Brenneman

Data Science and AI Group

Fall Technical Conference (October 2018)

About Me

- Amirhossein Tavanaei (Amir)
  - R&D Data Scientist at Procter and Gamble,
  - Dept. of Data Science and Artificial Intelligence.
- Education
  - Computer Science, PhD
  - AI, Machine Learning, Deep Learning
  - Bio-inspired Spiking Networks
- Hobby
  - Hiking
  - Movies

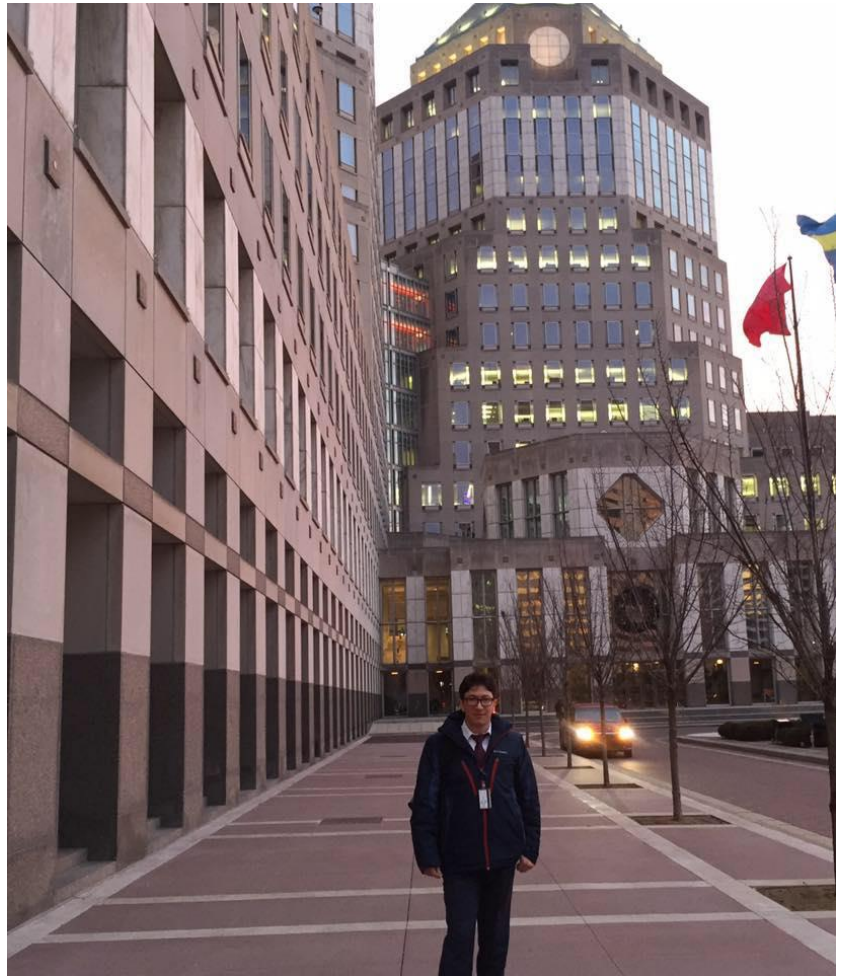

Agenda

- Introduction to neural networks and deep learning
- Deep learning architectures for different problems
- Data variations and requirements
- Deep learning customization
- Recent examples
- Summary

