A Brief Introduction to Deep Learning

Labradoodle or fried chicken

Puppy or bagel

Sheepdog or mop

Chihuahua or muffin

Barn owl or apple

Parrot or guacamole

But, we human actually lose!

- A demo that shows we, human, lose, on the classification task, we are proud of, we have been trained for millions of years!
- If we want to make it hard for bots, it has to be hard for human as well.

We human lose on Go!

We (will) lose on many specific tasks!

- Speech recognition
- Translation
- Self-driving
- …

- BUT, they are not AI yet…
- Don’t worry until it dates with your girl/boy friend…

Machine Learning vs Deep Learning

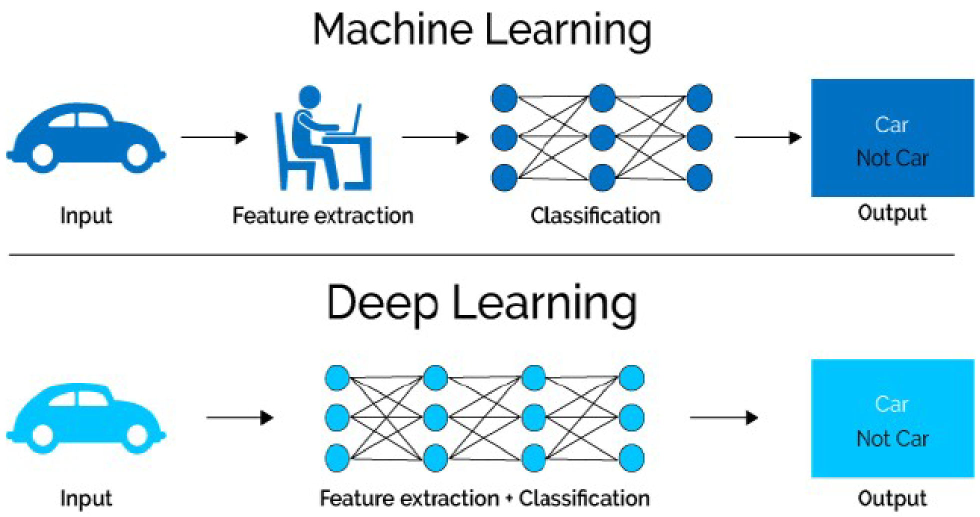

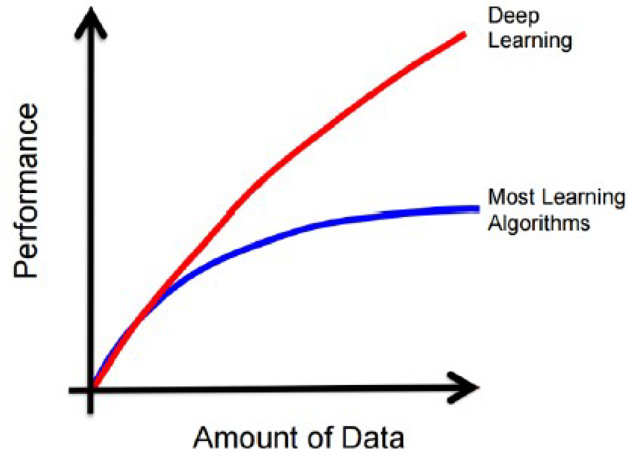

A Brief Introduction to Deep Learning

- Artificial Neural Networks
- Convolutional Neural Networks (CNN)
- Recurrent Neural Networks (RNN)
- AutoEncoder
- Generative Adversarial Networks (GAN)

