Hawaii Machine Learning Meetup

History and Applications of
Neural Networks

Table of Contents

Table of Contents

What is a neural network?

Definition: **A computational model inspired by biological neurons.**

Biological Neuron

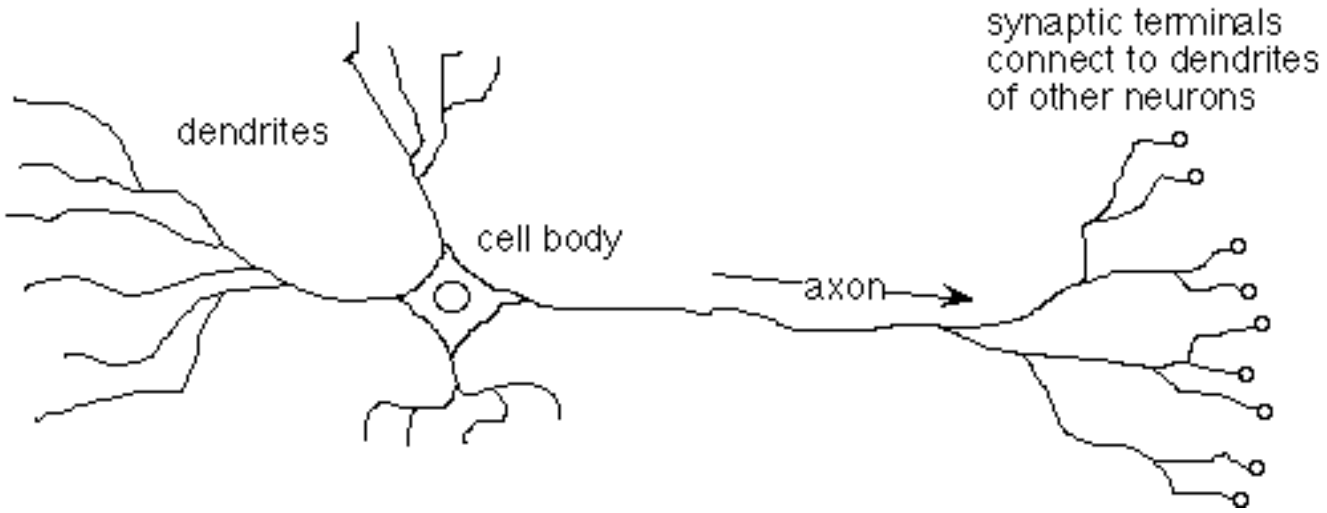

Artificial Neuron

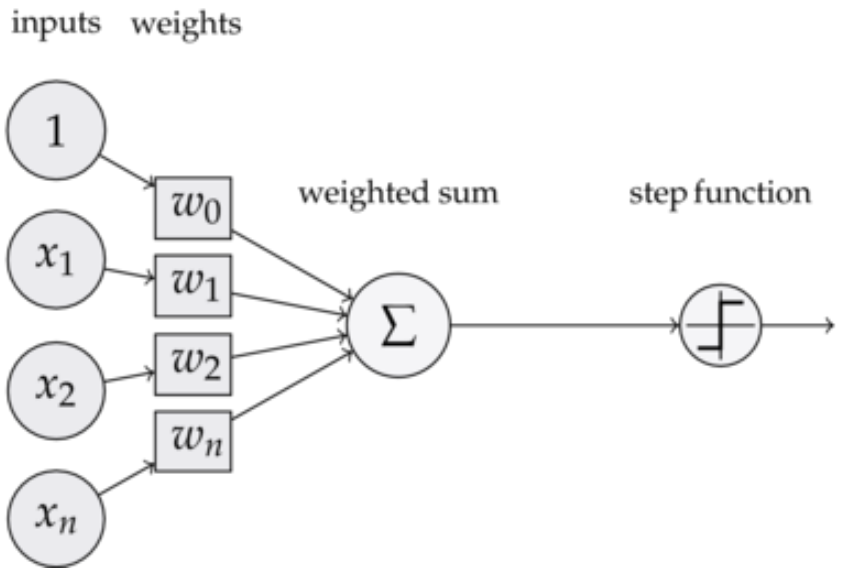

What is a neural network?

Pros

- State-of-the-art on many tasks
- Outperform humans
- Offload feature engineering
- Modular software (e.g., TensorFlow, pyTorch, Keras)

Cons

- Black box
- Requires lots of data and processing power
- Artificial neurons are unrealistically simple