Neural Networks and Deep Learning

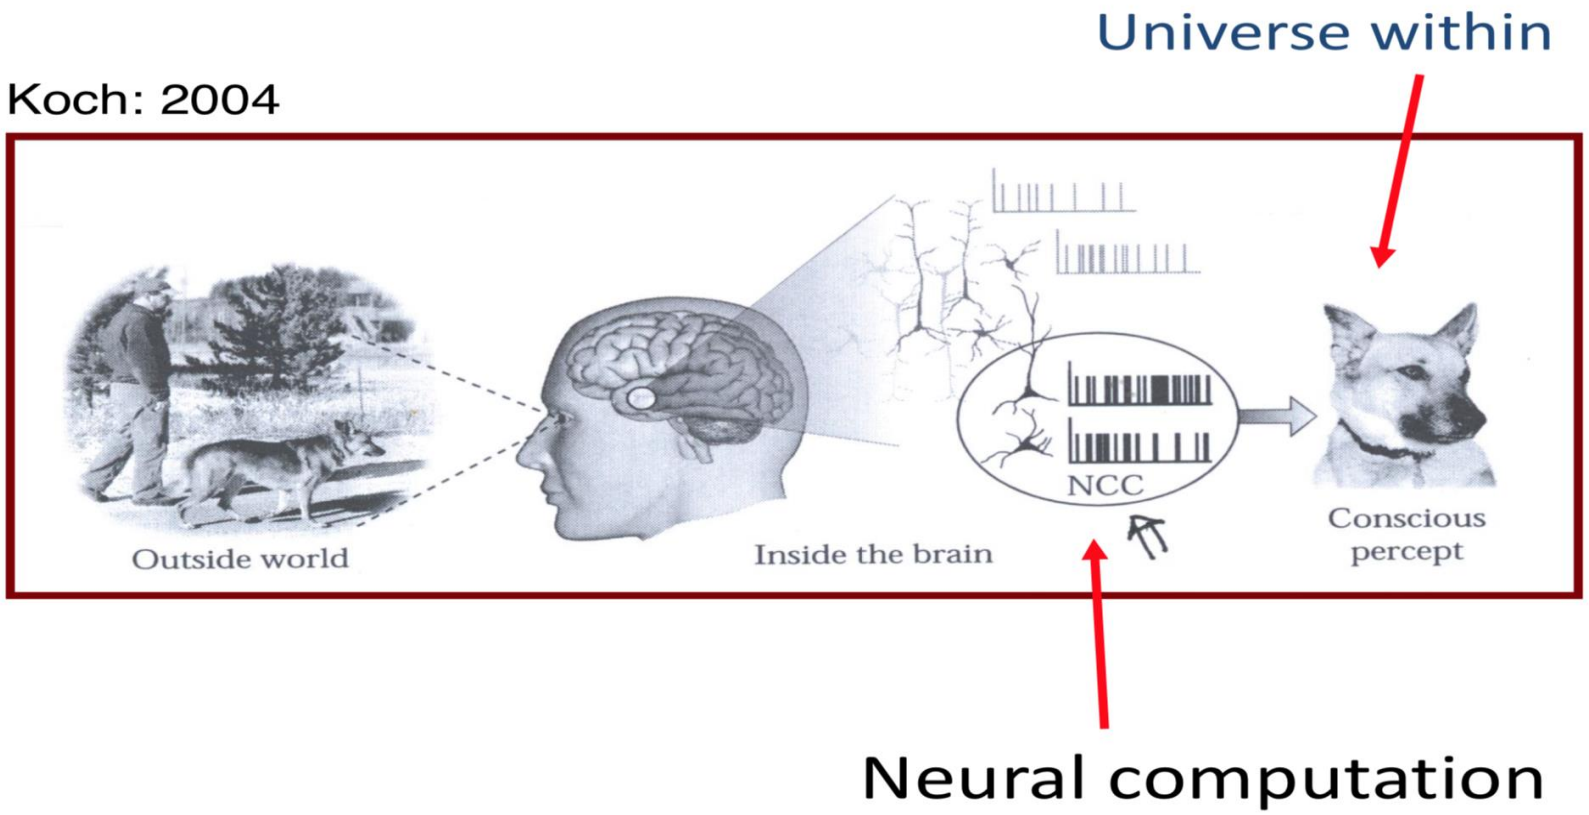

Neural Networks and Deep Learning

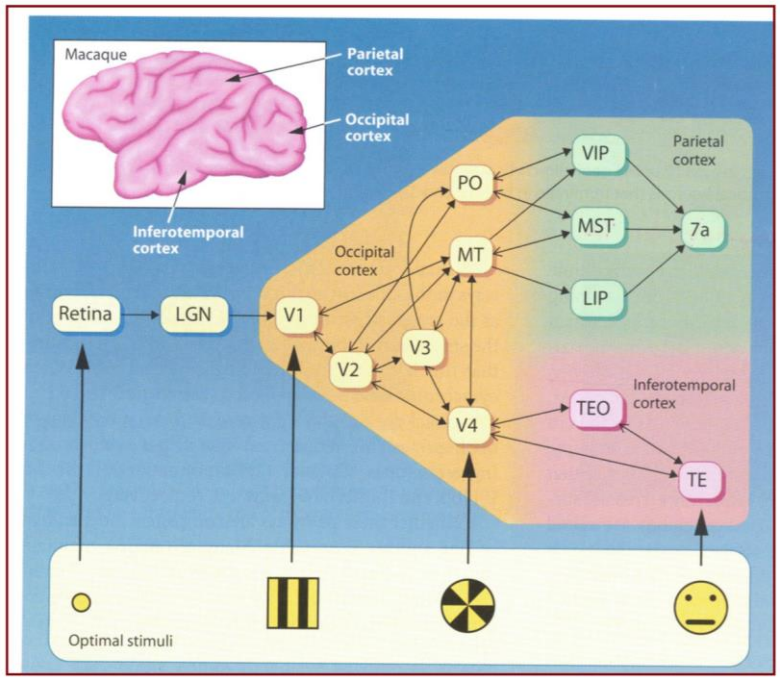

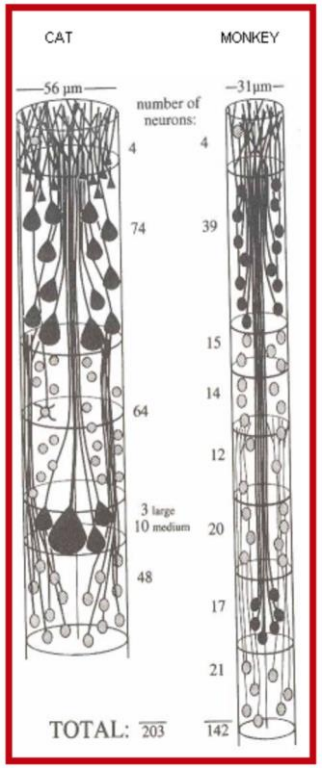

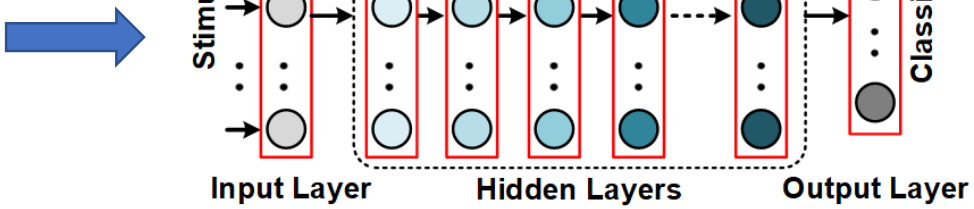

Bastos et al, *Neuron*, 2012

History, Big Data, GPU, and Break Through

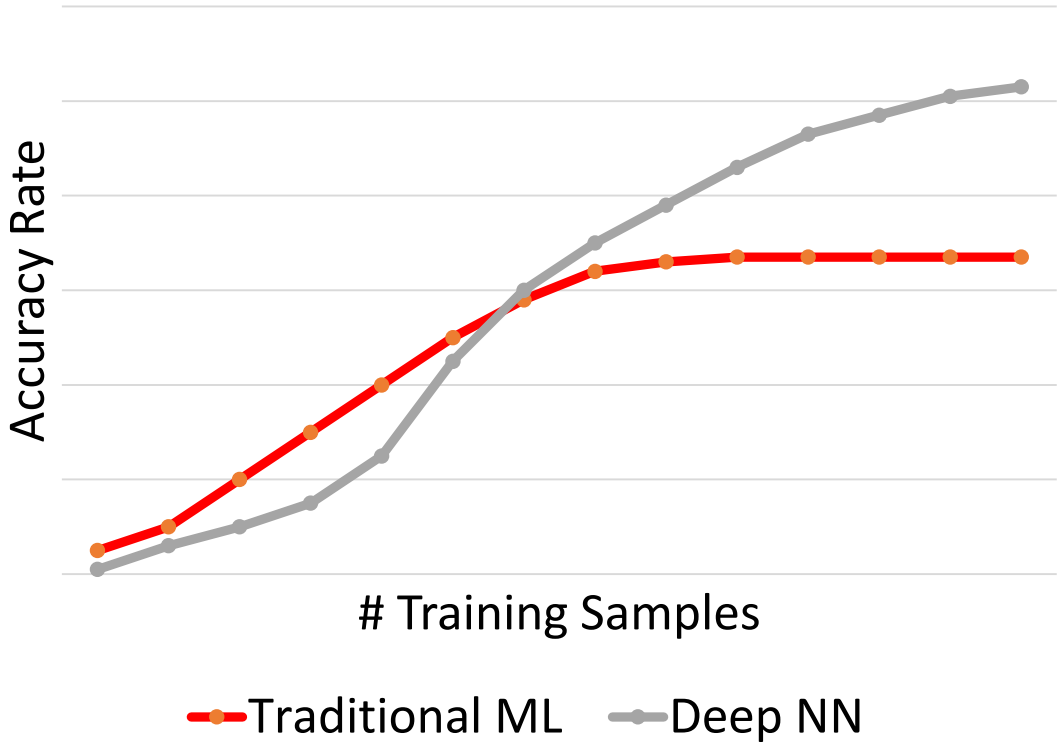

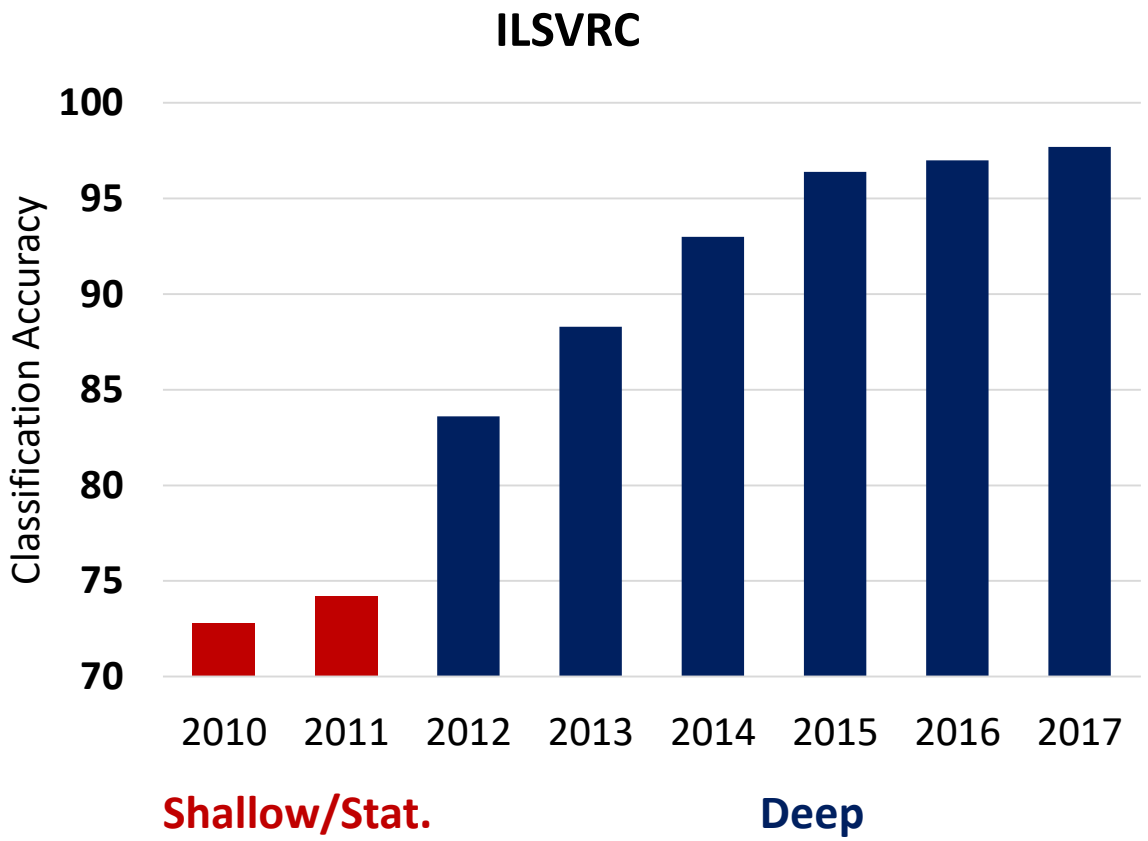