Artificial Neural Network

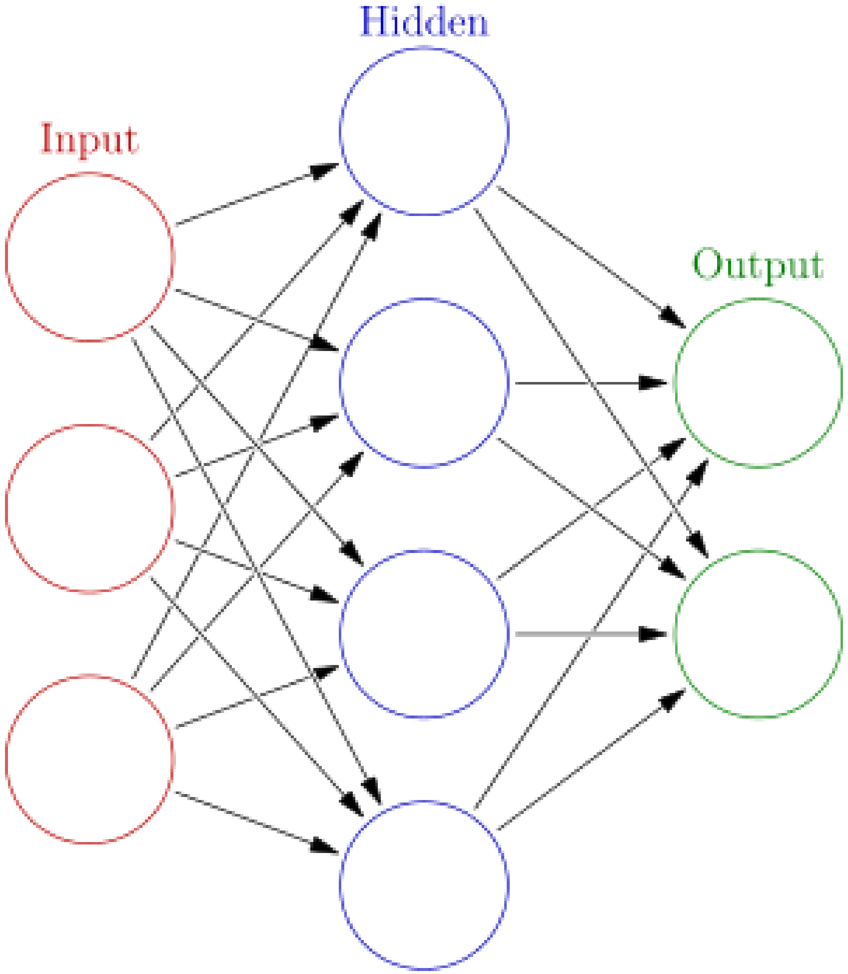

1. Activation function
2. Weights
3. Cost function
4. Learning algorithm

[Live Demo]