Table of Contents

A Brief History

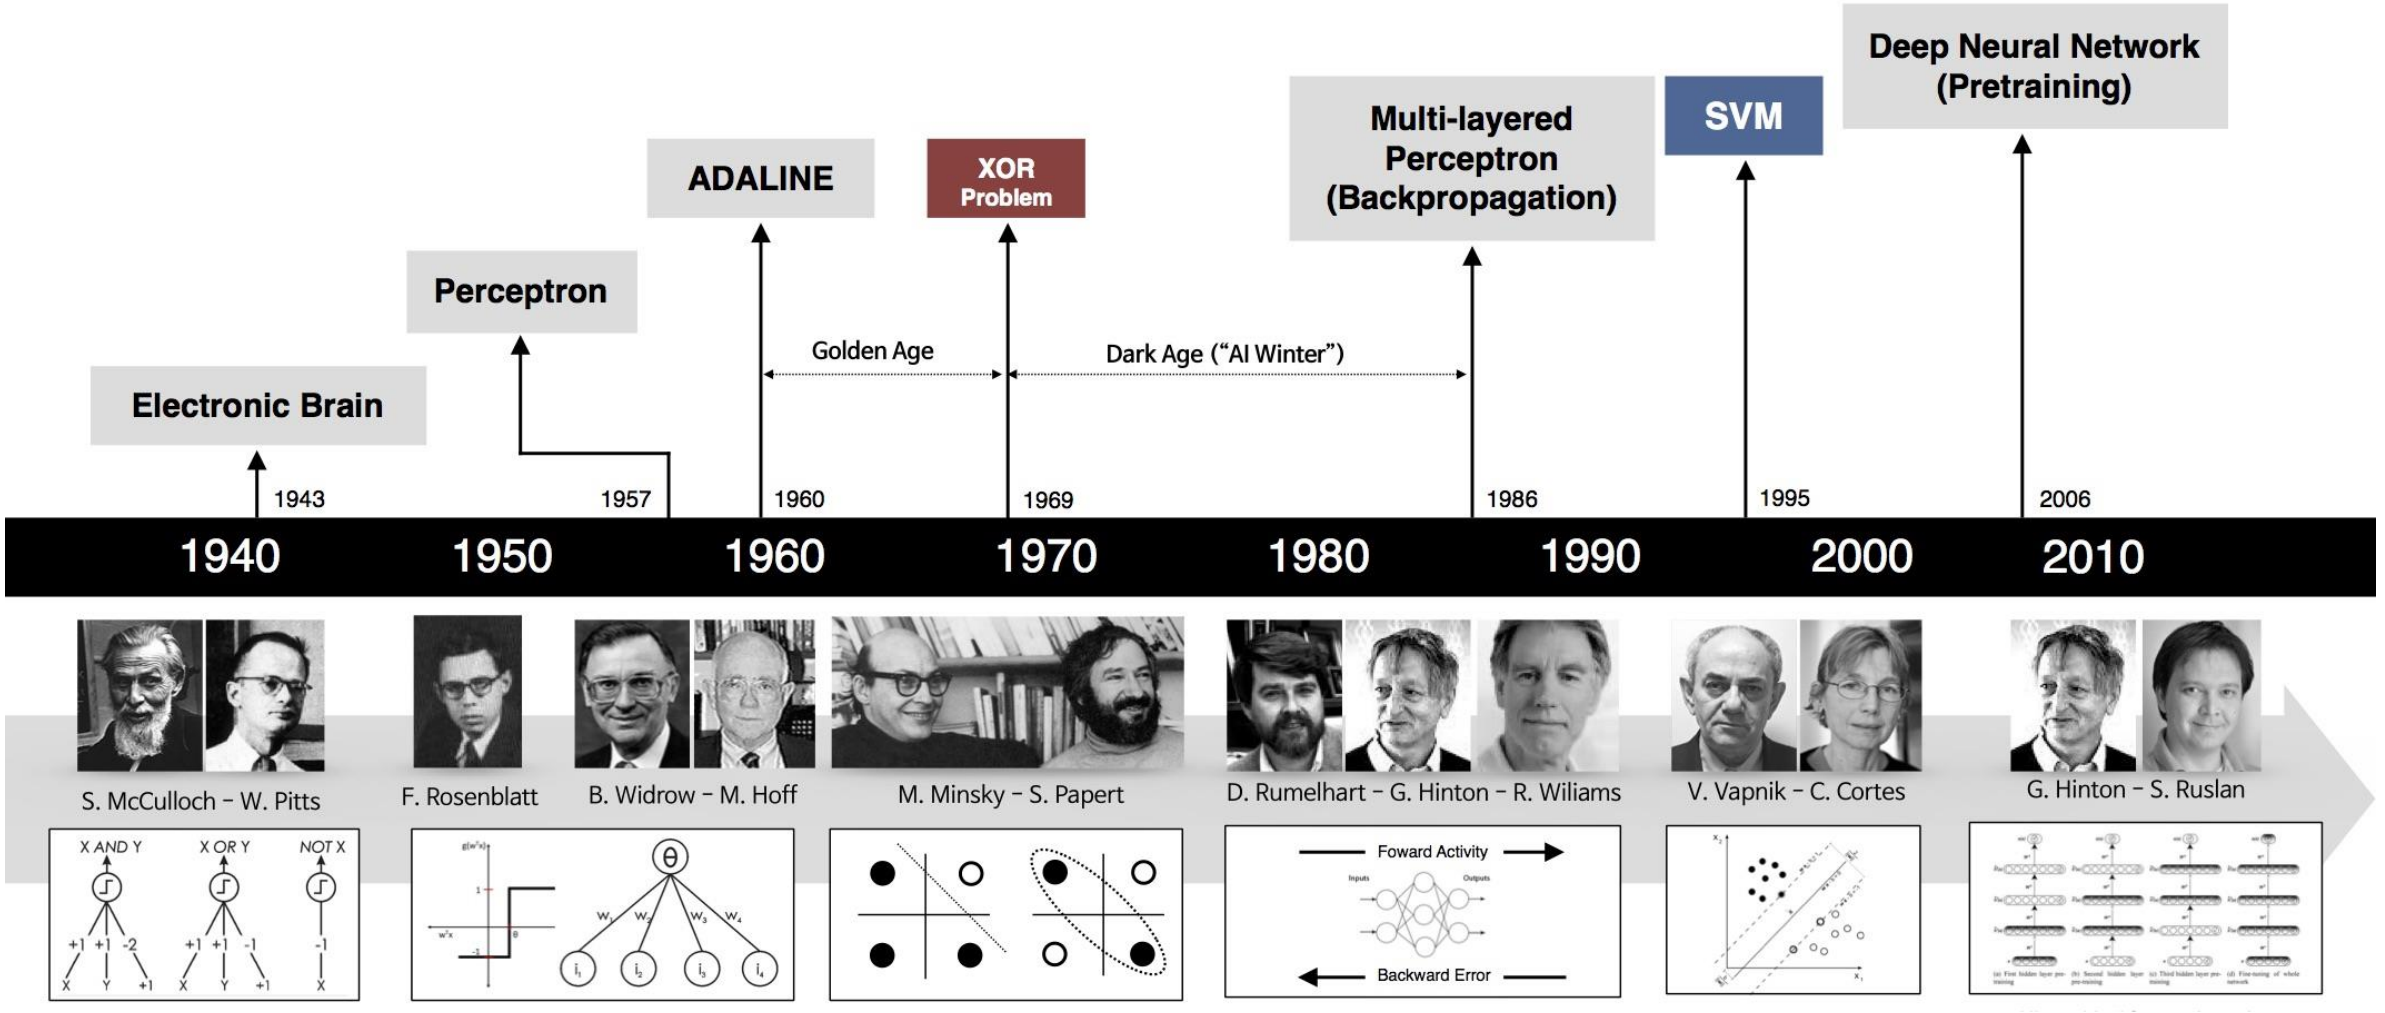

A Brief History

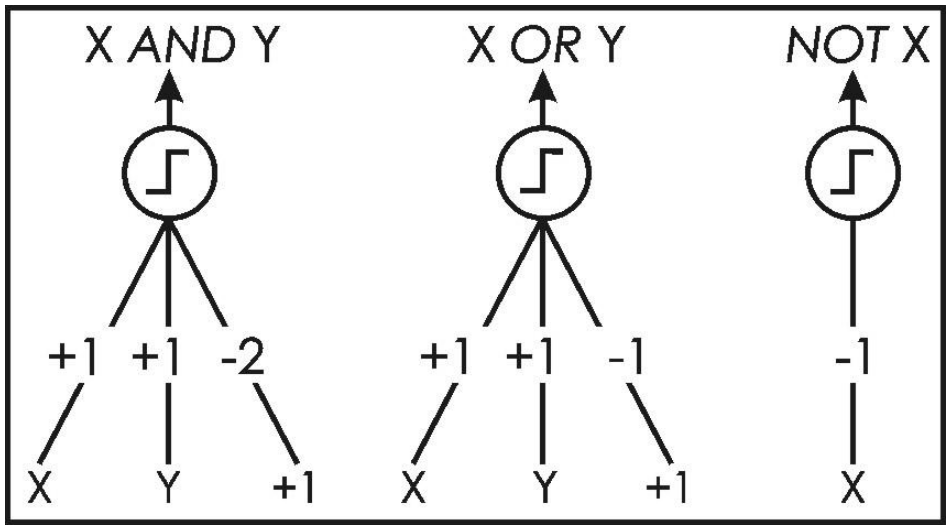

Electronic Brain

1943

1940 1950 1960 1970 1980 1990 2000 2010

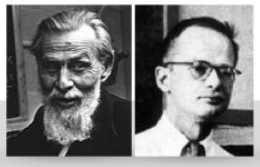

S. McCulloch – W. Pitts

F. Rosenblatt

B. Widrow – M. Hoff

M. Minsky – S. Papert

D. Rumelhart – G. Hinton – R. Wiliams

V. Vapnik – C. Cortes

G. Hinton – S. Ruslan

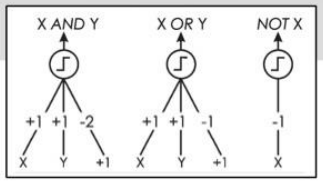

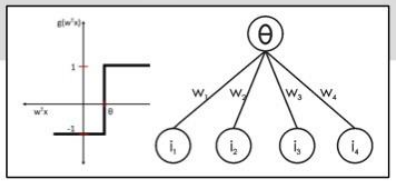

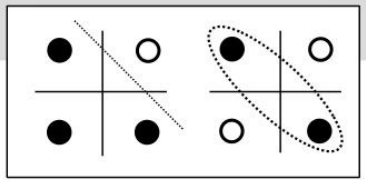

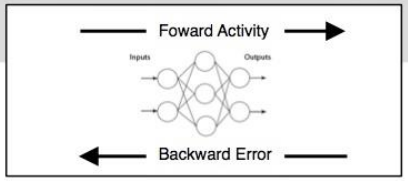

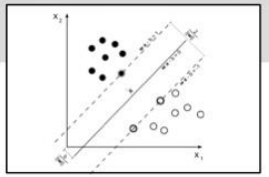

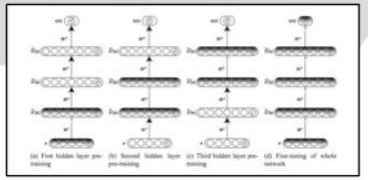

A Brief History

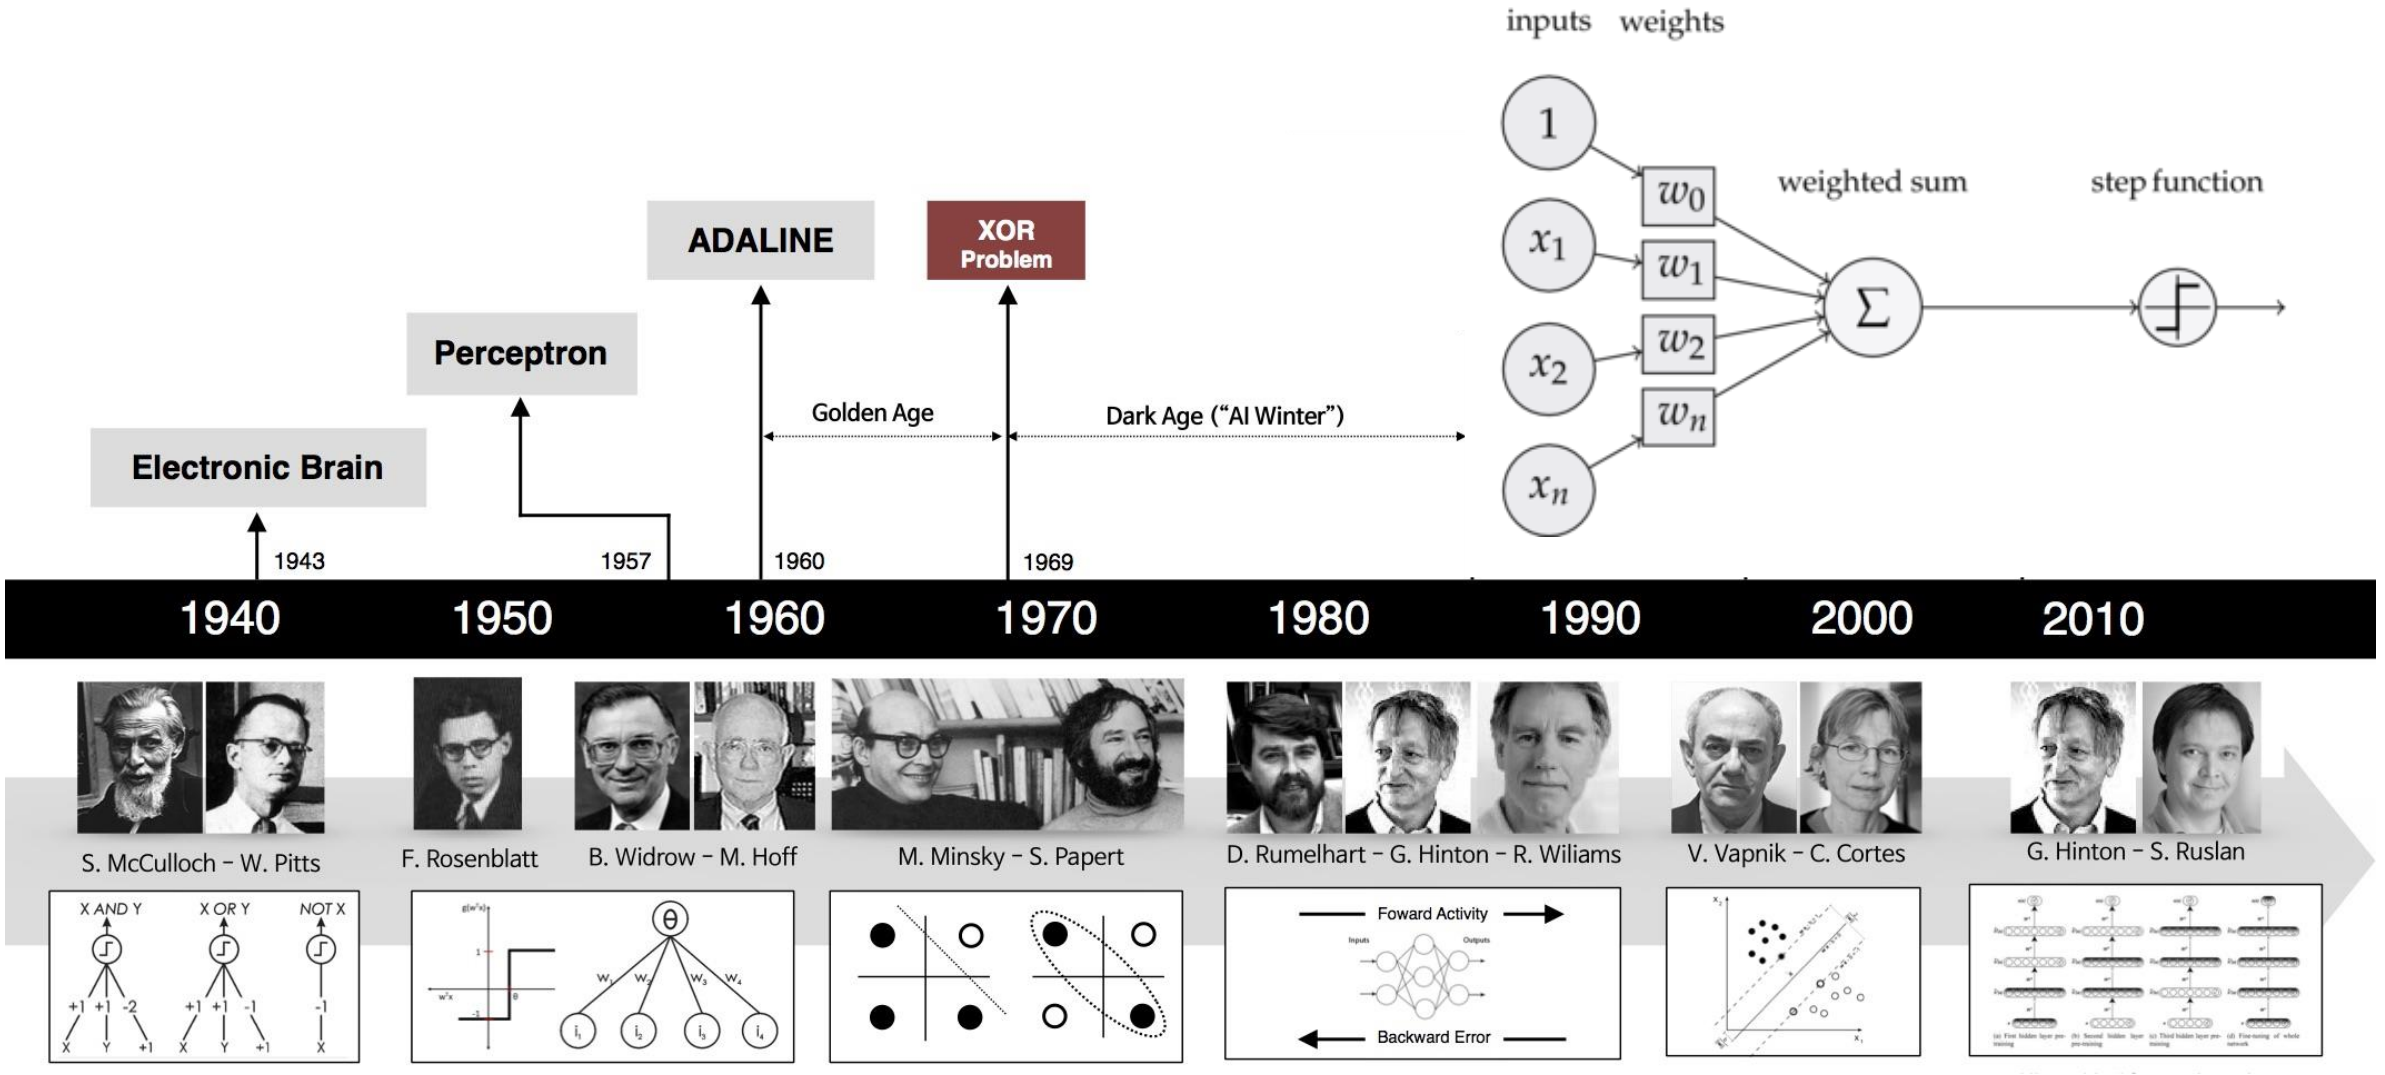

A Brief History

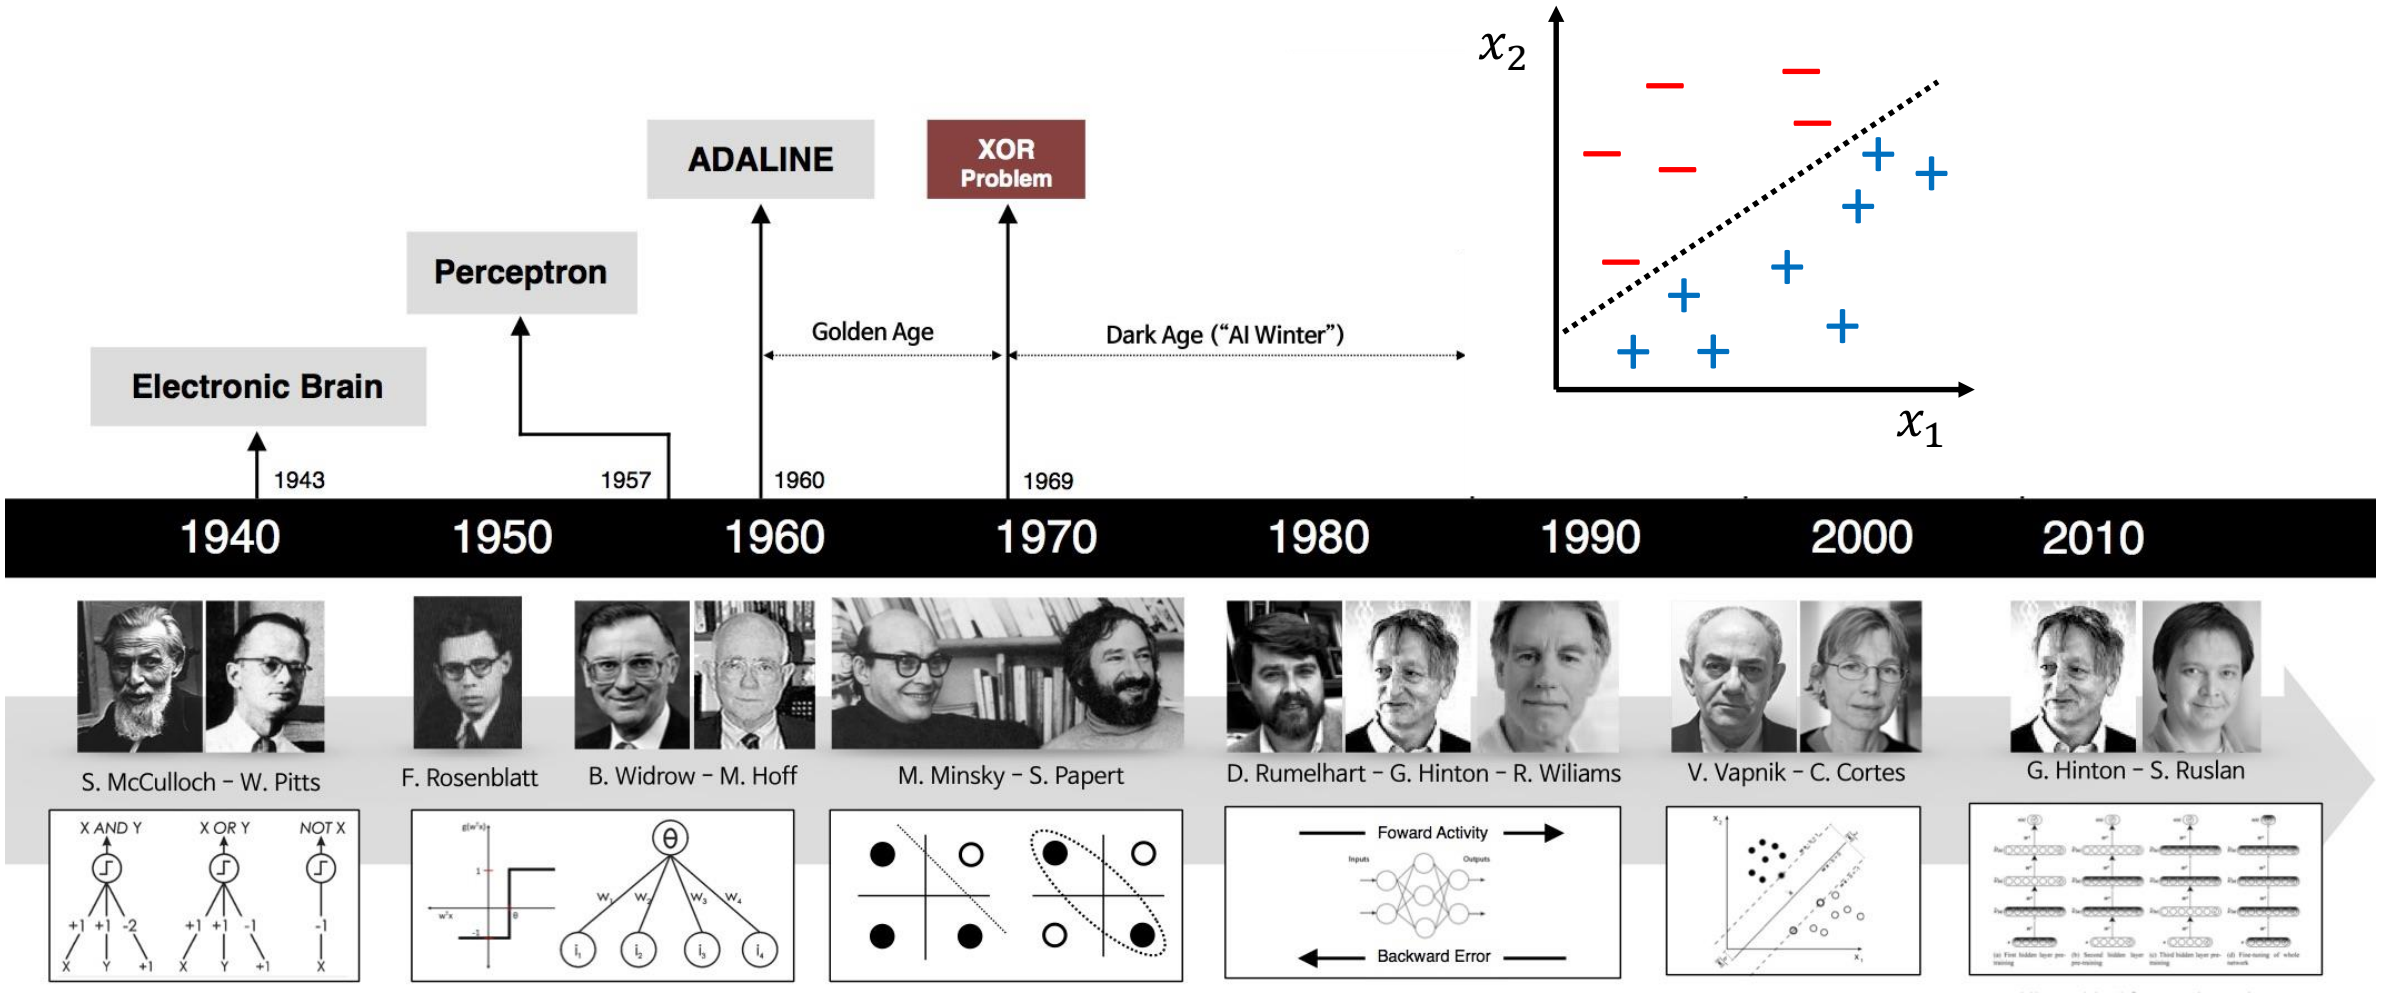

A Brief History

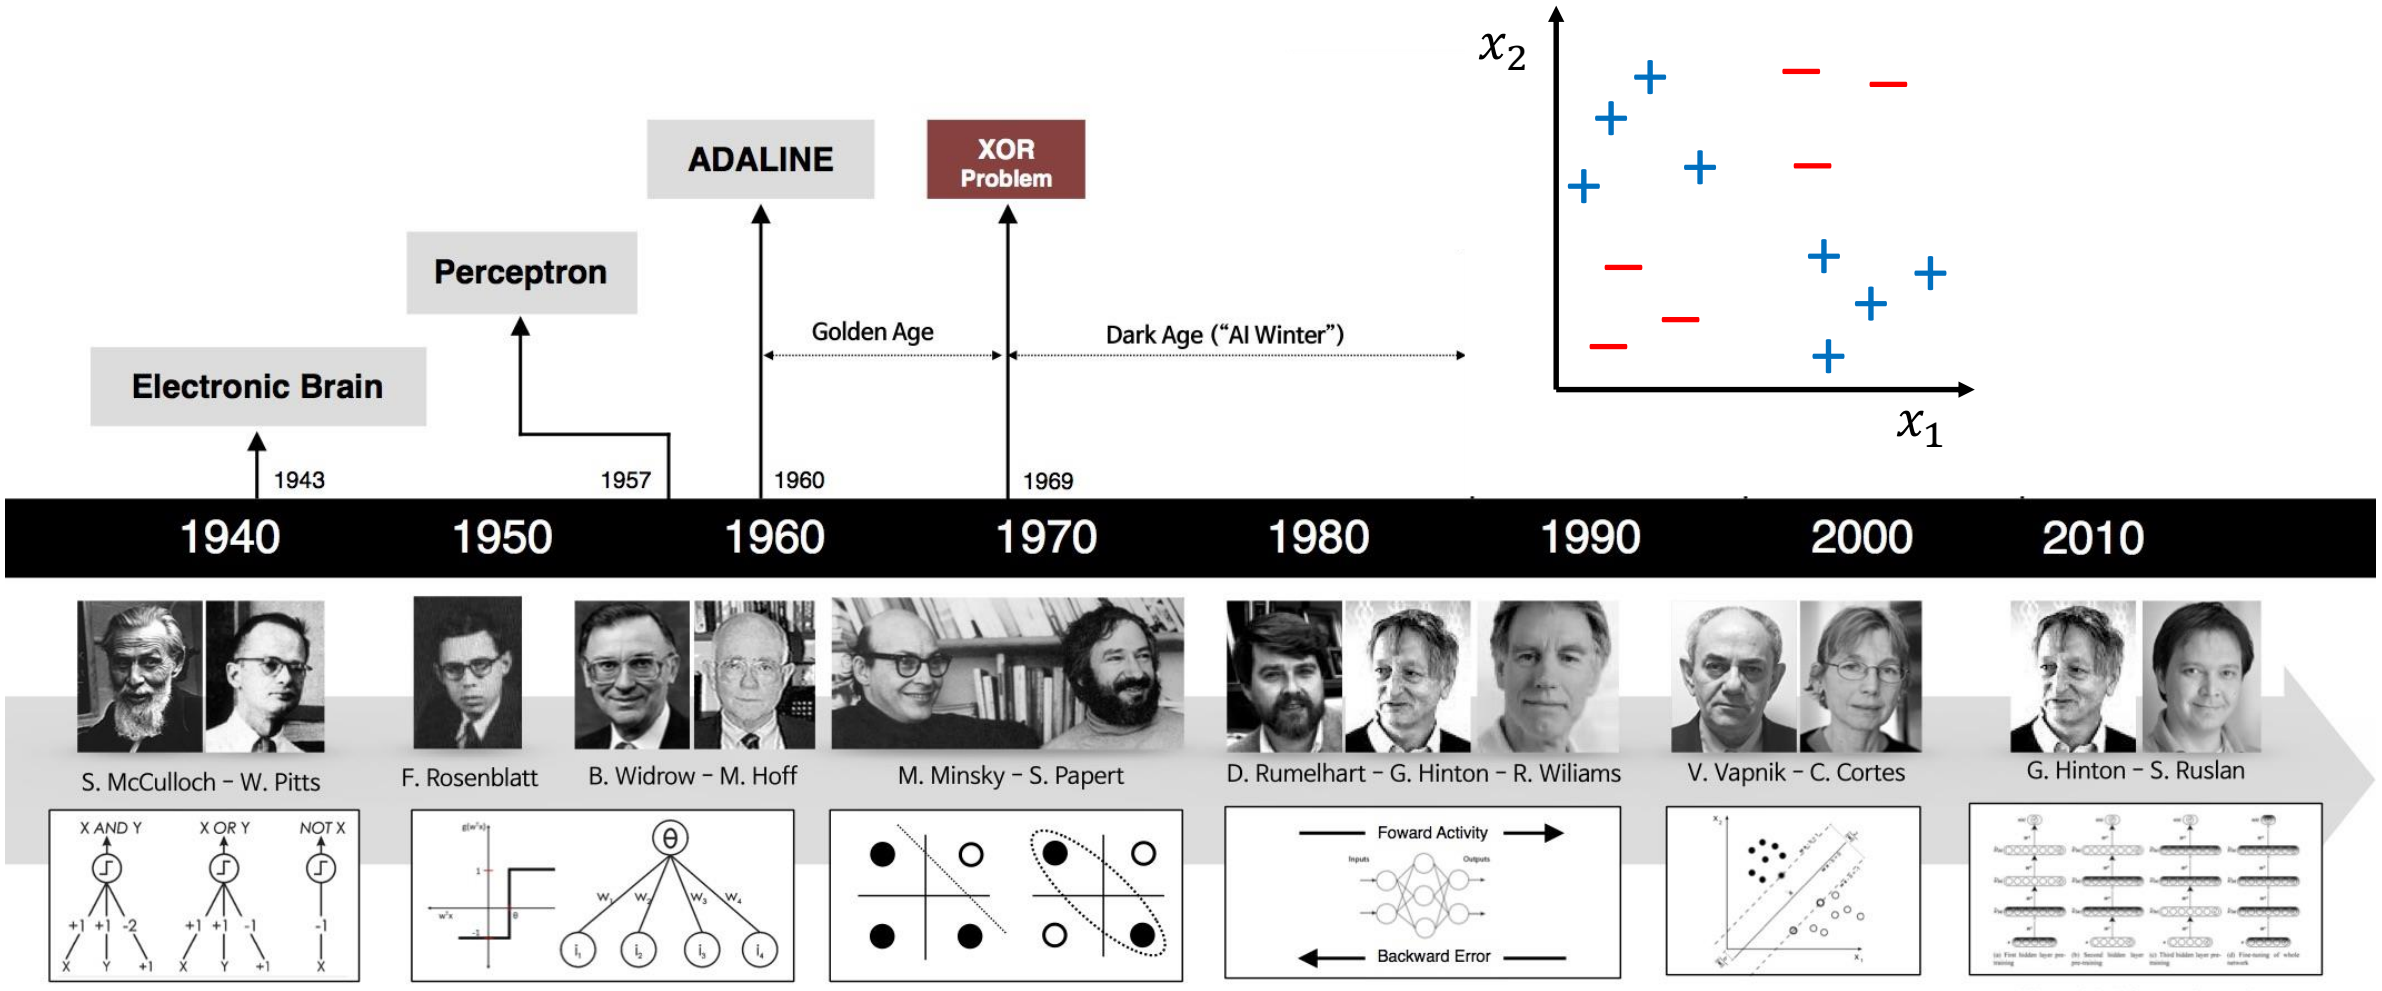

A Brief History

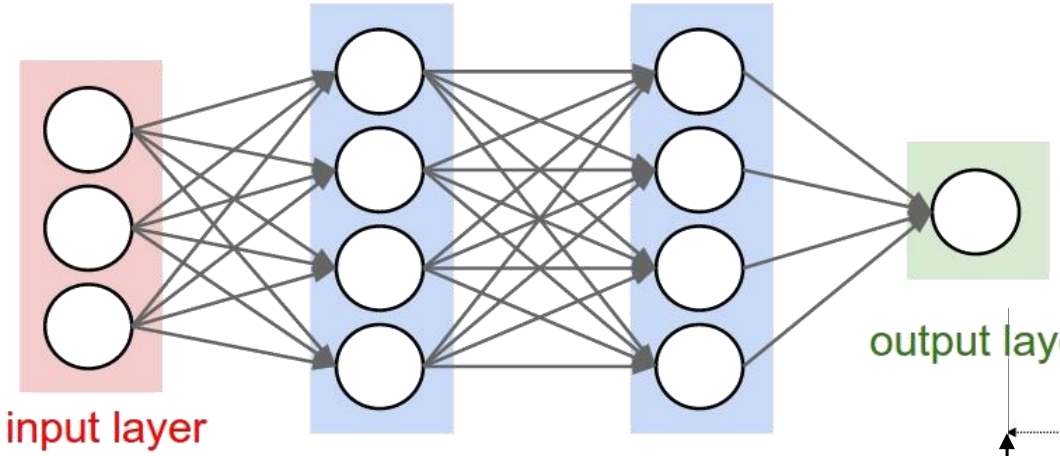

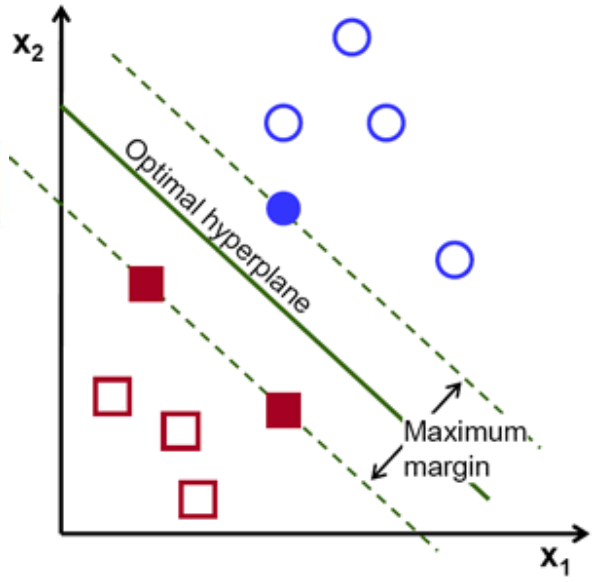

Dark Age ("AI Winter")