Machine Learning Models

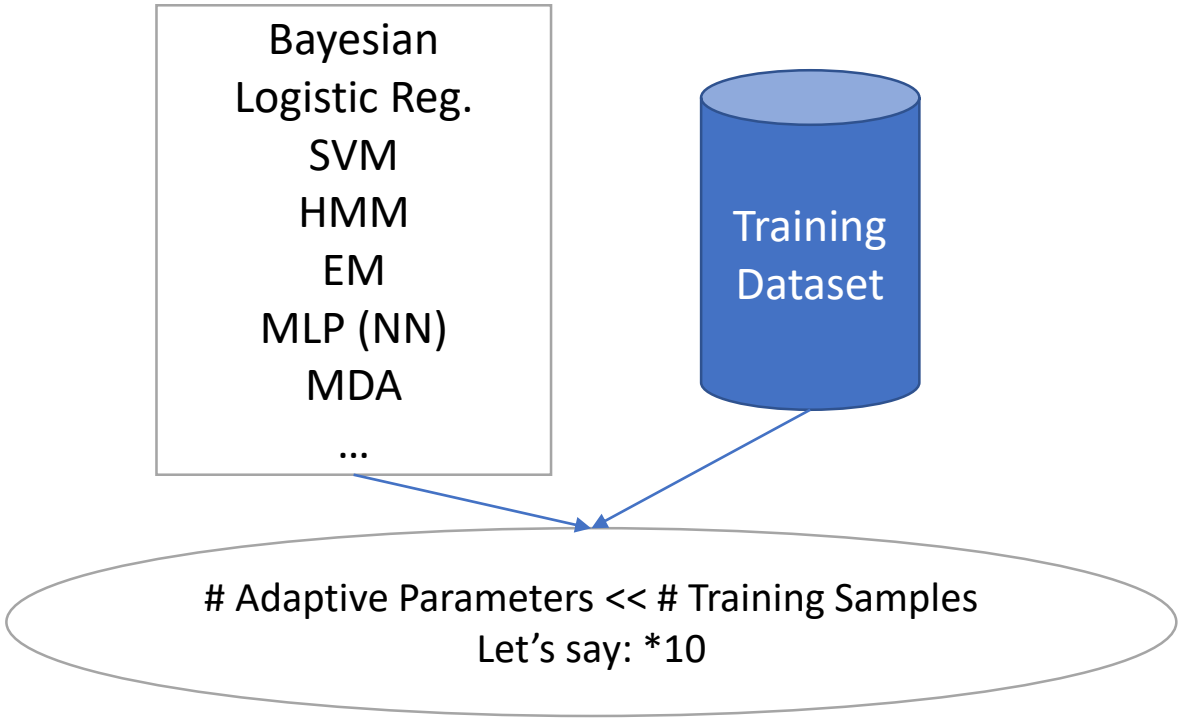

NNs are easily stuck in local minima

Deep NNs have **many** parameters to train

Thanks to Big Data for managing and providing **enough data** for training deep neural networks with many parameters. Enough data points make the NN enable to handle saddle areas and local minima. Thanks to GPU for speeding up the massive computations.