Now, serious stuff, a bit…

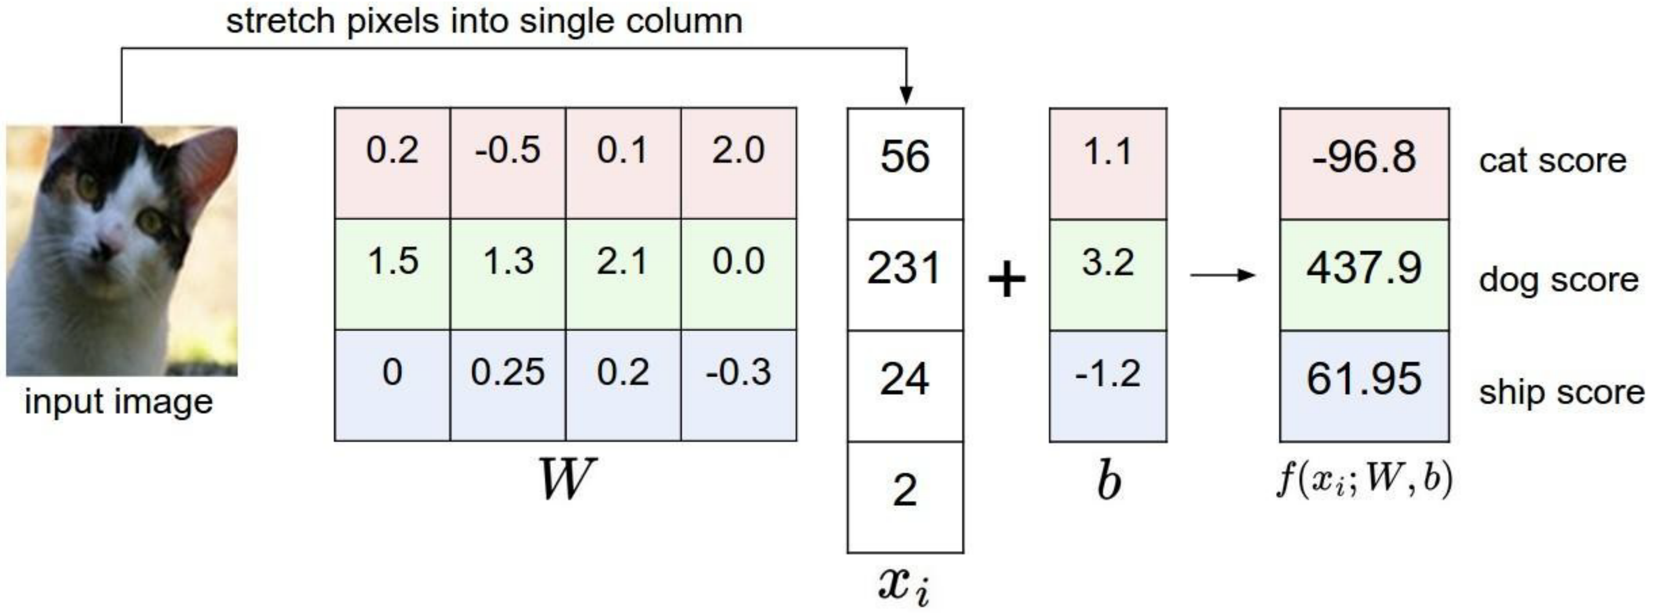

Fully Connected Layers

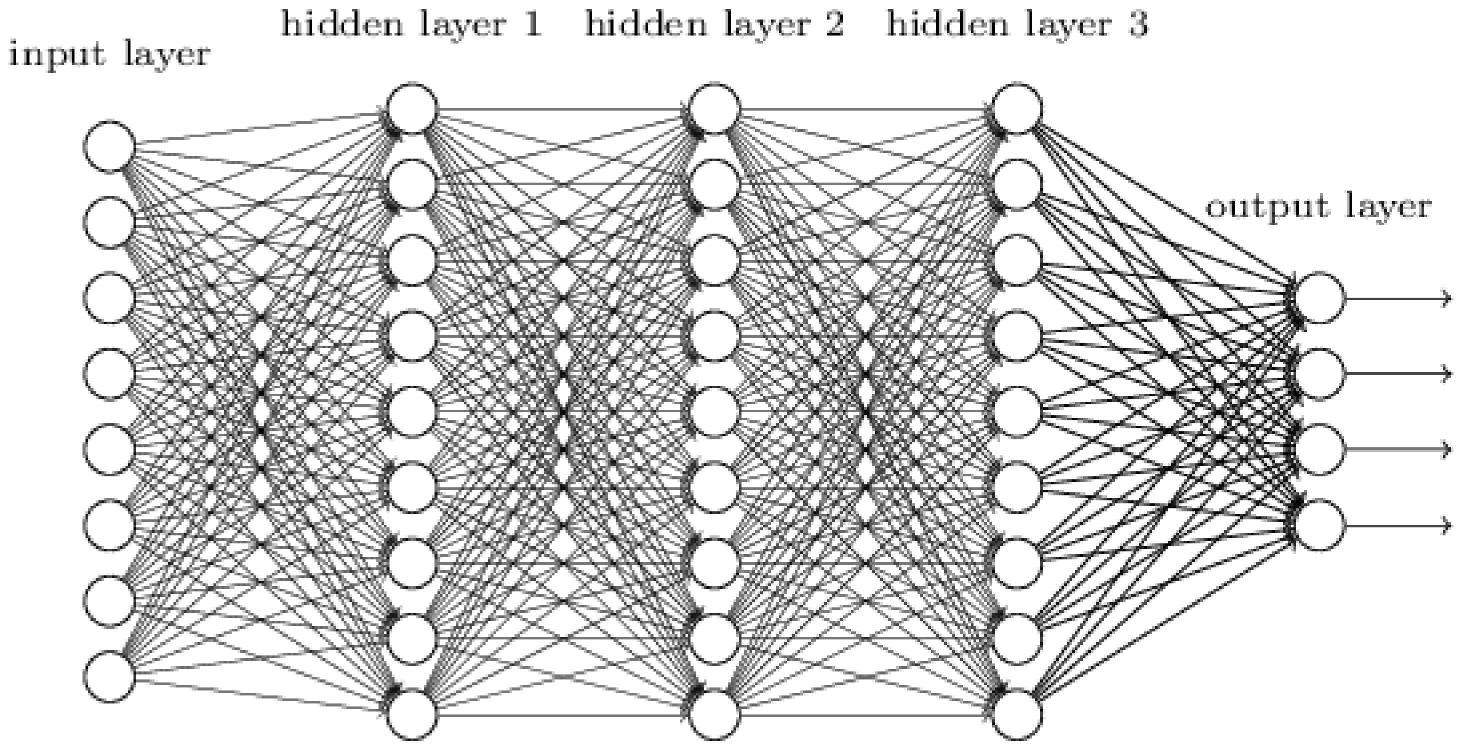

CNN

- Typical application: computer vision
- Intrinsic characteristic: chapter space-related features