Multi-layered Perceptron (Backpropagation)

SVM

1969 1986 1995

1940 1950 1960 1970 1980 1990 2000 2010


S. McCulloch – W. Pitts


F. Rosenblatt


B. Widrow – M. Hoff


M. Minsky – S. Papert


D. Rumelhart – G. Hinton – R. Wiliams


V. Vapnik – C. Cortes


G. Hinton – S. Ruslan

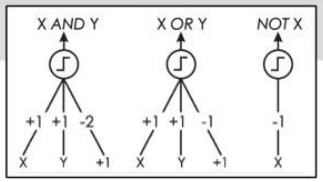

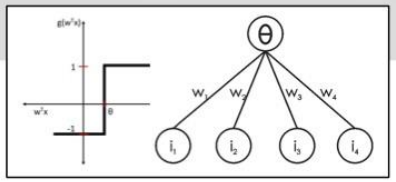

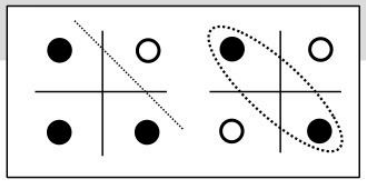

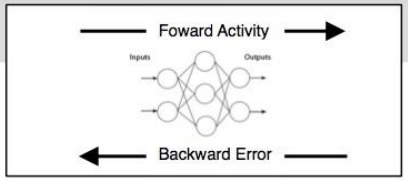

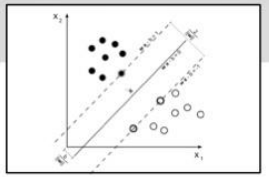

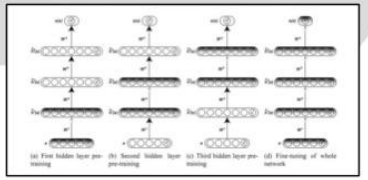

A Brief History

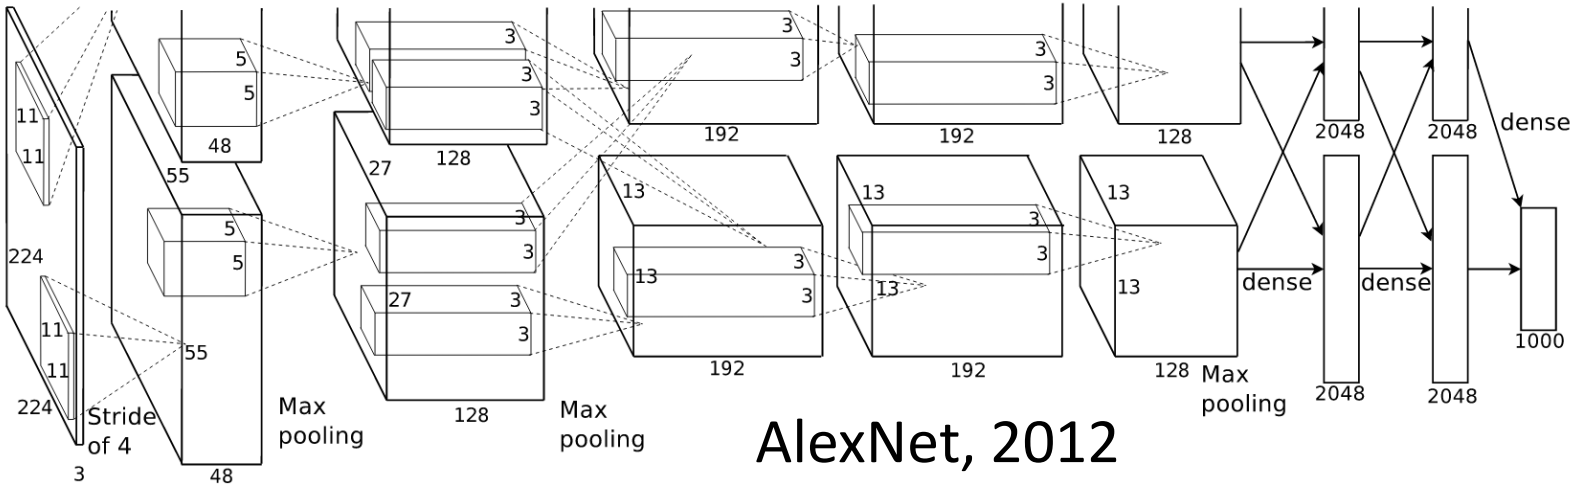

AlexNet, 2012

Table of Contents

Applications – Image Classification

Convolutional Neural Network

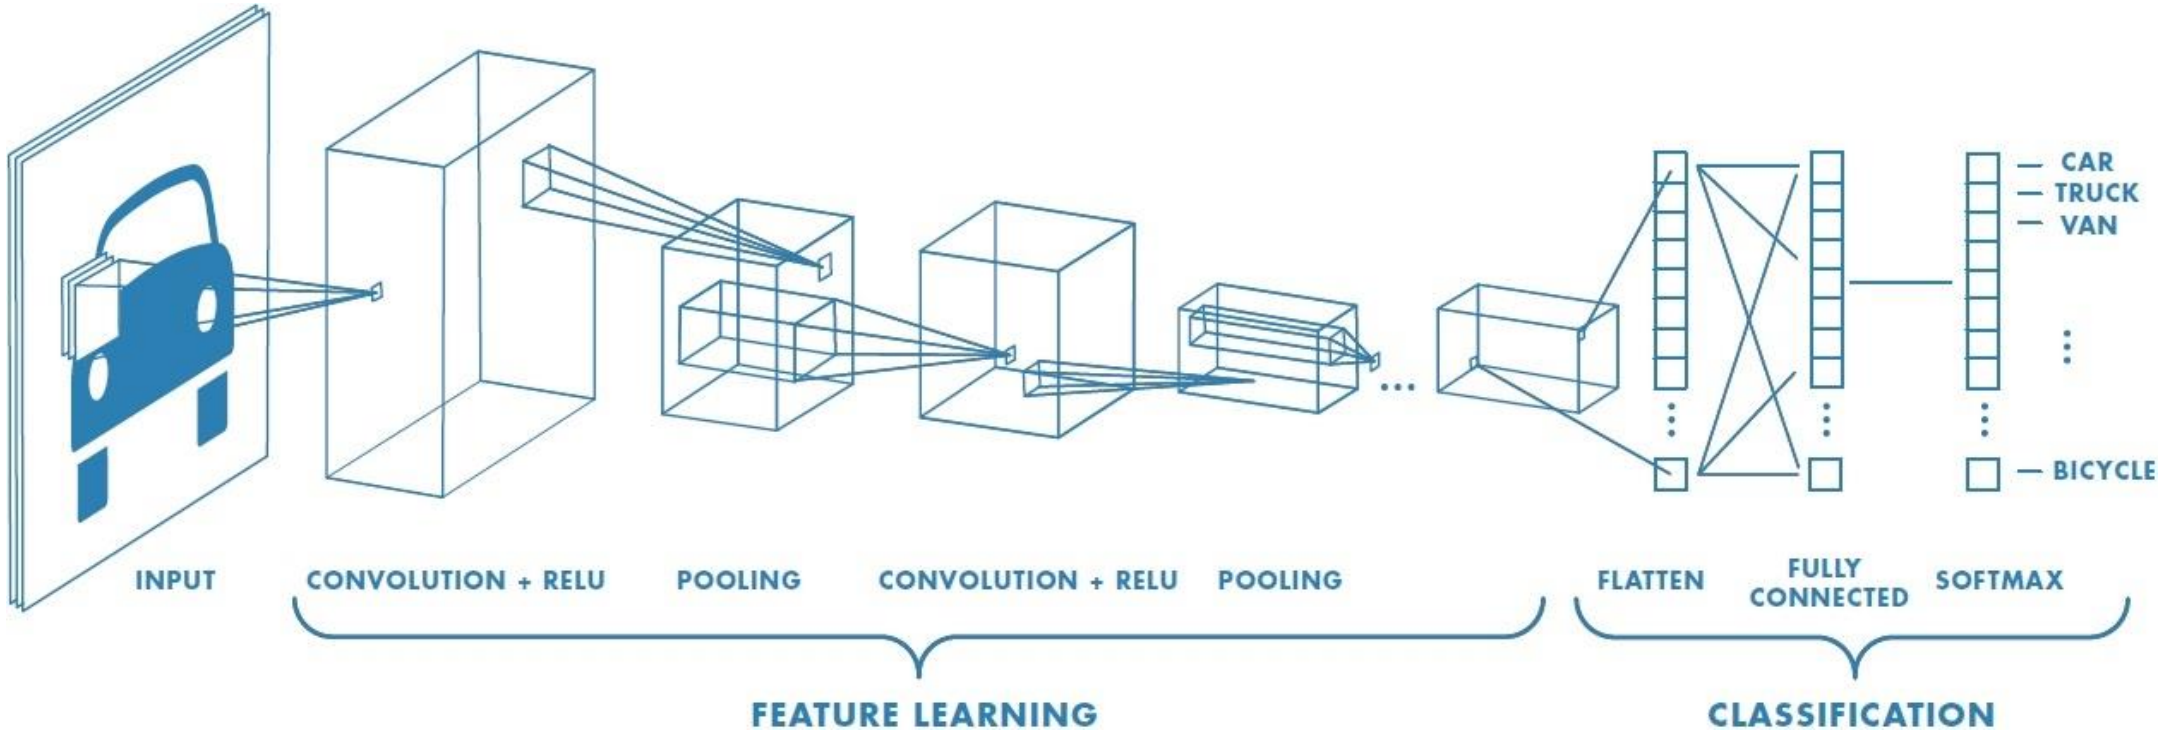

Applications – Image Classification

Convolutional Neural Network

- Convolution

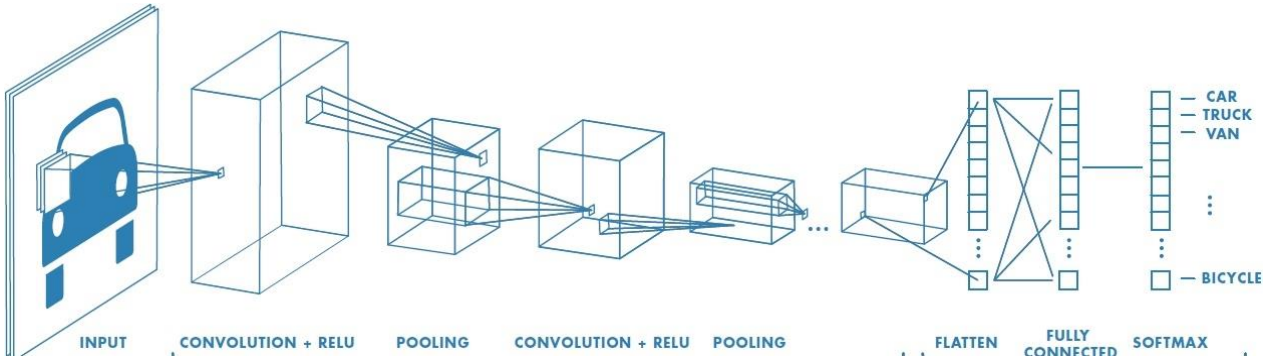

0	0	0	0	0	0	0
0	0	21	0	0	0	0
0	85	71	0	0	0	0
0	250	231	127	63	3	0
0	250	252	250	209	56	0
0	250	252	250	250	83	0
0	0	0	0	0	0	0