DS/AI Project Management

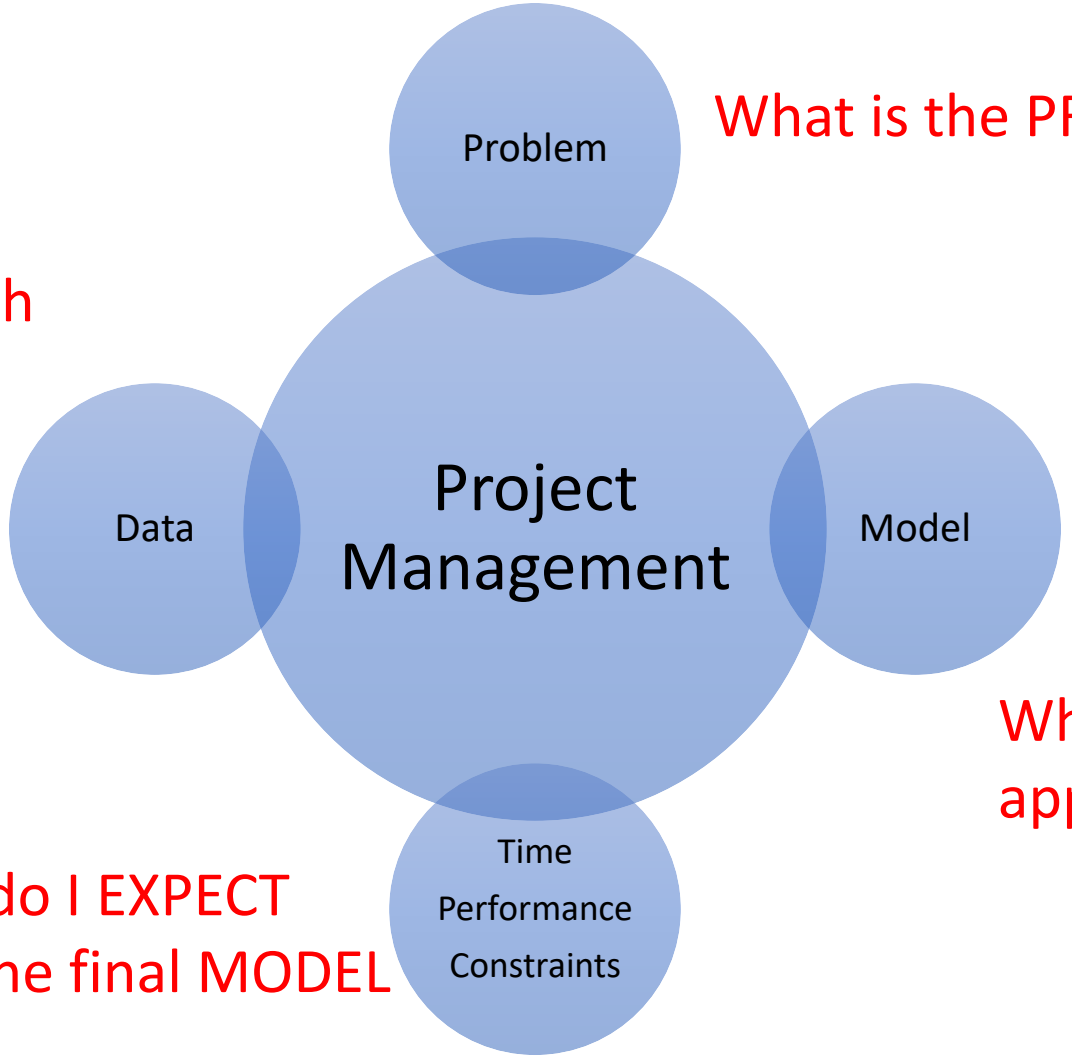

What is the PROBLEM?

Do I have enough Data to train my MODEL

Which MODEL is appropriate?

What do I EXPECT from the final MODEL

Data/Model Variations

Data

- Structured Data
- Image
- Voice
- Video
- Experimental
- Info. Record
- Clinical
- Sensory Signals
- Security and Data Generation
- ...

Type

- N-D
- 2D Pixels
- 1D Frequency (Sequential)
- 2D Pixels (Sequential)
- Multimodal
- Multimodal (Sequential)
- 1D/2D/3D Complex structure
- N-D Numerical (Sequential)
- Max-Min Problem
- ...

Model

- Fully Connected NN
- CNN
- RNN, LSTM
- Convolutional LSTM/RNN
- Multimodal DBN
- Multimodal Pred. Coding
- 3D CNN, RNN, Multimodal
- RNN, Stacked Autoencoder
- GANs, Generative NNs
- ...

Common Mistakes in Applying ML/DL

You won't have a good AI model if you do:

- not aware of data distribution, skewness, balance/imbalance, …
- not have enough and good (clean) data for training
- not choose proper ML model/structure
- not have a strategy to build your model (e.g. # of layers for NN)
- use one algorithm for everything
- not consider generalization and regularization
- not consider required equipment
- not know programing and the theory behind the ML model

Convolutional Neural Network (CNN)

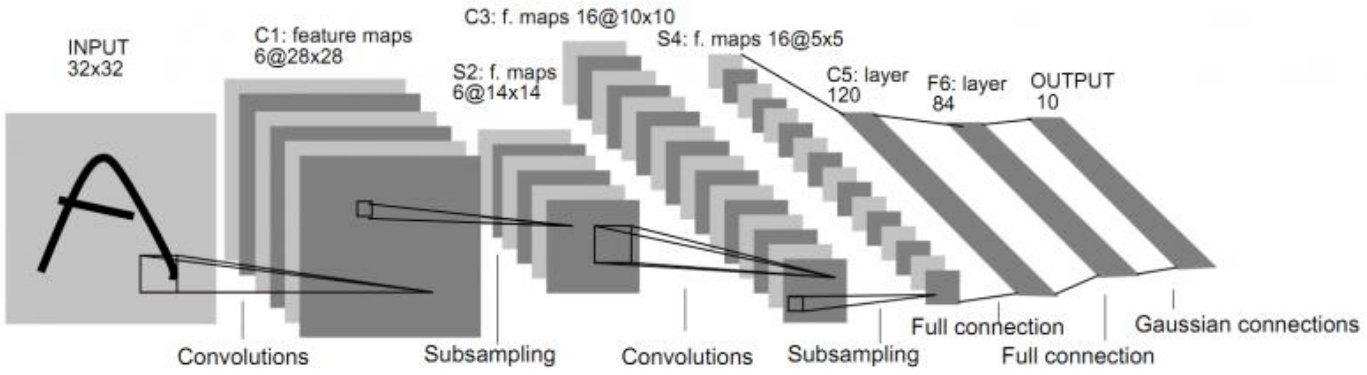

Feature Extraction

Classifier

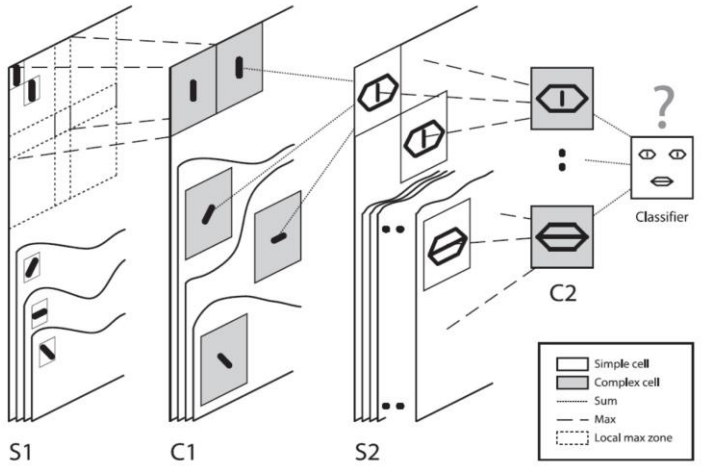

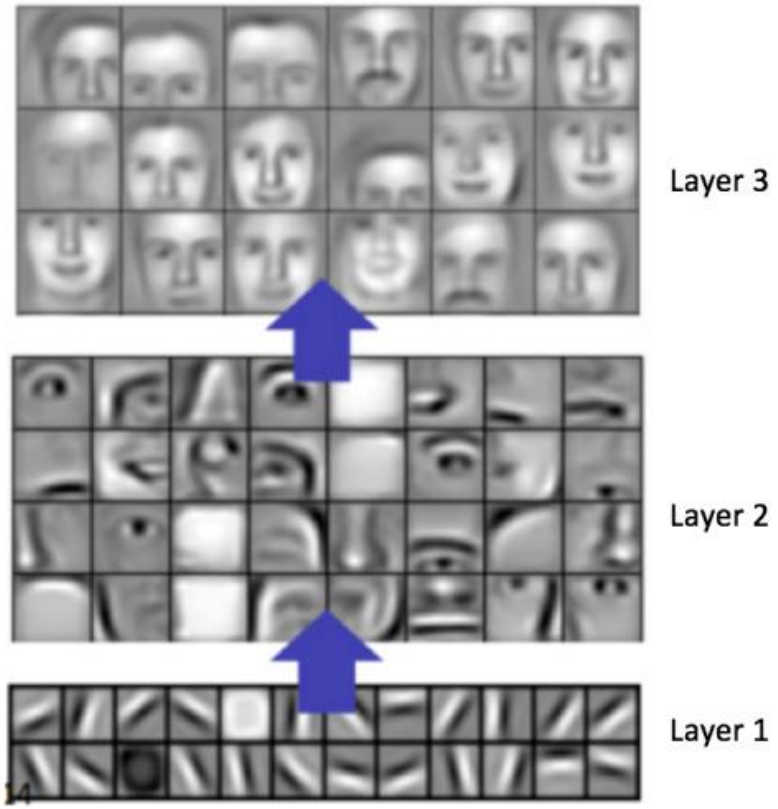

LeNet: Lecun et. al.

Tavanaei et. al. (2016)

Lee et. al. (2009)

CNN Example: Bioinformatics and Biomedicine

Do not ignore the Preprocessing Phase

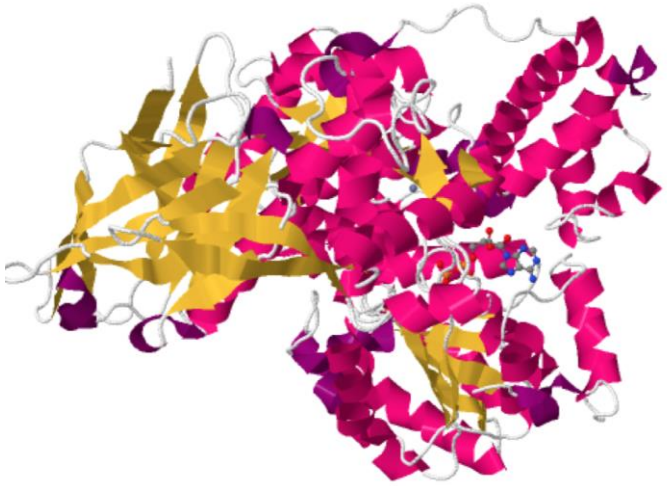

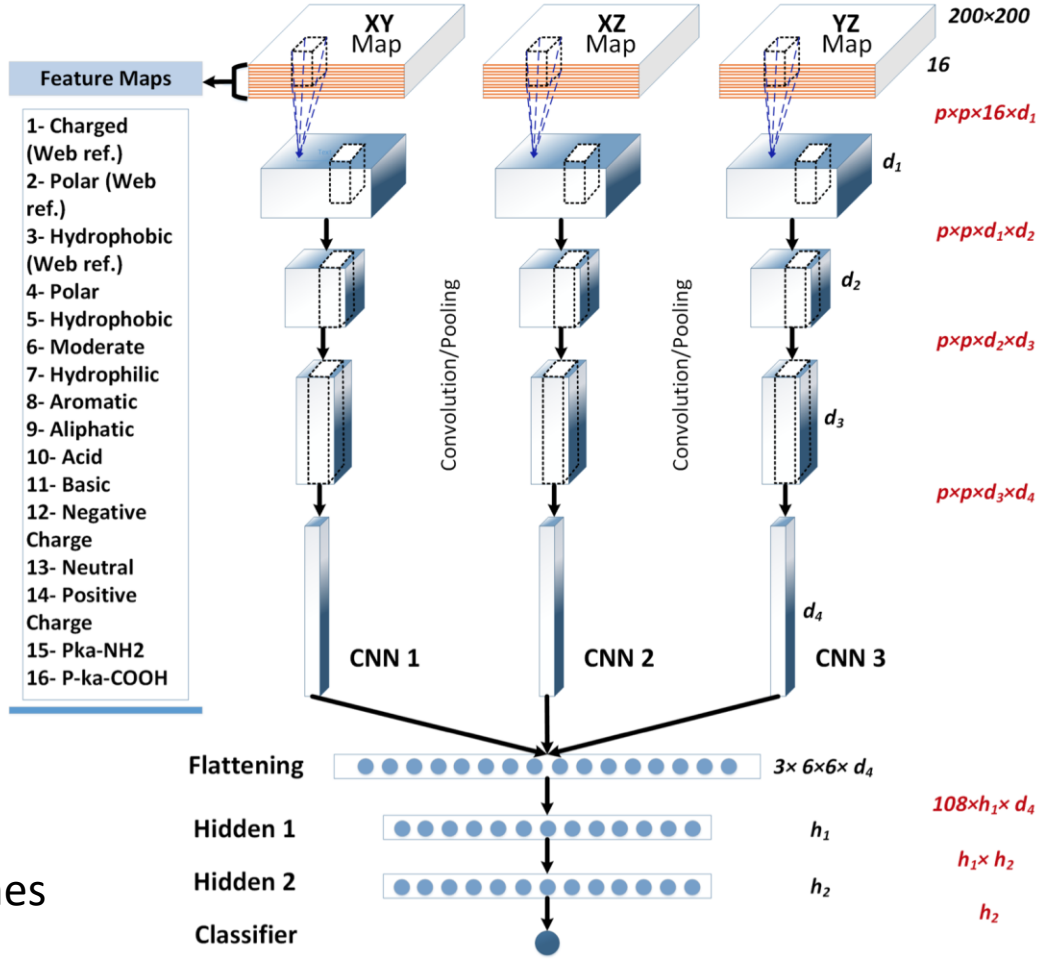

Detecting Tumor Suppressor Genes and Onco-Genes

Tavanaei et. al. (2017)