Image

○

0	0	1
0	1	0
1	0	0

Kernel

→

0	106			

Feature map

Applications – Image Classification

Convolutional Neural Network

- Convolution
- Nonlinear Activation

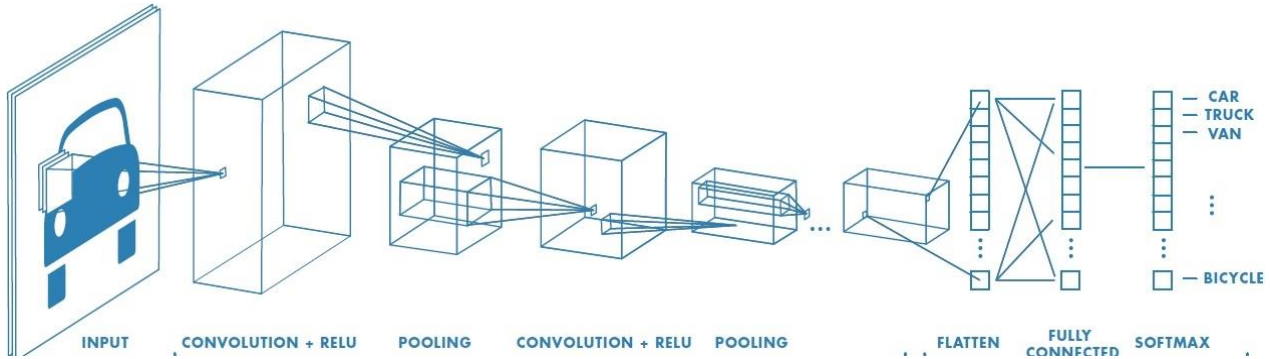

Sigmoid

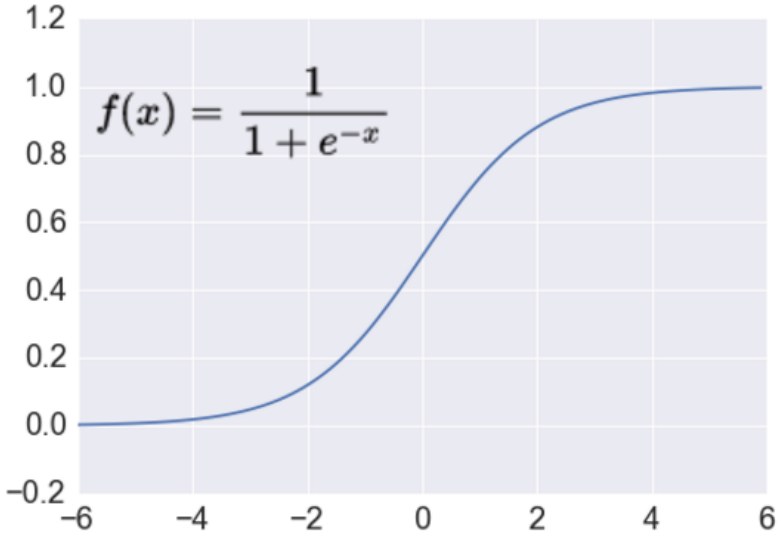

TanH

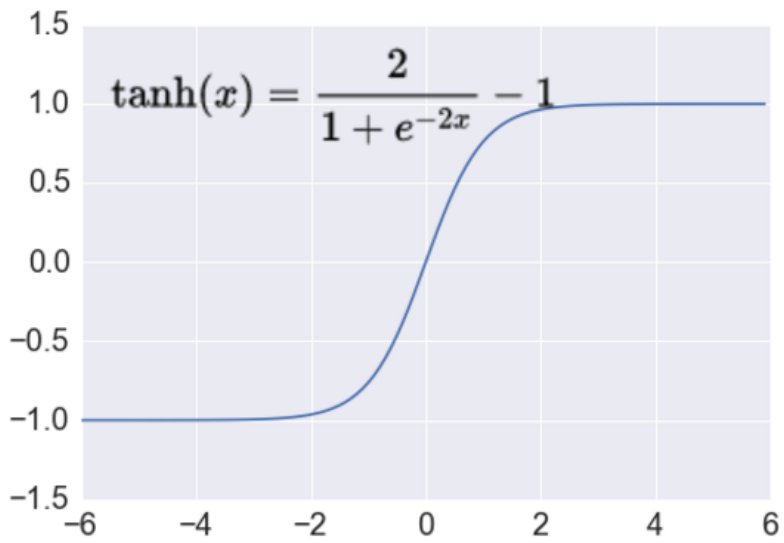

ReLU

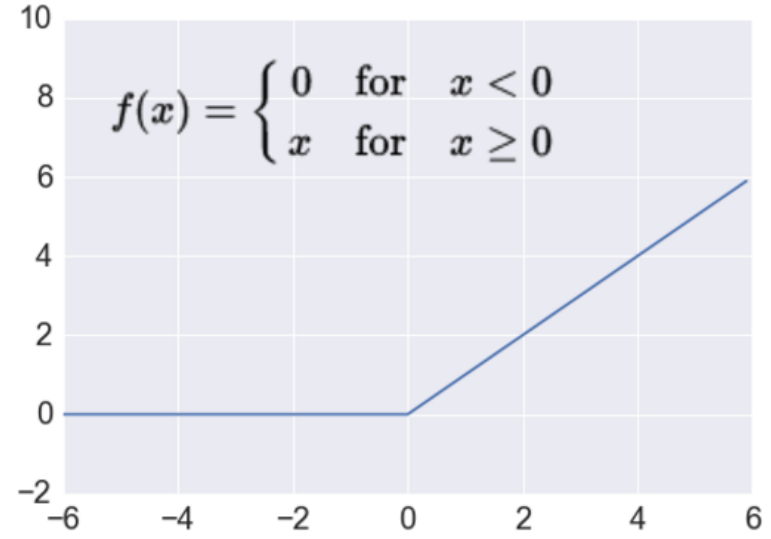

Applications – Image Classification

Convolutional Neural Network

- Convolution
- Nonlinear Activation
- Pooling (down sampling)

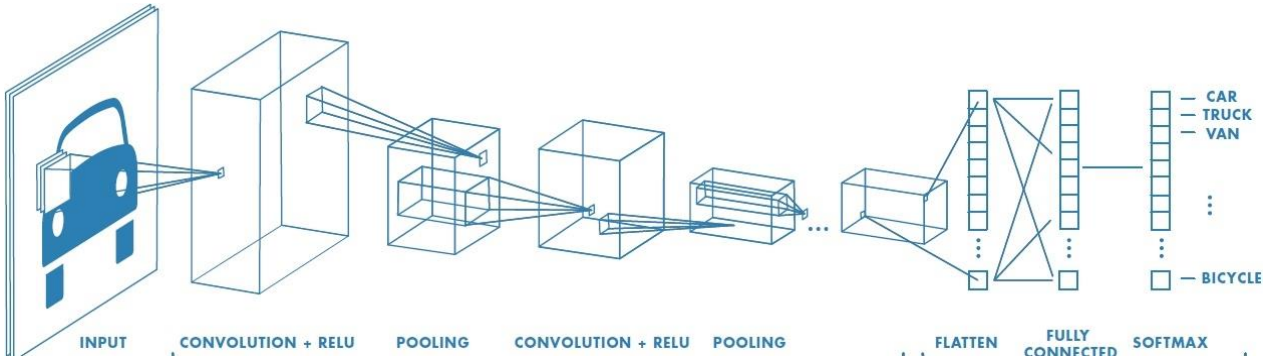

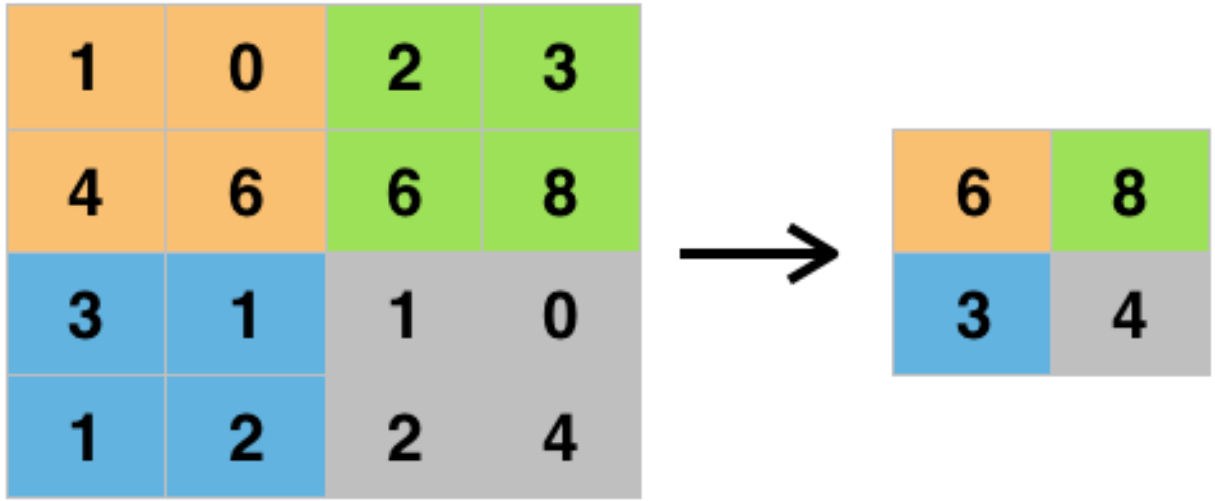

Applications – Image Classification

Compete against machine learning to predict member's age

Convolutional Neural Network

IMDB-WIKI Dataset (500k+ facial images)

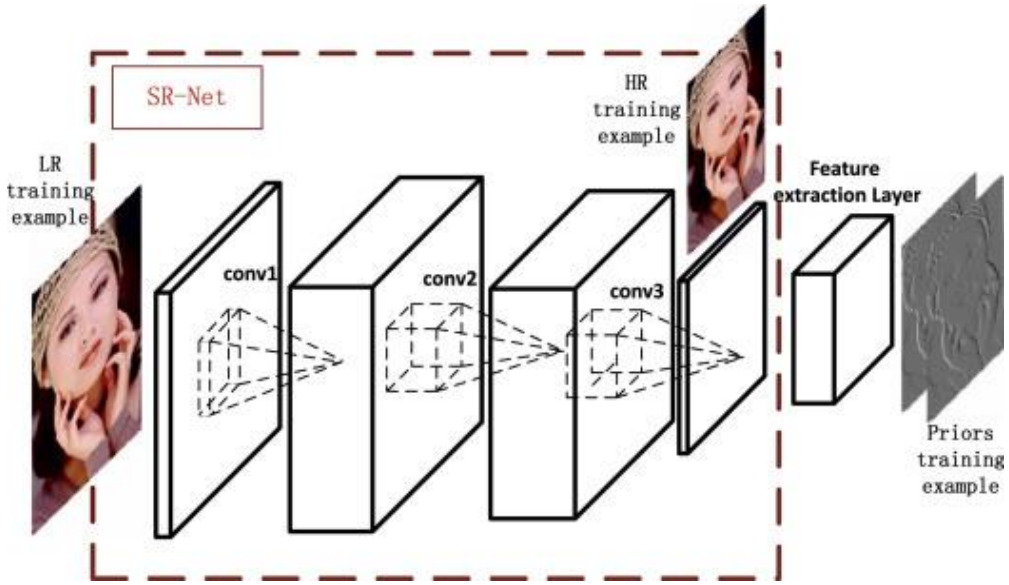

CatBoost

Speed Dating Data (2002-2004, 8000+ records)

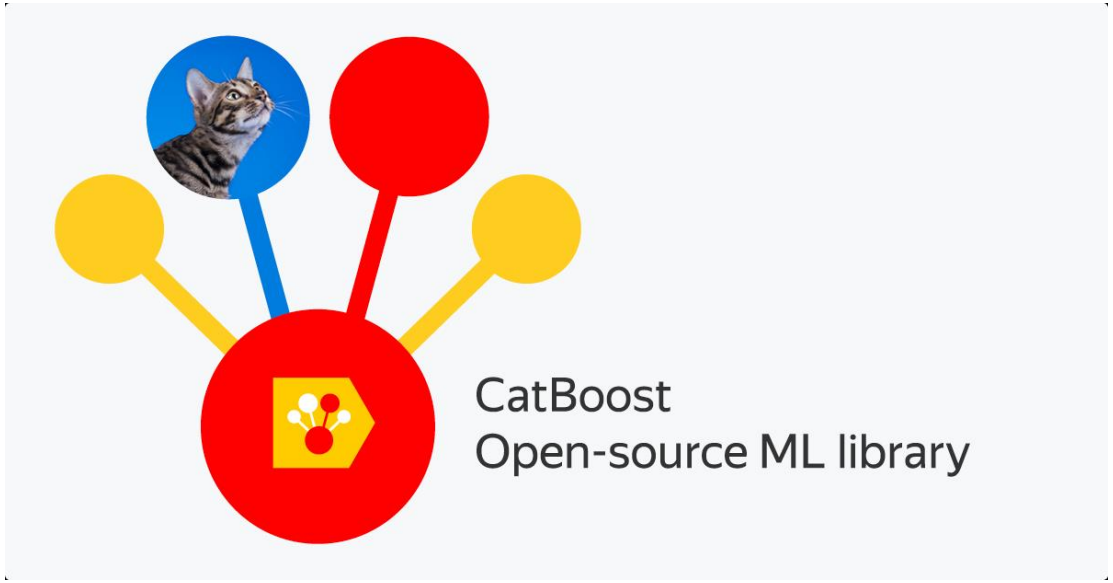

Applications – Image Classification

Compete against machine learning to predict member's age

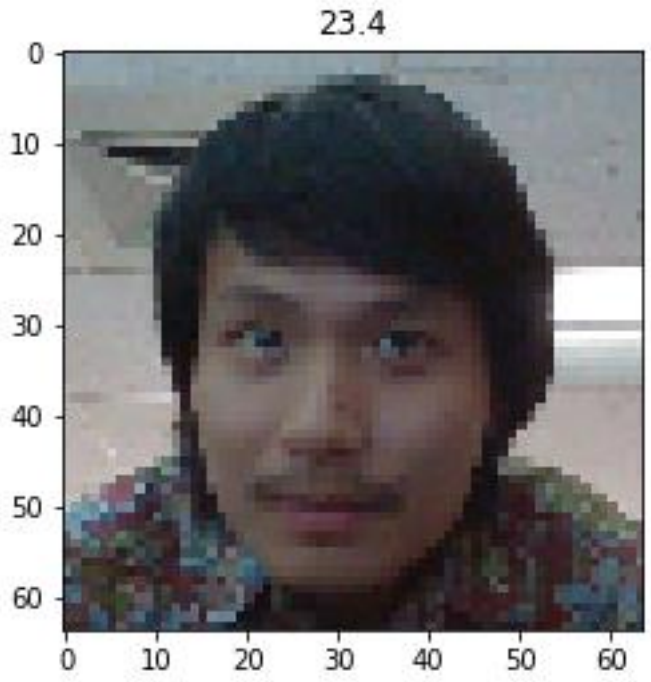

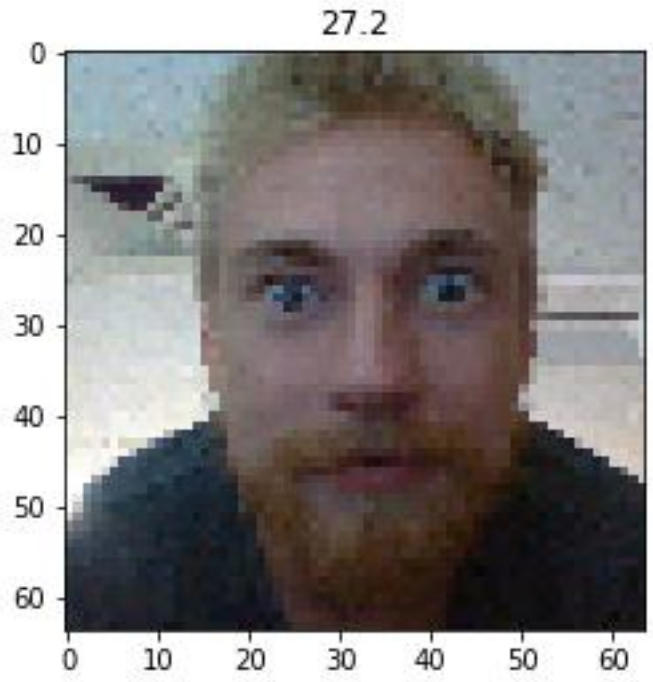

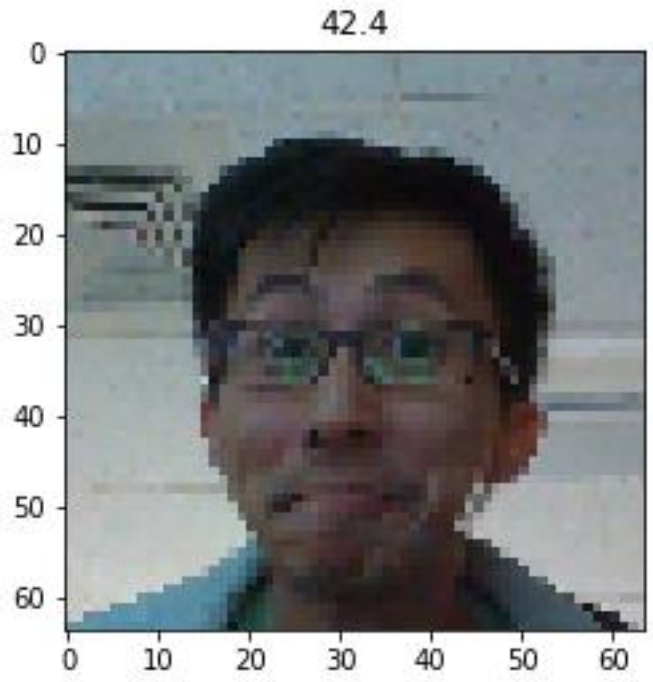

Applications – Image Classification

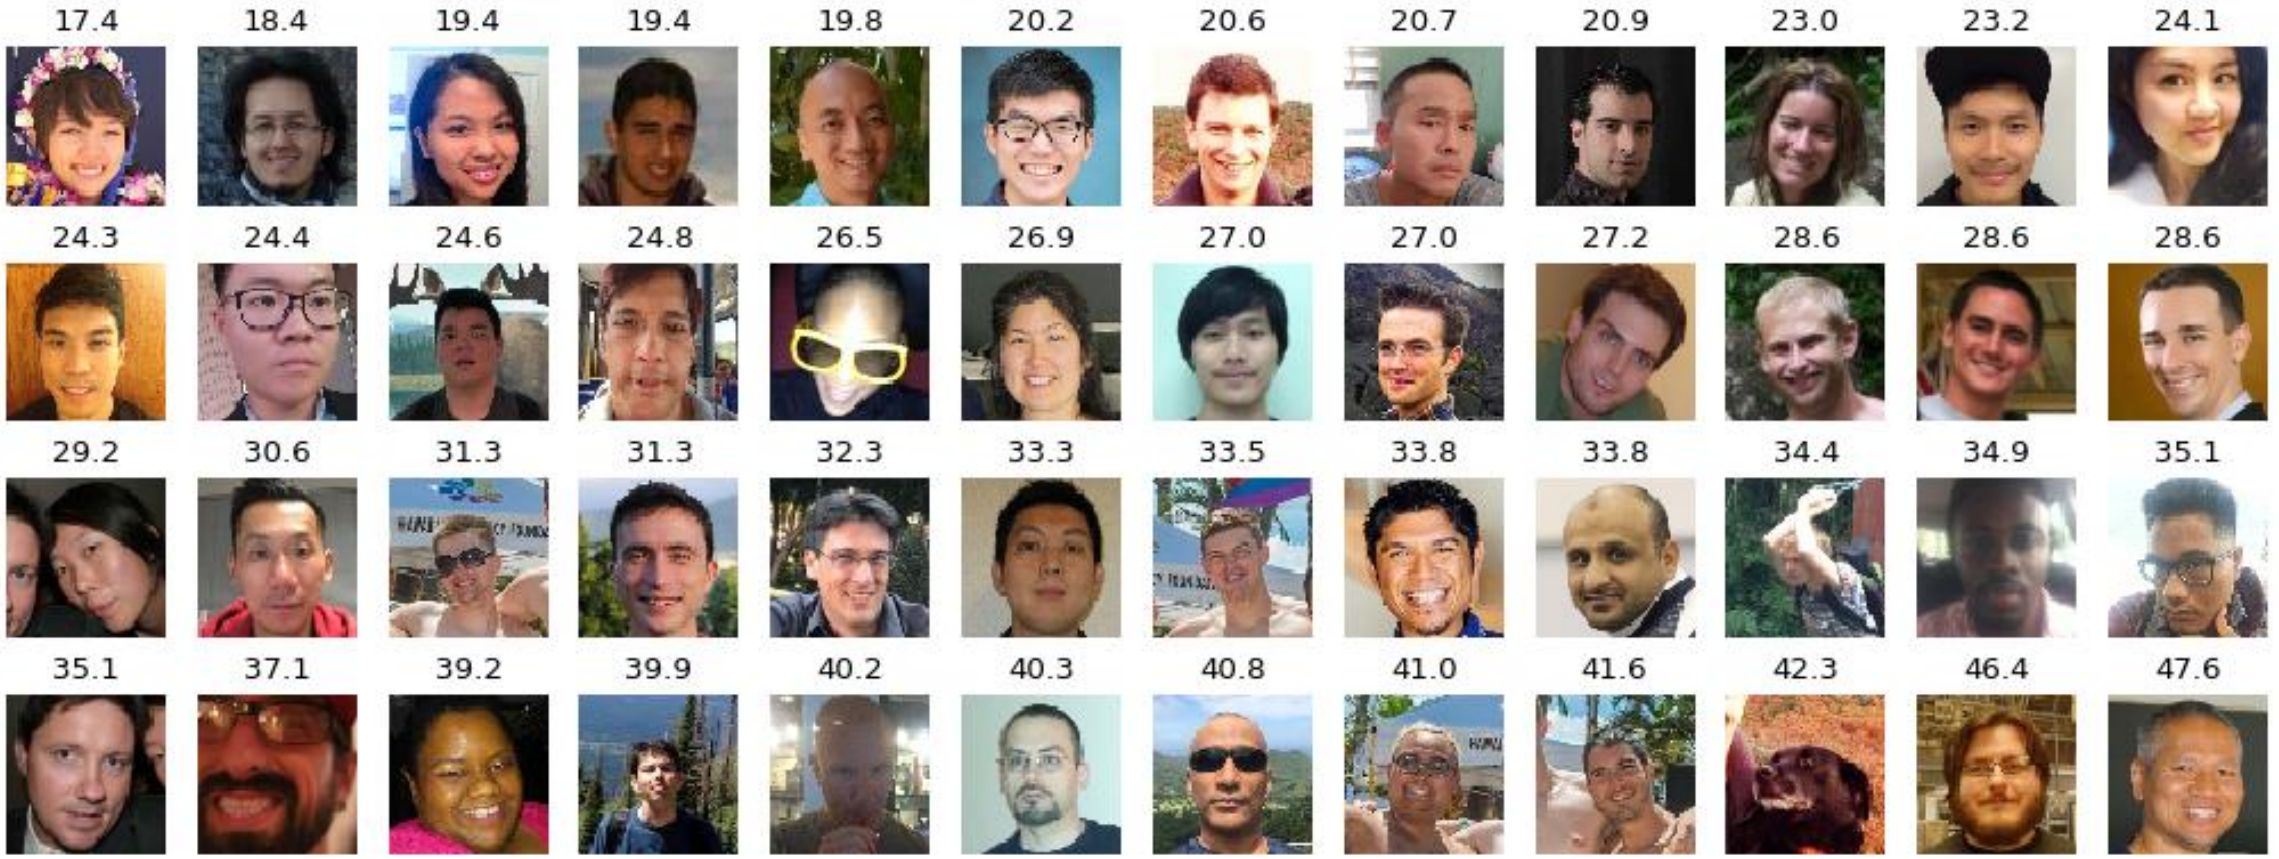

Applications – Neural Style Transfer

Deep Art - https://deepart.io

https://deepart.io/img/3AMd62EG/

content

+

style

=

Applications – Neural Style Transfer

Problem Formulation

$x_c = content\ image$

$x_s = style\ image$

$\hat{x} = output\ image$

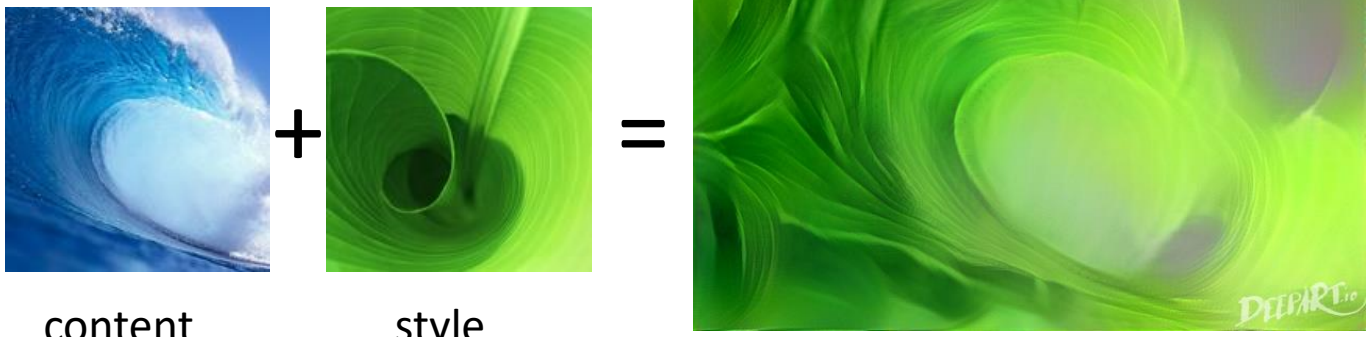

content style output

$$\underset{\hat{x}}{\text{minimize}}\, L_{total}(x_c, x_s, \hat{x})$$

$$L_{total}(x_c, x_s, \hat{x}) = L_{content}(x_c, \hat{x}) + L_{style}(x_S, \hat{x})$$

minimize with gradient descent: $\hat{x} \leftarrow \hat{x} + \alpha \nabla L_{total}(x_c, x_s, \hat{x})$

Applications – Neural Style Transfer

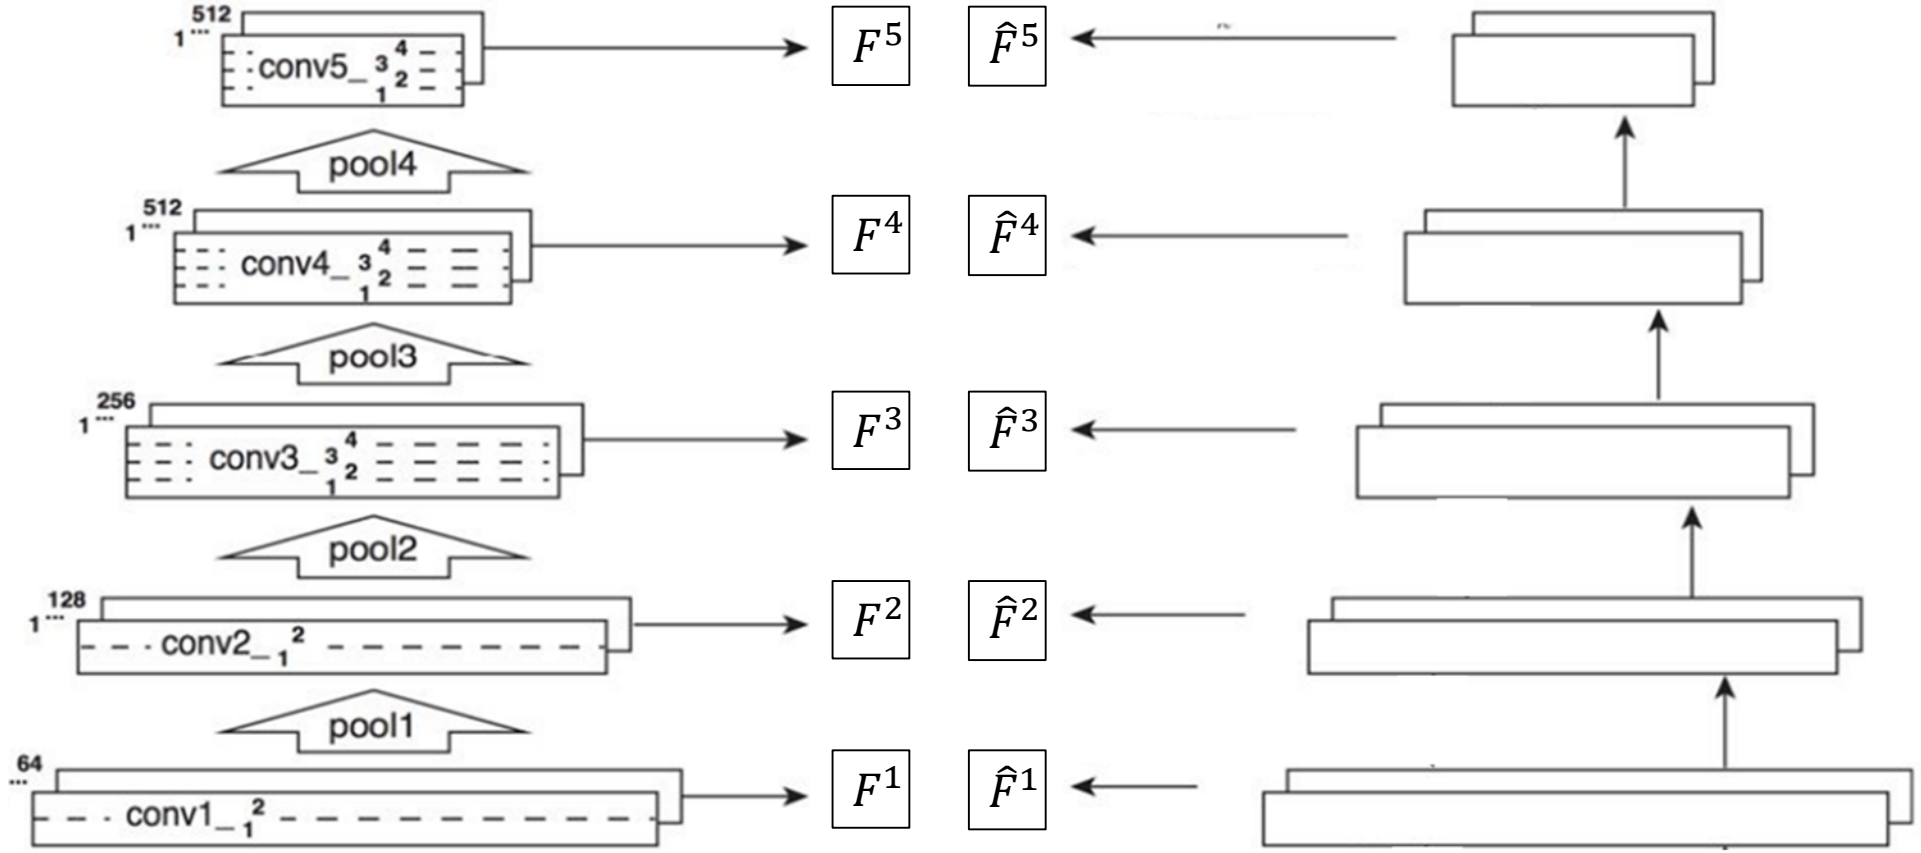

content

$$L_{content}(x_c, \hat{x}) = \sum_{l=1}^{5} \left\| F^l - \hat{F}^l \right\|^2$$

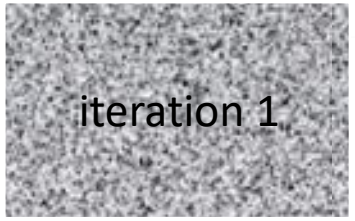

output

Applications – Neural Style Transfer

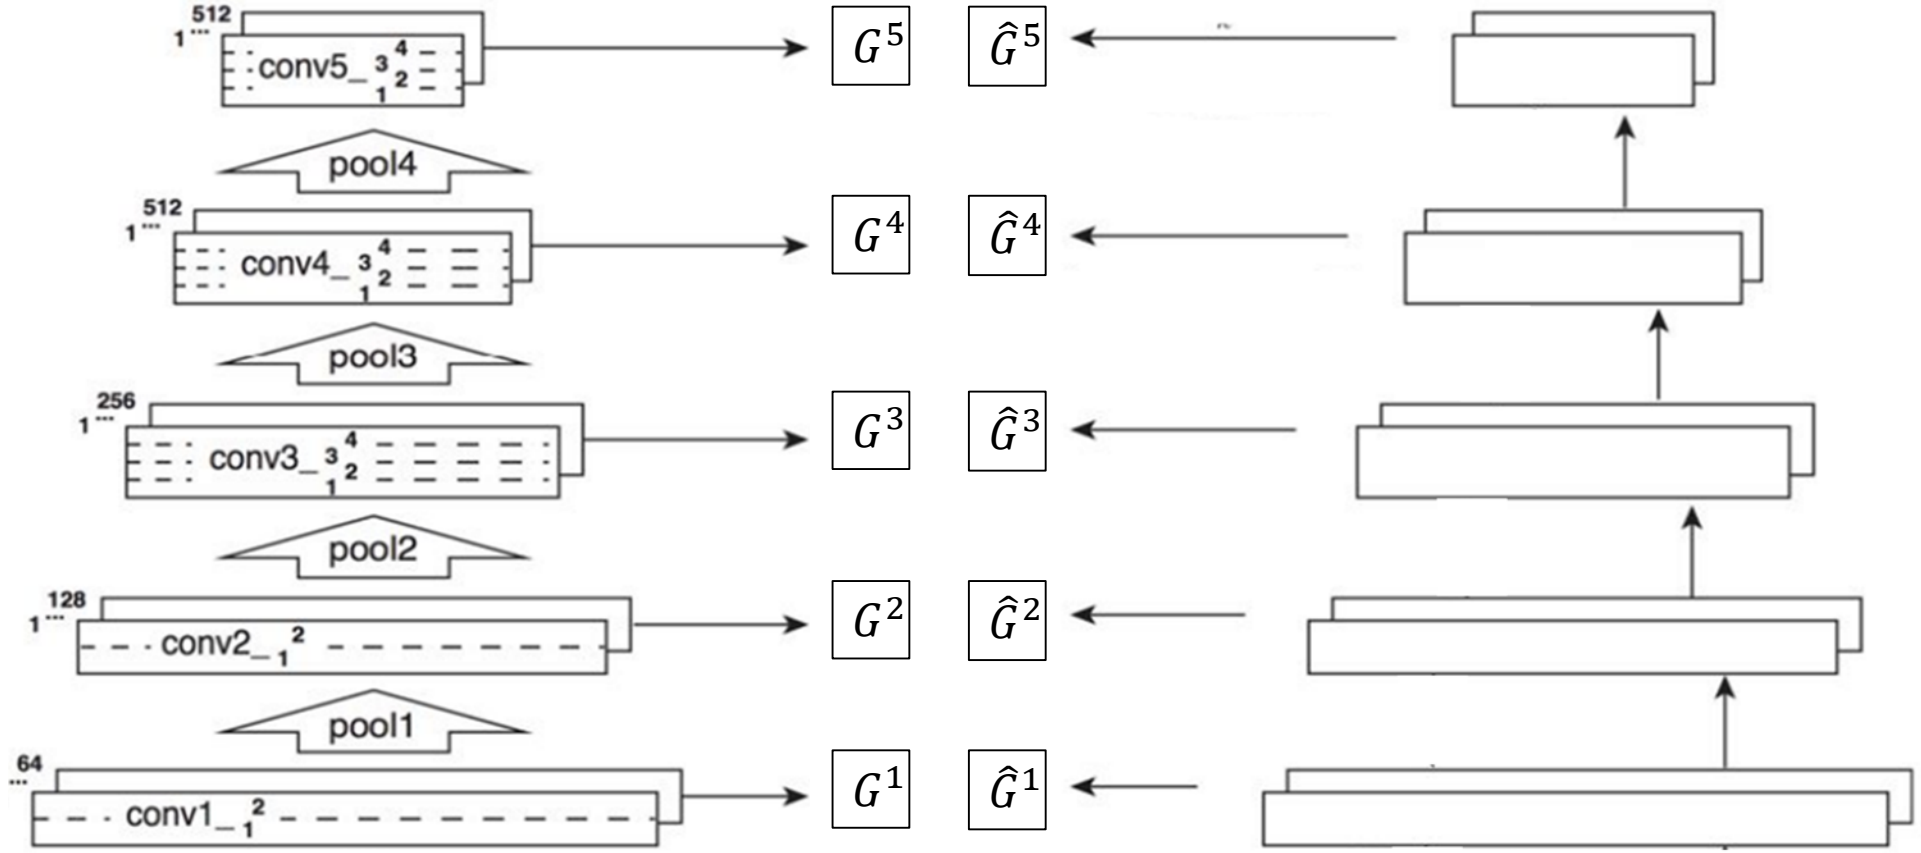

style

$$L_{style}(x_s, \hat{x}) = \sum_{l=1}^{L} w_l \left\| G^l - \hat{G}^l \right\|^2$$