CNN Example: P&G Skin Advisor

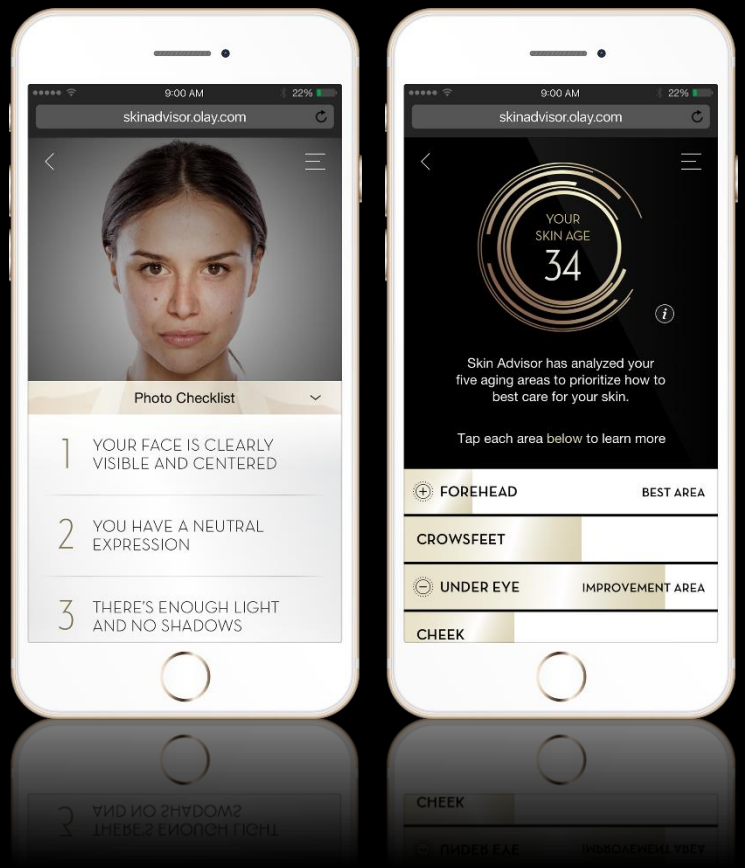

Deep Learning

Image Processing

Cloud Computing

Problem Detection

OLAY®

SKIN ADVISOR

THE SCIENCE BEHIND
OLAY SKIN ADVISOR
skinadvisor.olay.com

CNN Example: P&G Skin Advisor

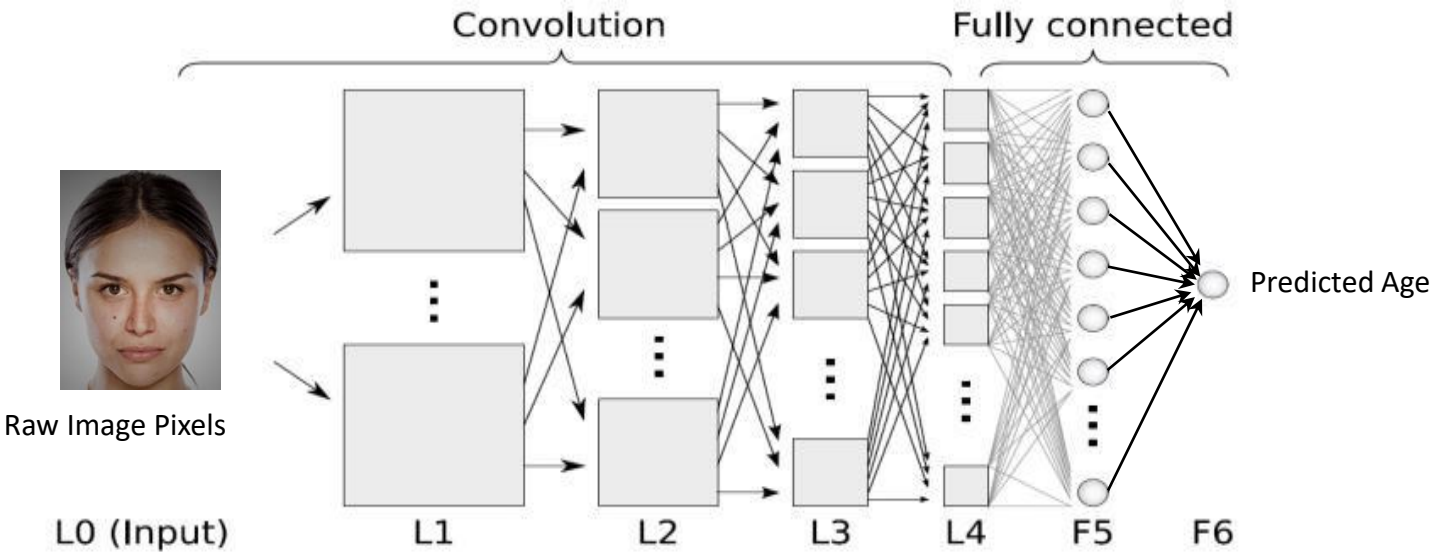

- An input image was forward propagated through the model to obtain a predicted age.

P&G Skin Advisor Demo

- Age is predicted
- A heat map was created.
- The heat map localizes pixel differences of a subject's image relative to younger than their predicted age

Recurrent Neural Network (RNN)