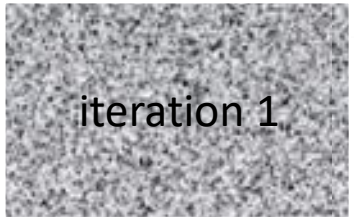

output

Applications – Neural Style Transfer

Gradient Descent

$$\hat{x} \leftarrow \hat{x} + \alpha\left[\nabla L_{content}(x_c, x_s, \hat{x}) + \nabla L_{style}(x_s, \hat{x})\right]$$

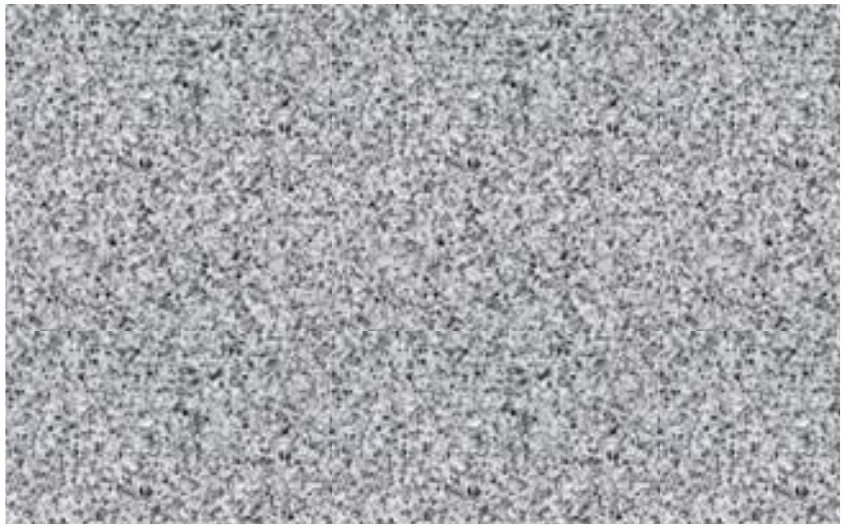

iteration: 1

...

iteration: 100

Applications – Neural Style Transfer

High Level Theory

$$L_{content}(x_c, \hat{x}) \quad + \quad L_{style}(x_s, \hat{x}) \quad = \quad L_{total}(x_c, x_s, \hat{x})$$

Applications – Neural Style Transfer

Deep Art - https://deepart.io

https://deepart.io/img/f0rqrVR8/

content

+

style

=

Applications – Neural Style Transfer

Deep Art - https://deepart.io

https://deepart.io/img/JYHGsxl41/

content

+

style

=

Applications – Neural Style Transfer

Deep Art - https://deepart.io

https://deepart.io/img/VNuNaYbU/

content

+

style

=

Applications – Neural Style Transfer

Deep Art - https://deepart.io

https://deepart.io/img/4woyLt0B/

content

+

style

=

Applications – Neural Style Transfer

Deep Art - https://deepart.io

https://deepart.io/img/fREjk8nP/

content

+

style

=

Applications – Neural Style Transfer

Deep Art - https://deepart.io

https://deepart.io/img/G5CWntgd1/

content

+

style

=

Applications – Neural Style Transfer

Deep Art - https://deepart.io

https://deepart.io/img/ZkujbVDb/

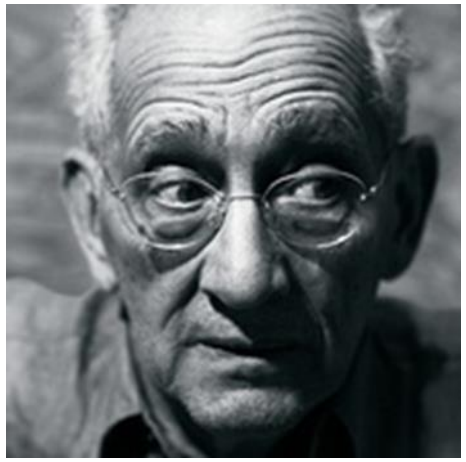

content

+

style

=

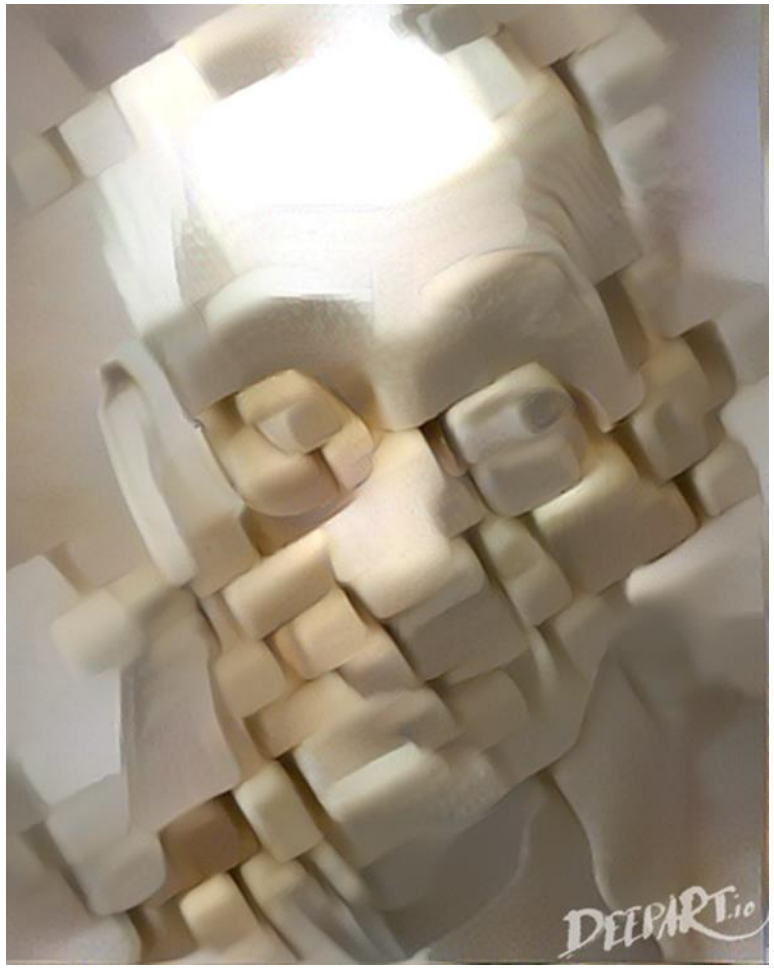

Applications – Language Modeling

Recurrent Neural Networks (RNN)

Many to One

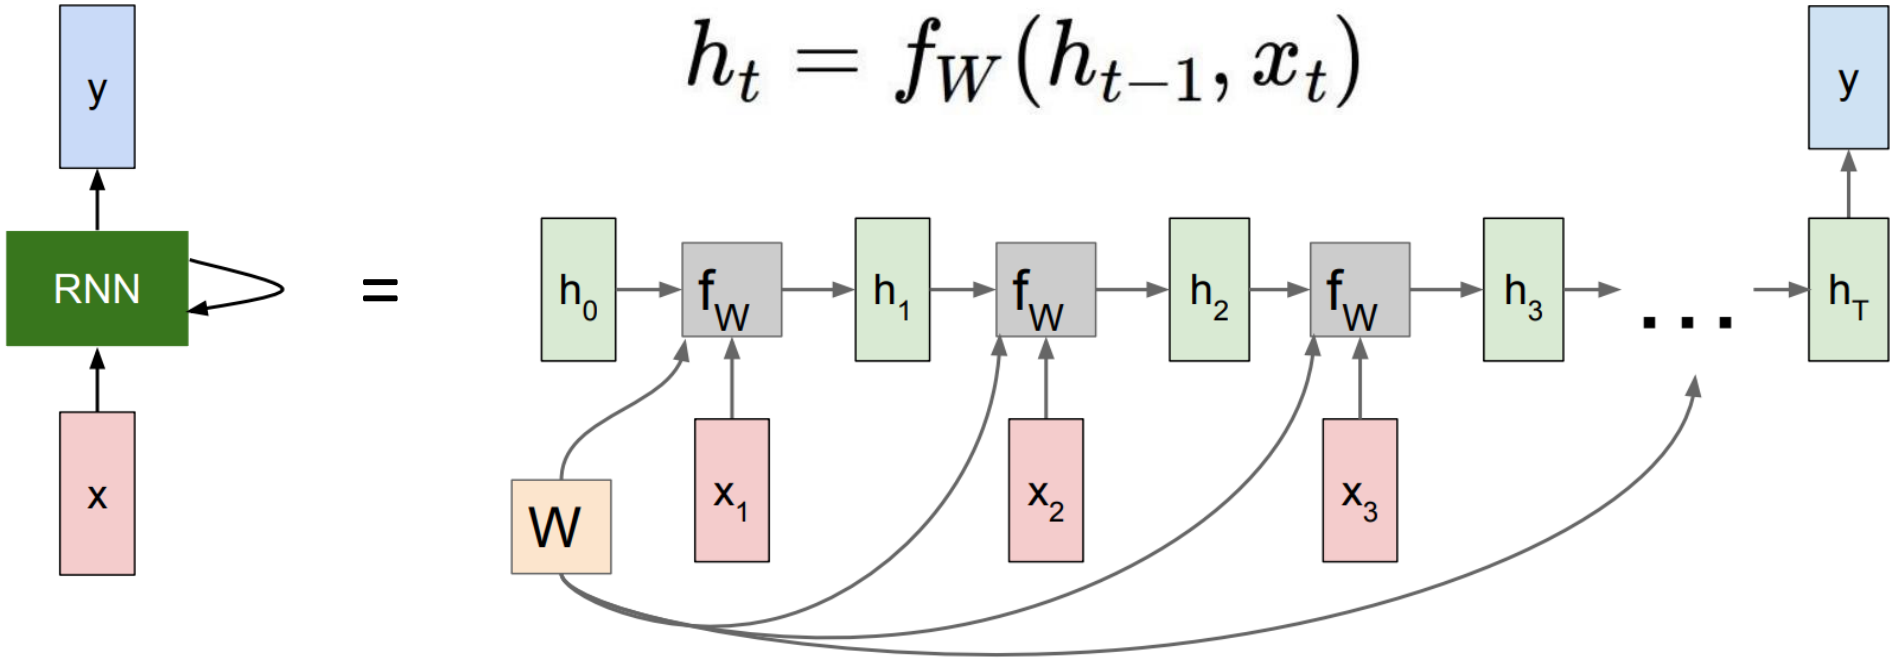

Applications – Language Modeling

Recurrent Neural Networks (RNN)

Many to Many

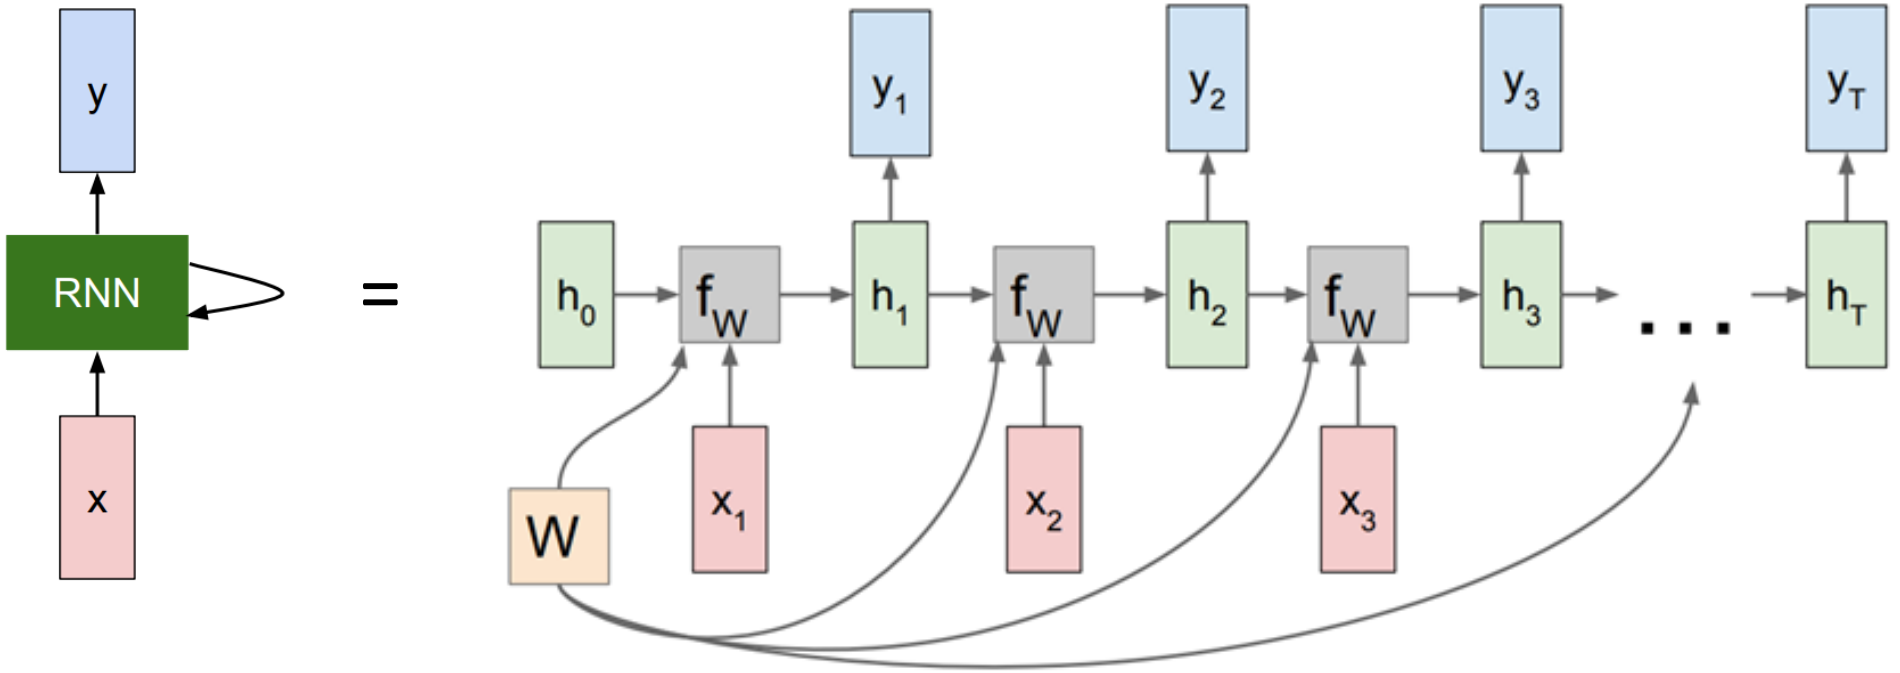

Applications – Language Modeling

Recurrent Neural Networks (RNN)

One to Many

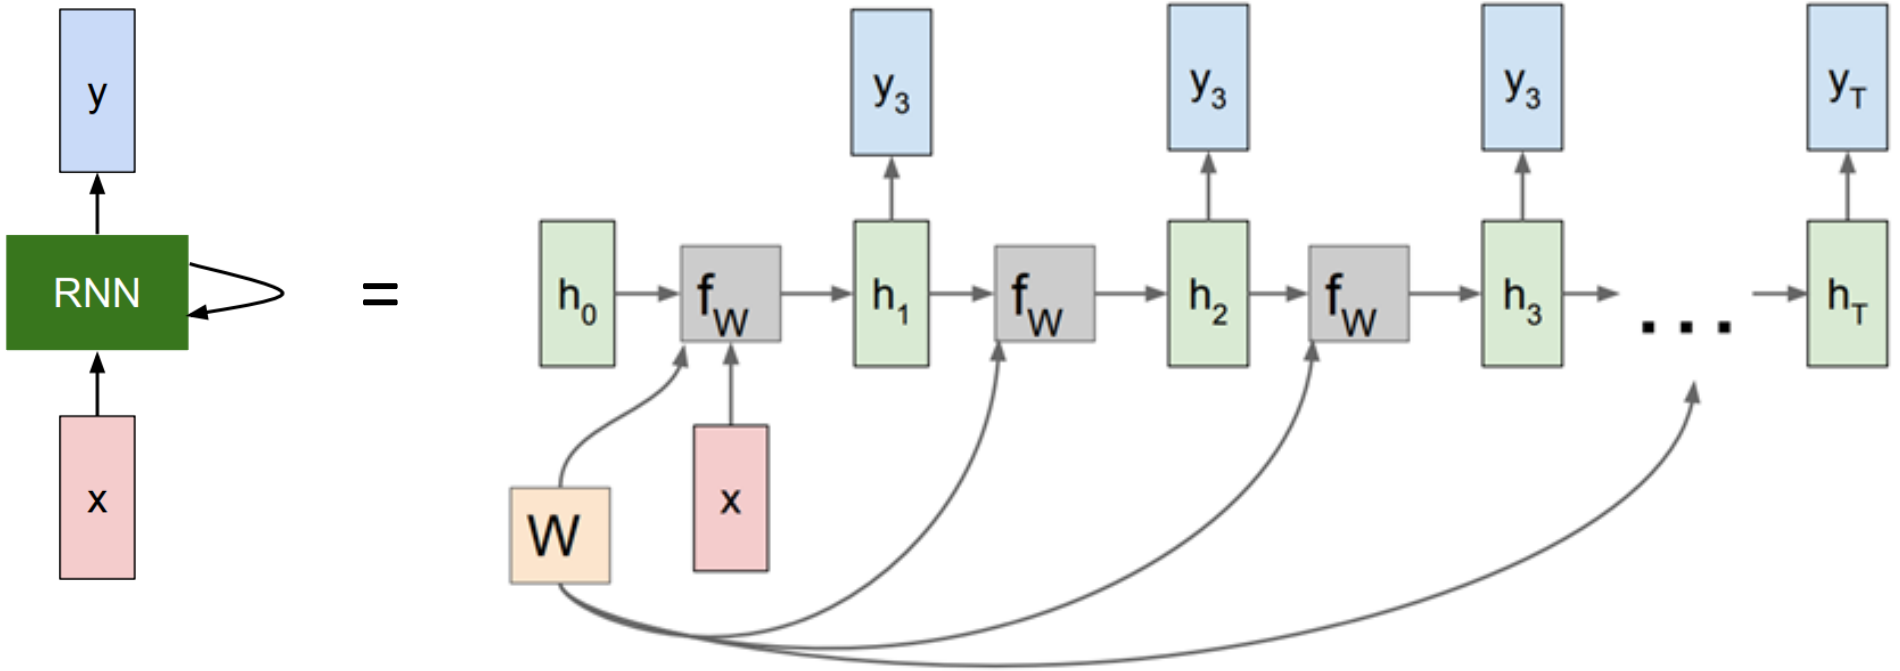

Applications – Language Modeling

Long Short Term Memory (LSTM)

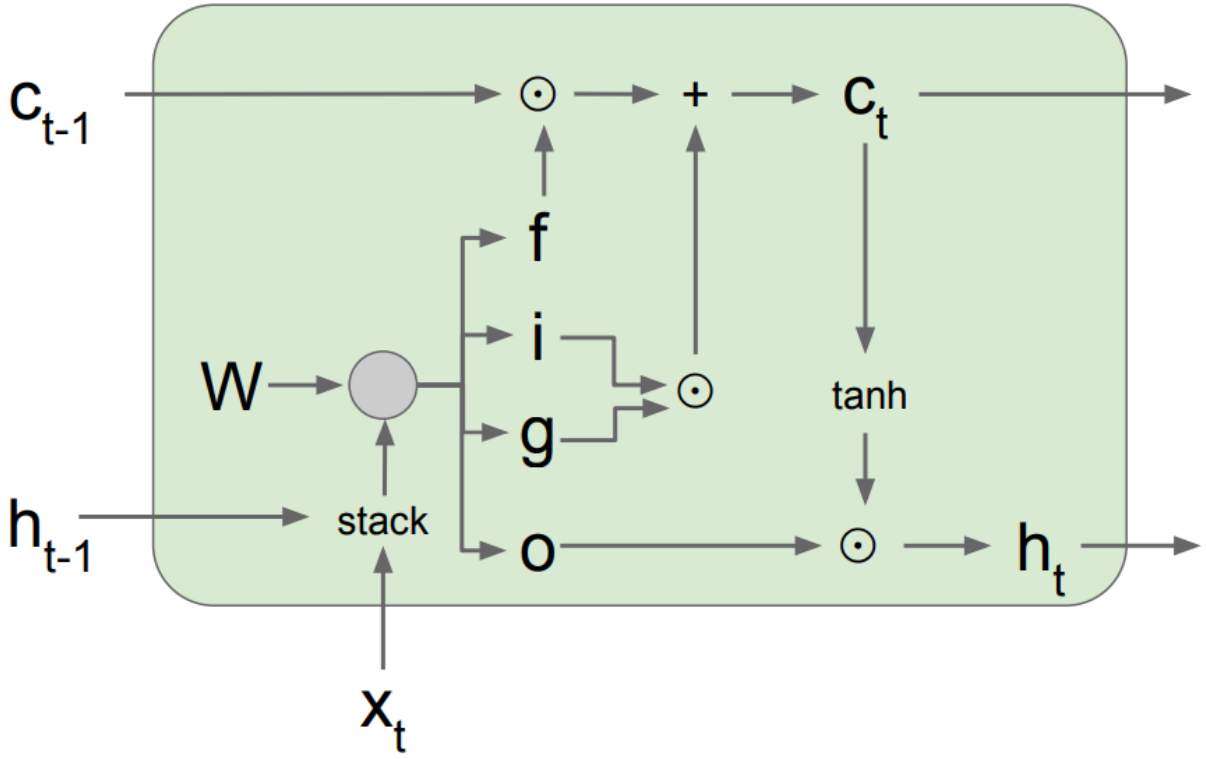

$$\begin{pmatrix} i \\ f \\ o \\ g \end{pmatrix} = \begin{pmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{pmatrix} W \begin{pmatrix} h_{t-1} \\ x_t \end{pmatrix}$$

$$c_t = f \odot c_{t-1} + i \odot g$$

$$h_t = o \odot \tanh(c_t)$$

Applications – Language Modeling

Pidgin Bible

- http://www.pidginbible.org/Concindex.html
- 15,891 verses 2.5 million characters (59 unique characters)
- Sample verses:

  02.7.3 *But firs I goin make da Pharaoh guy hard head. Even I goin do plenny awesome tings fo erybody see, hea inside da Egypt land.*

  16.7.72 *Da odda peopo give 20,000 gold coin, 2,500 pound silva, an 67 robe fo da pries guys.*

Applications – Language Modeling

Generated Pidgin Bible Text

i wen to da hawaii machine learning meetup to see the peopo dat talk about da bad kine stuff dat god wen tell um fo do. da guy dat goin come back alive on top da boy dat tell um, “you guys goin go look jalike da boss jesus.

...hand to da sky, fo make um come dark all ova da egypt land, i goin stay ready fo do da right ting an say, “i no can do real nice fo da guys dat stay alive, an da big strong peopo dat wen go back an go outside da sky, an same ting dat get plenny love an aloha...

Applications – Other Applications

More Applications & Techniques

- Machine translation
- Image segmentation
- Autoencoders
- Word embeddings
- Image captioning

More Applications & Techniques

- Machine translation

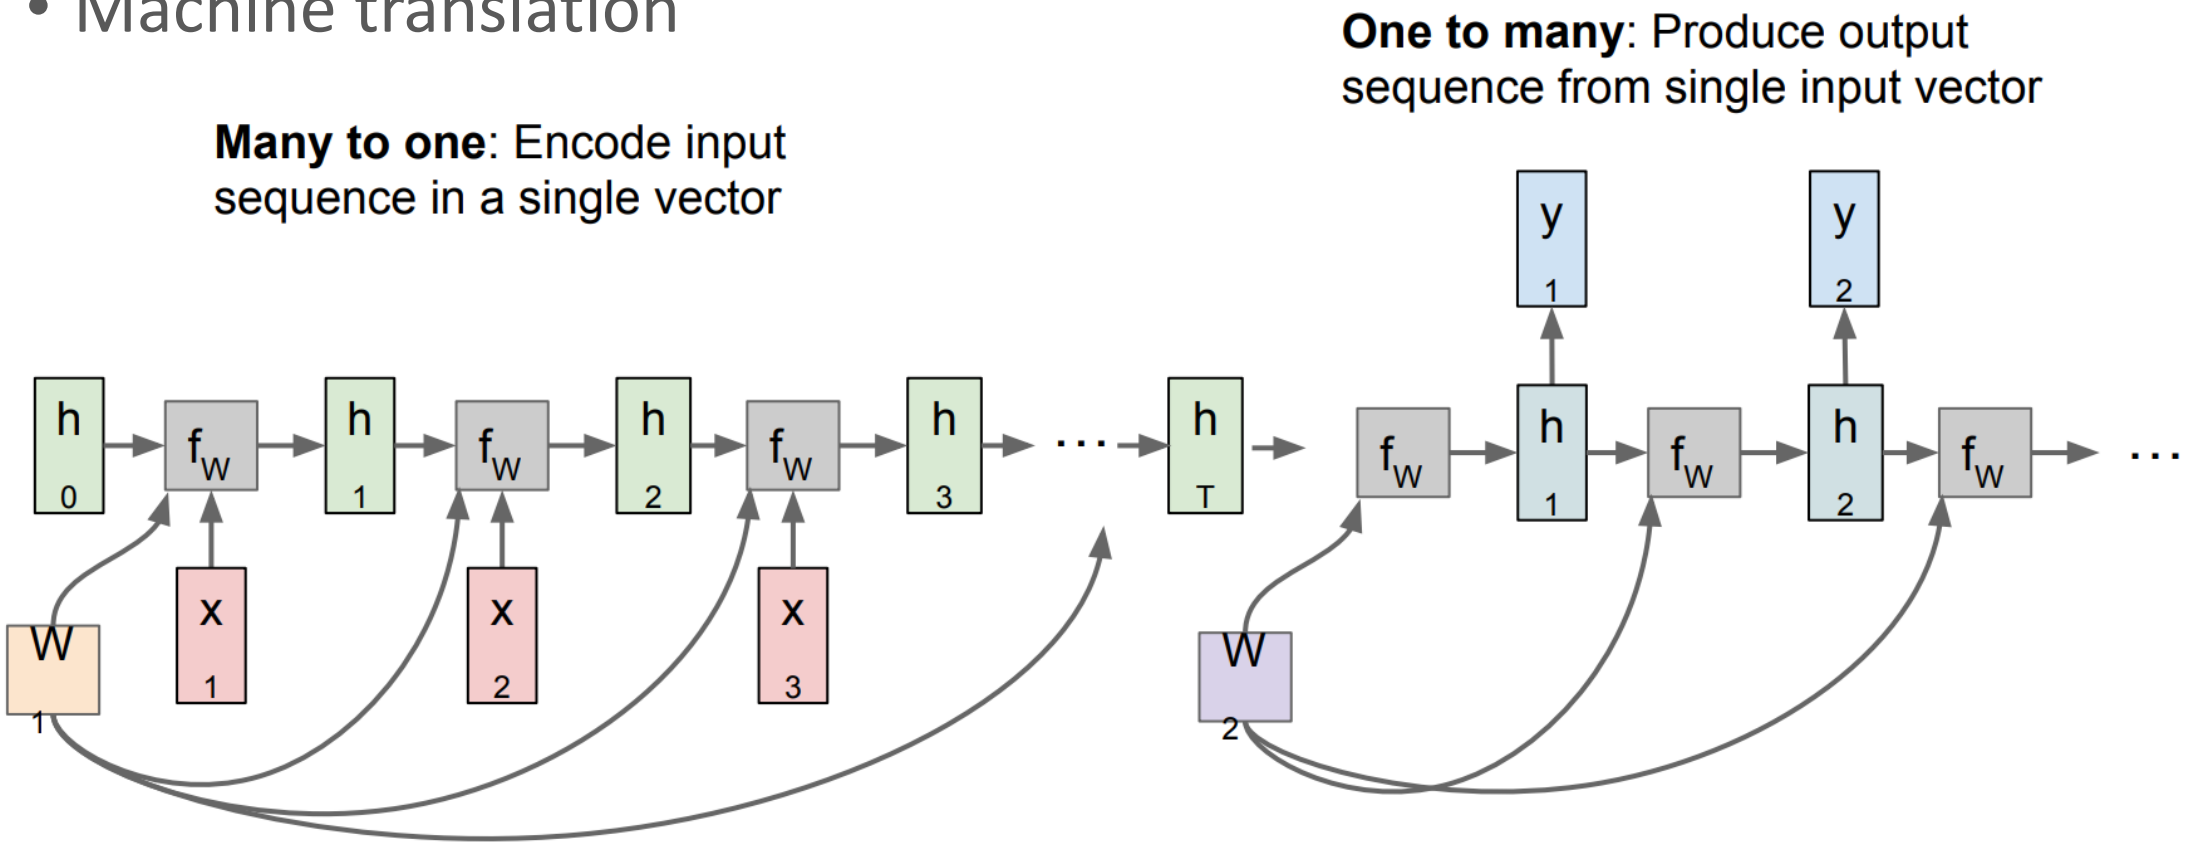

Applications – Other Applications

More Applications & Techniques

- Machine translation
- Image segmentation (masking)

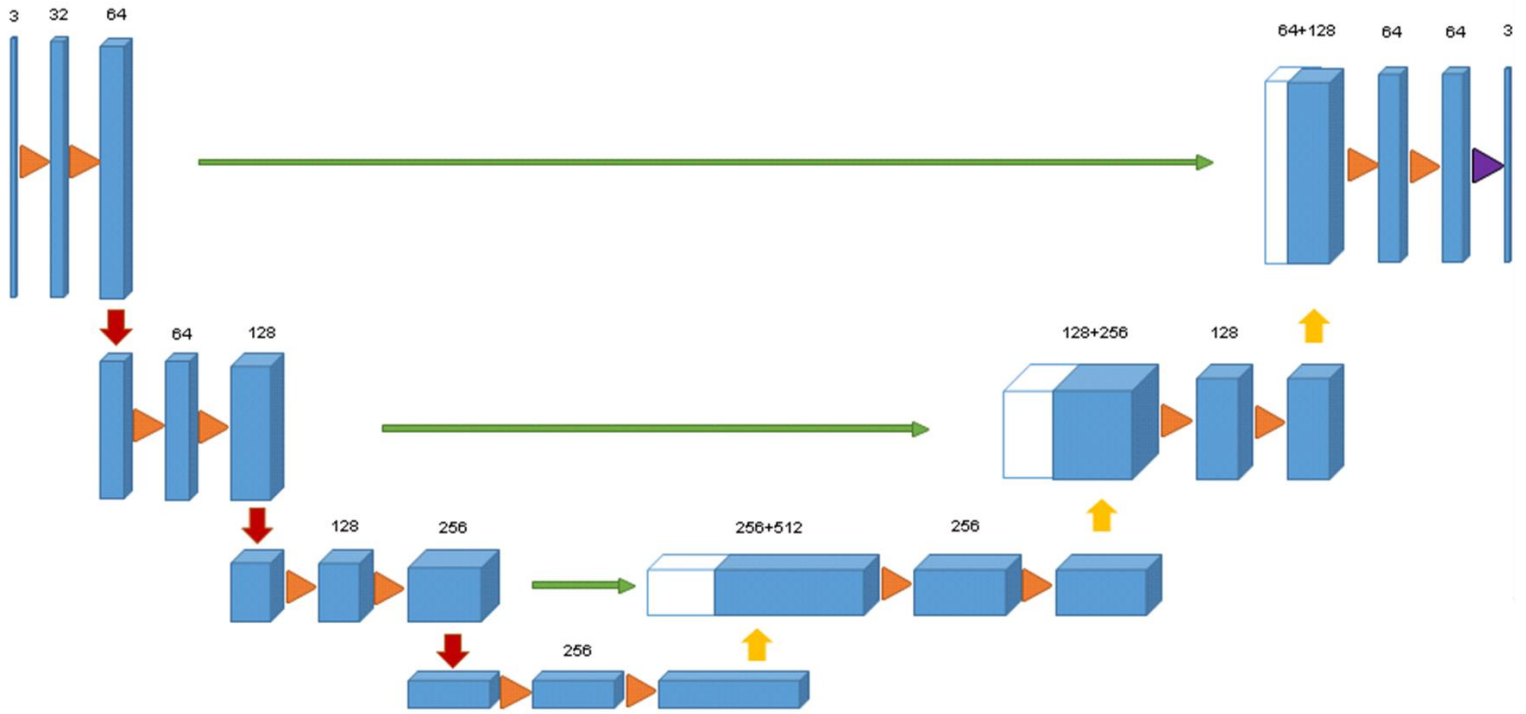

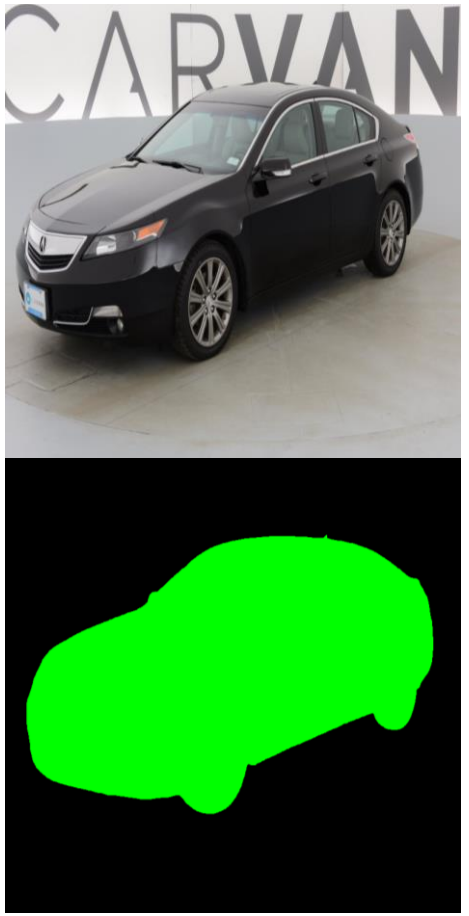

Applications – Other Applications

More Applications & Techniques

- Machine translation
- Image segmentation

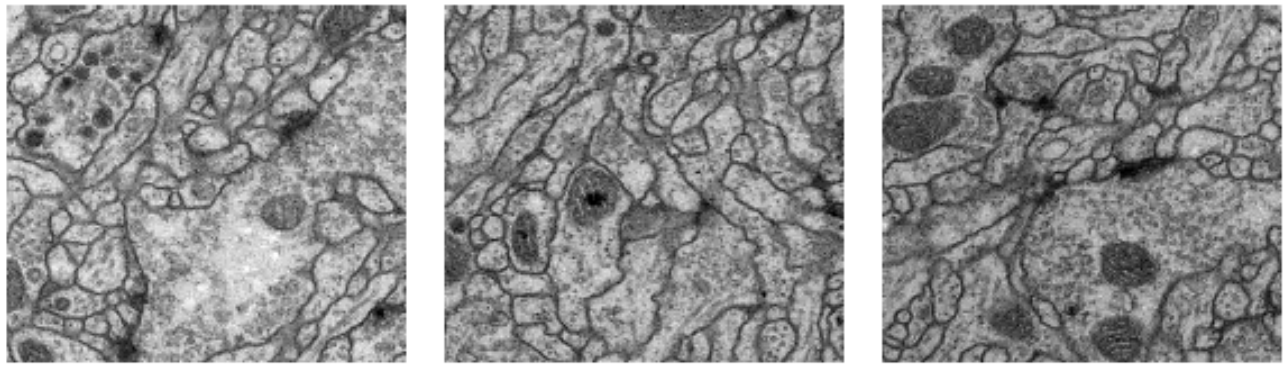

(a) Original images

(b) Segmentation results

Applications – Other Applications

More Applications & Techniques

- Machine translation
- Image segmentation
- Autoencoders
  - Feature engineering
  - Dimensionality reduction
  - Anomaly detection
  - Binary classification