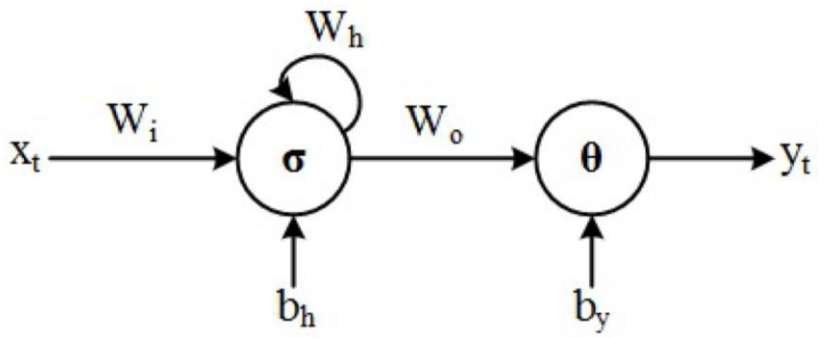

Recurrent Neural Network

Memory gate

Learn to forget

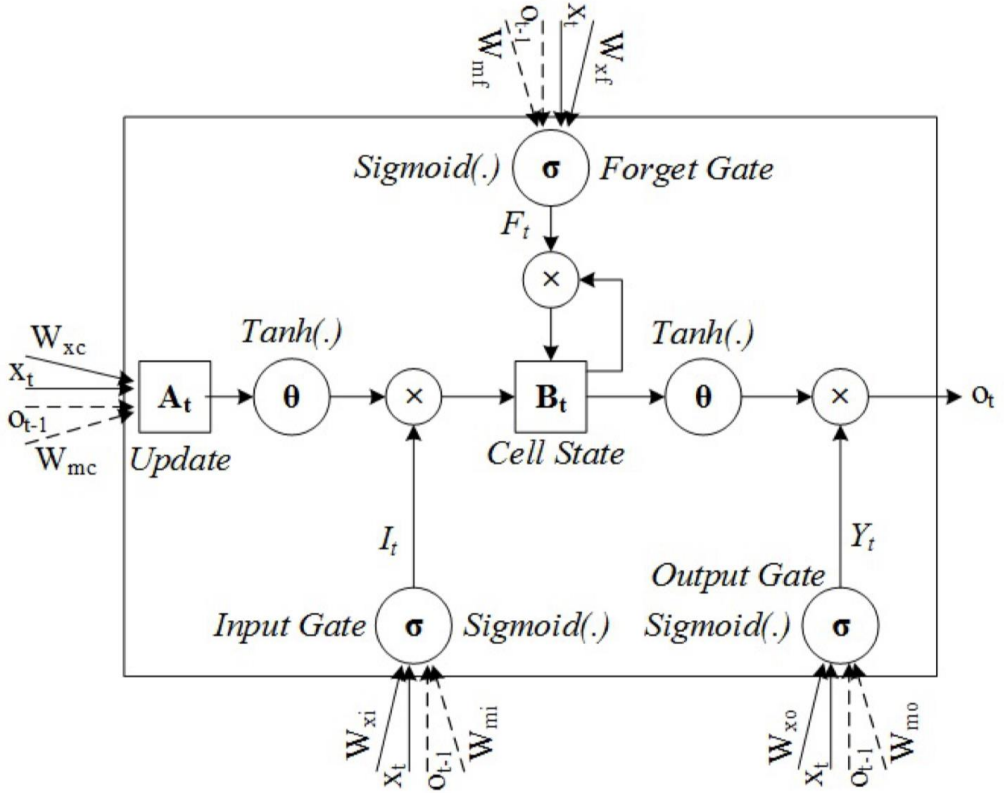

Long Short Term Memory (LSTM)

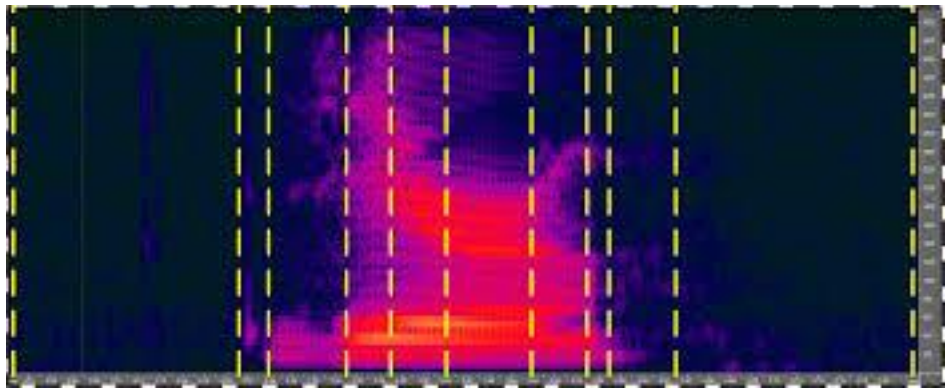

Sequential Data, e.g. Speech

RNN/LSTM Example: Influenza Count Prediction

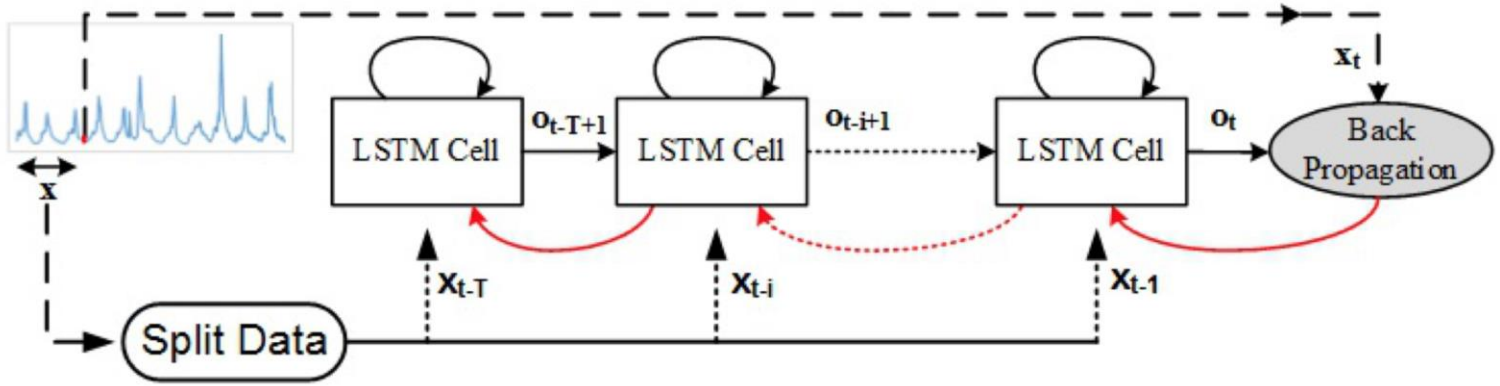

Weeks	1-week			5-weeks			10-weeks			15-weeks		
MODEL	MAPE	RMSPE	RMSE	MAPE	RMSPE	RMSE	MAPE	RMSPE	RMSE	MAPE	RMSPE	RMSE
LSTM	21.38	29.31	0.26	57.09	80.66	0.58	62.32	78.82	1.59	70.05	96.18	1.46
LSTM+CI	21.13	29.17	0.25	57.1	80.96	0.58	62.2	78.58	1.61	69.67	94.76	1.46
LSTM+CI+SA	***16.69***	***23.13***	***0.22***	***51.49***	***72.58***	***0.55***	***60.47***	***76.28***	***1.54***	***65.86***	***87.93***	***1.41***
ARIMA	44.69	83.58	0.3	68.99	95.75	0.68	79.89	100.46	1.78	109.6	154.32	1.93
ARIMA+CI	45.2	83.5	0.3	69.11	95.51	0.68	80.06	100.74	1.76	110.85	156.85	1.94
ARIMA+CI+SA	45.73	86.02	0.28	62.03	85.25	0.67	77.82	97.76	1.71	103.15	143.13	1.88

- Influenza peak prediction
- Emergency management
- Vaccination
- Controlling the influenza trend

Venna, Tavanaei, et. al. (2017)

Generative Adversarial Network (GAN)