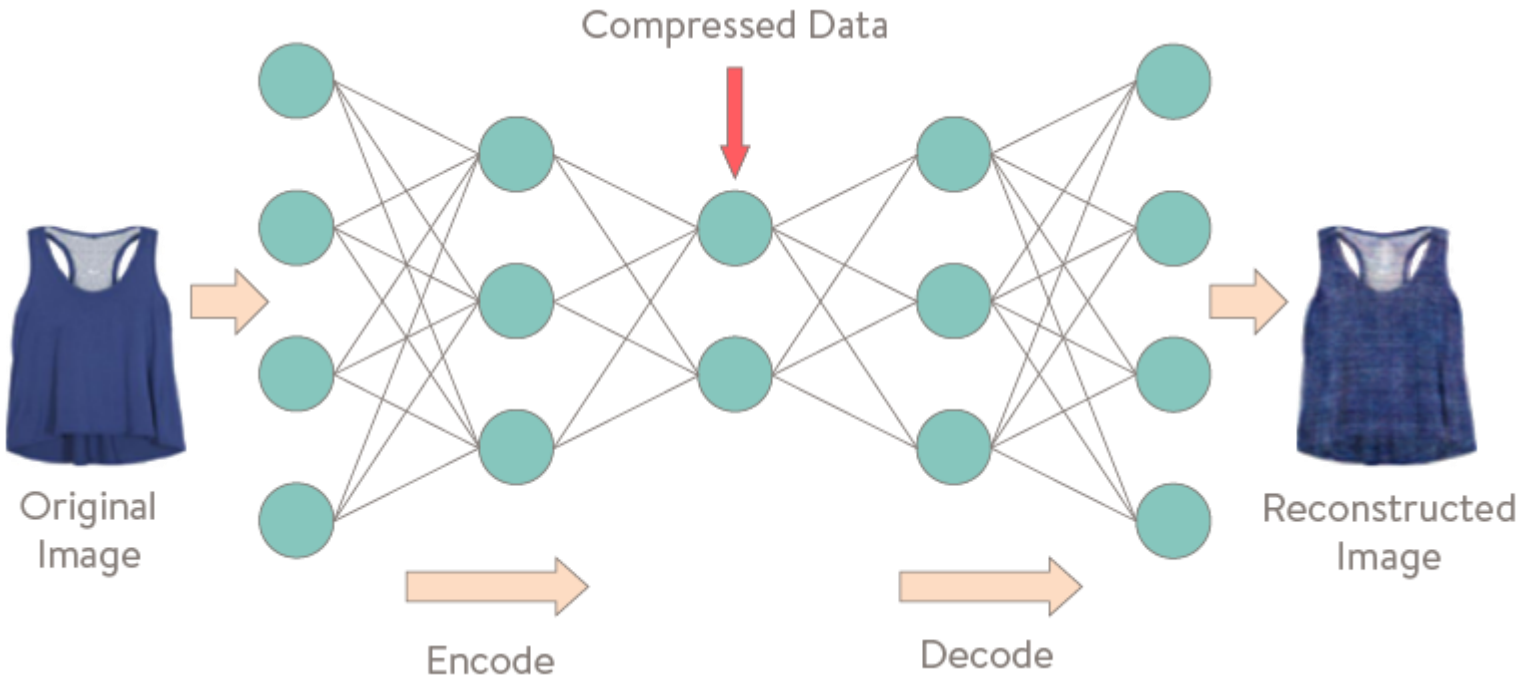

Applications – Other Applications

More Applications & Techniques

- Machine translation
- Image segmentation
- Autoencoders
- Word embeddings

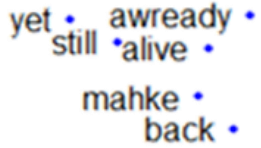

Text → Numeric

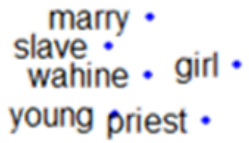

Applications – Other Applications

More Applications & Techniques

- Machine translation
- Image segmentation
- Autoencoders
- Word embeddings
- Image captioning

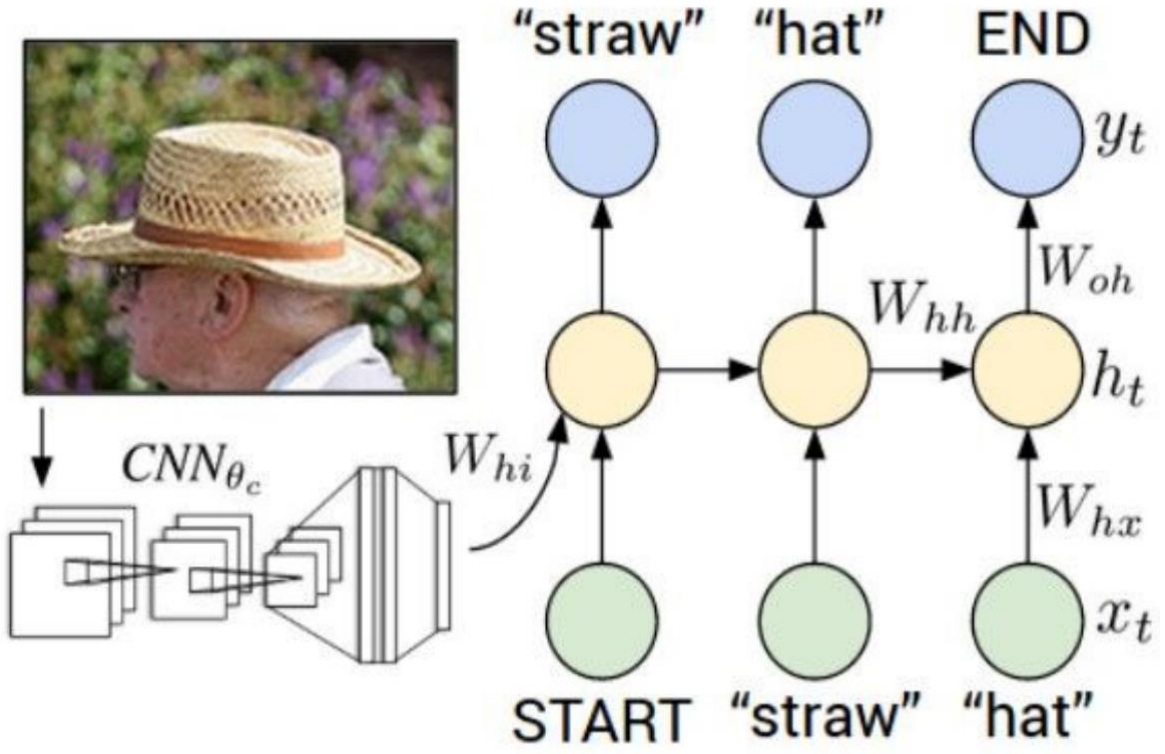

Table of Contents

Resources

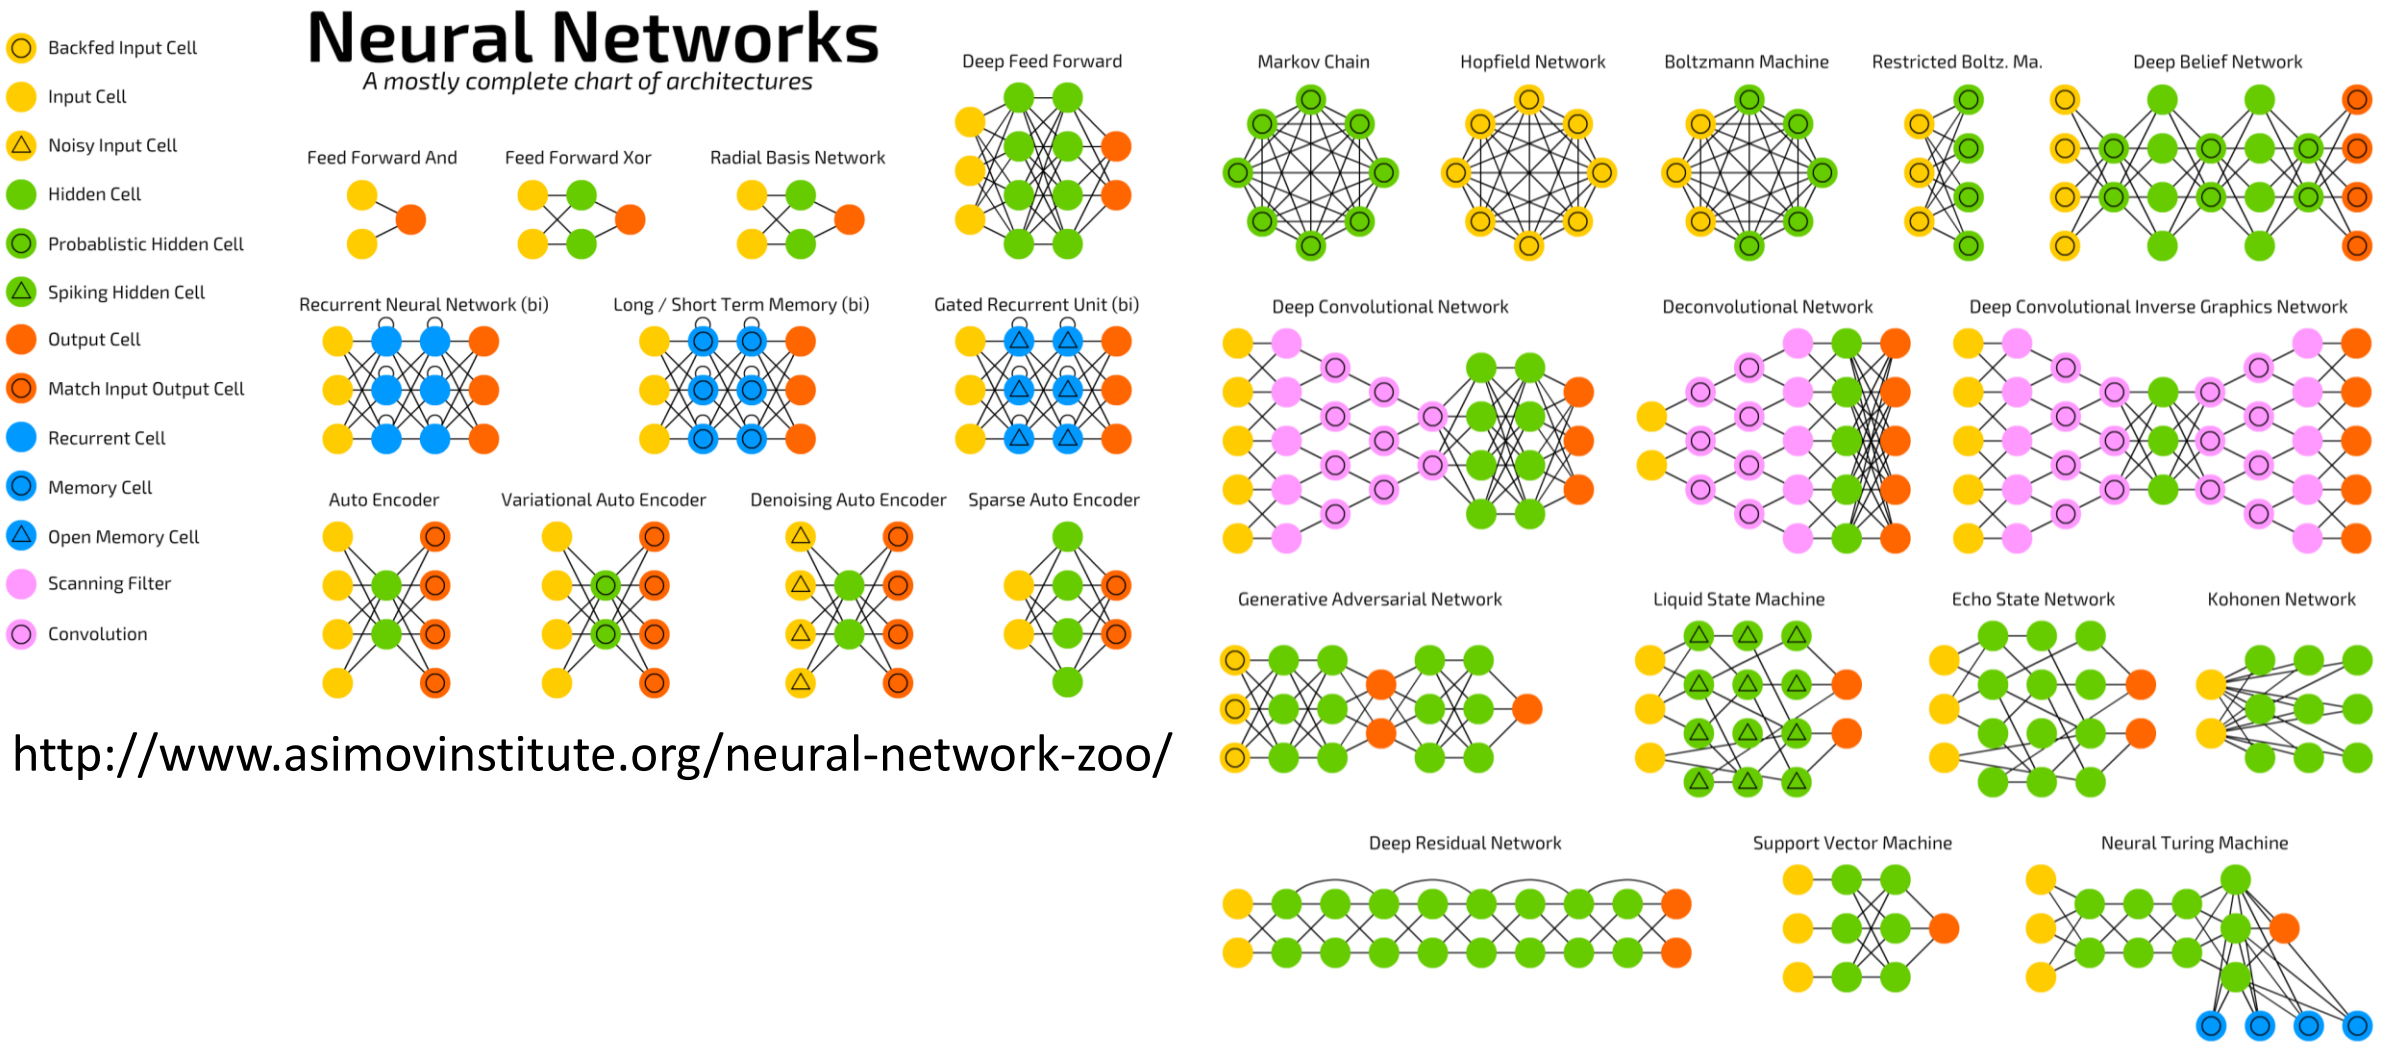

http://www.asimovinstitute.org/neural-network-zoo/

Resources

Mostly Free Online Courses

- **Coursera:**
  Neural Networks for Machine Learning by Geoff Hinton
  Neural Networks and Deep Learning by Andrew Ng (*not free*)
- **YouTube:**
  Deep learning at Oxford 2015 by Nando de Freitas
  Convolutional Neural Networks for Visual Recognition by Stanford
  Natural Language Processing with Deep Learning by Stanford
- **Udacity:** Deep Learning by Google
- **fast.ai:** Deep Learning Part 1,2 by Jeremy Howard

Resources

- **Convolutional Neural Networks**
  http://cs231n.github.io/convolutional-networks/
  https://github.com/asmith26/wide_resnets_keras
  https://lmb.informatik.uni-freiburg.de/people/ronneber/u-net/
- **Neural Style Transfer**
  https://arxiv.org/abs/1508.06576
  https://github.com/lengstrom/fast-style-transfer
- **Recurrent Neural Networks**
  http://karpathy.github.io/2015/05/21/rnn-effectiveness/
  https://github.com/tensorflow/nmt/blob/master/README.md
  http://colah.github.io/posts/2015-08-Understanding-LSTMs/

Resources

- Hawaii Machine Learning Website
  - https://hawaii-ml.github.io/
  - Previous meetup's slides
- hawaii-ml Slack
- Future Meetups
  - Convolutional Neural Networks
  - Recurrent Neural Networks
- PyHawaii TensorFlow presentation

Conclusion

Share your Experiences

- Do you have experience using neural nets?
- Do you know of any interesting applications not mentioned yet?
- What applications of neural nets do you want to hear about at future meetups?

Machine Learning in R

Chris Sugai

- Why not Excel?
- Tidyverse
- data.table
- R Markdown

Machine Learning in Python?