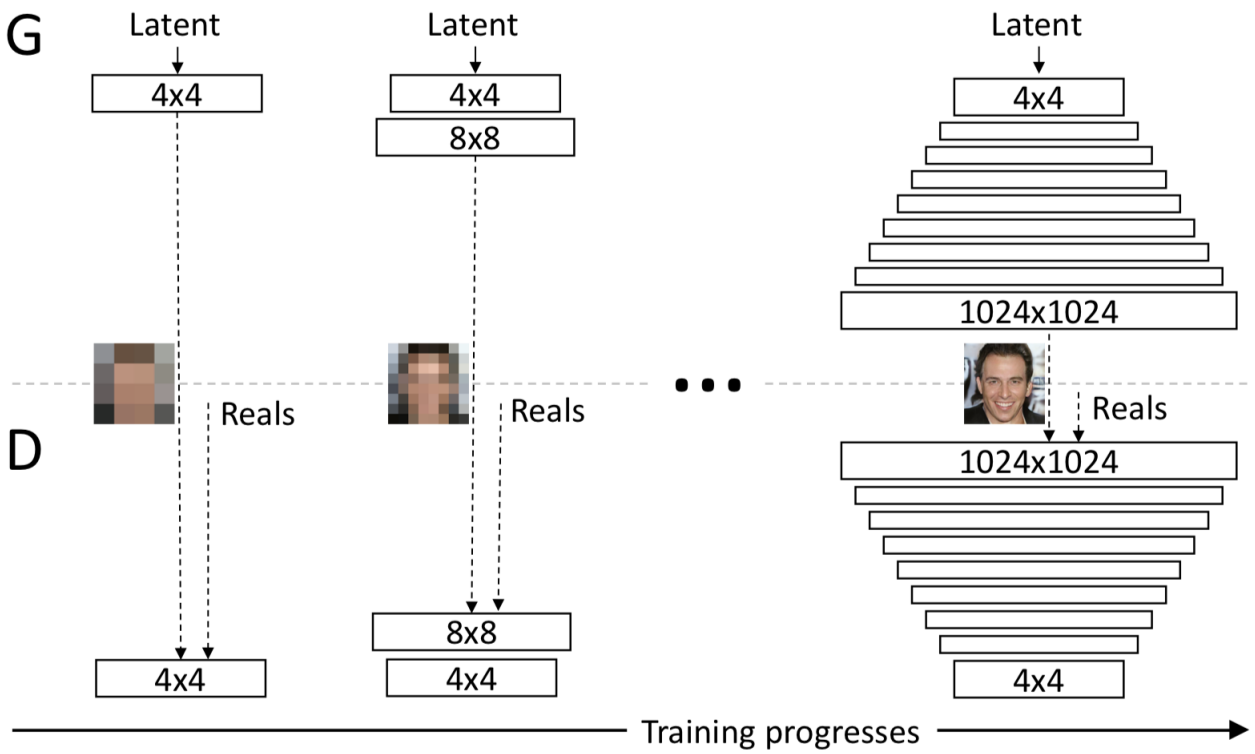

- Virtual image generation
- Security
- Data/Formula generation (e.g. Medicine discovery)
- Automatic data augmentation
- etc.

G: Generator

D: Discriminator

Karras et. al. (2018), NVIDIA Research

Summary

Do I have enough Data to train my deep learning model?
Which type of data do we have?
Can I acquire or augment the data?

What do we want to do?

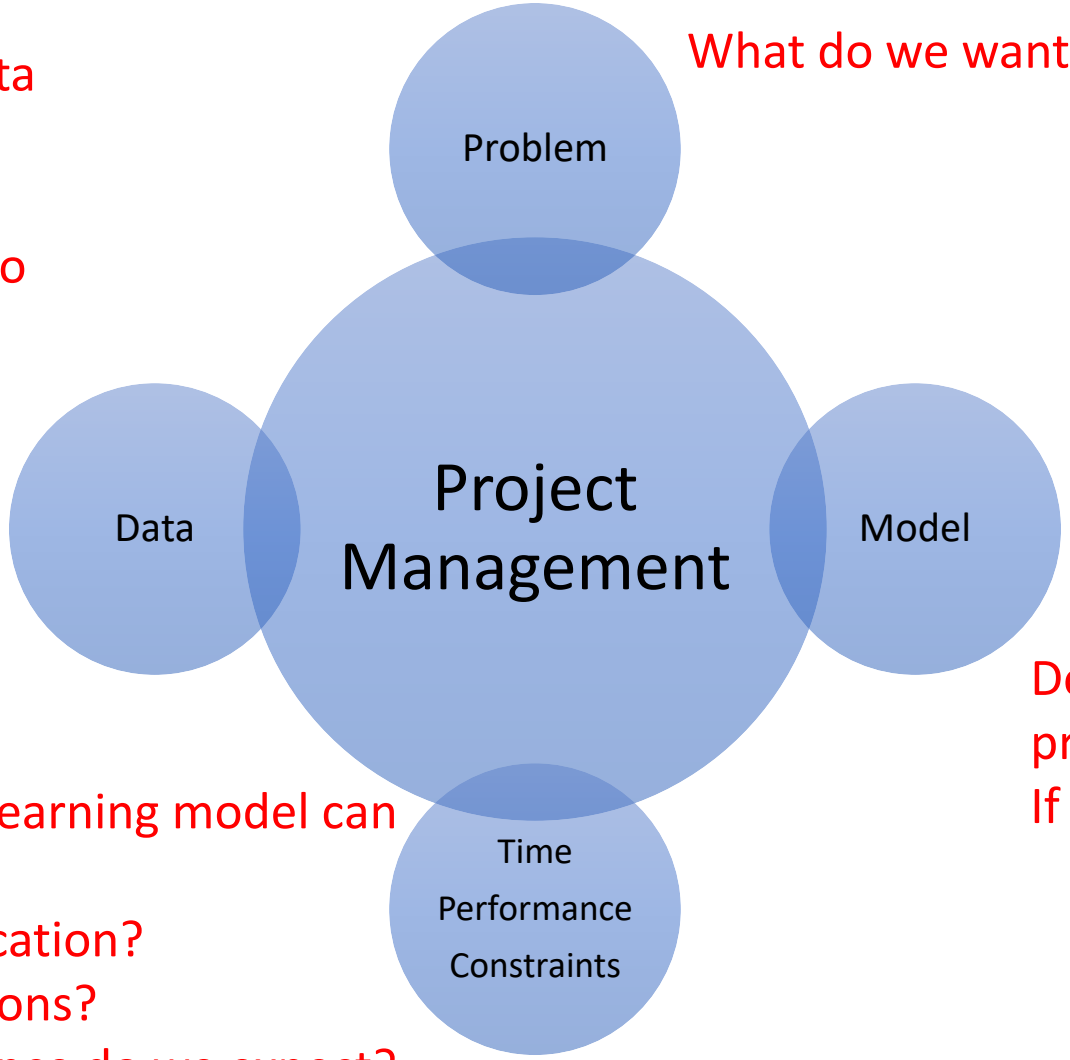

Does deep learning is a proper solution?
If yes, which architecture?

How large the deep learning model can be?
Is it a real time application?
What are the limitations?
How much performance do we expect?

Question?

Thank you!

Contact Info:

tavanaei.a@pg.com

Deep Learning is not a hammer
It is a HUGE toolbox