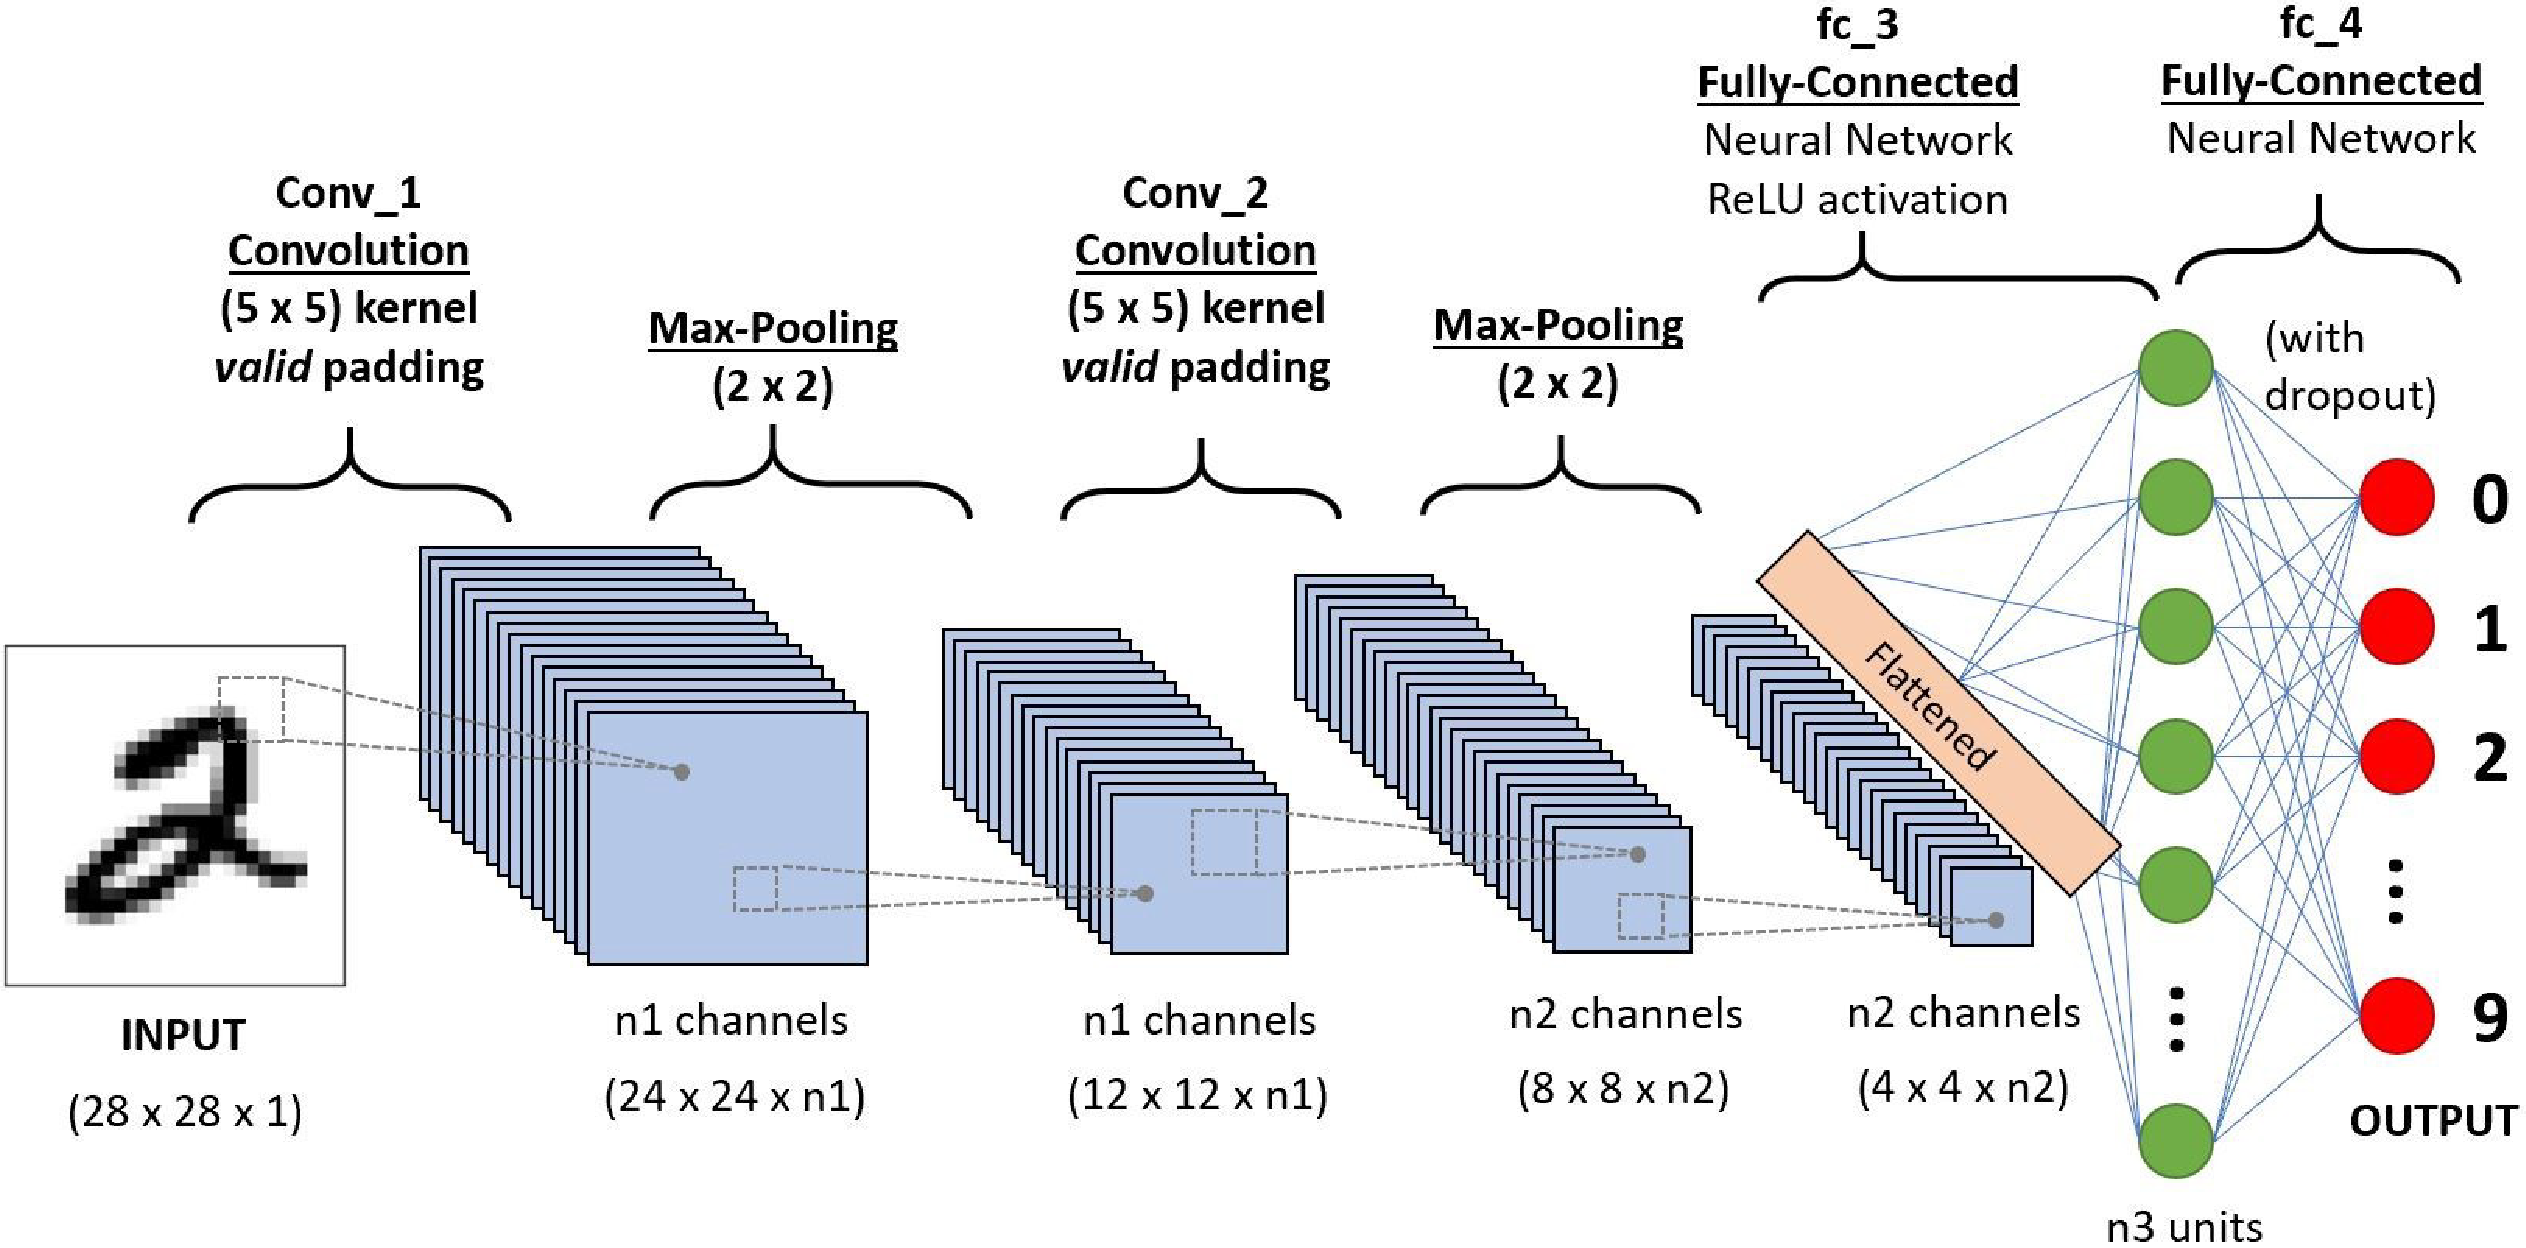
fc_3
Fully-Connected
Neural Network
ReLU activation
fc_4
Fully-Connected
Neural Network
Conv_1
Convolution
(5 x 5) kernel
valid padding
Max-Pooling
(2 x 2)
Conv_2
Convolution
(5 x 5) kernel
valid padding
Max-Pooling
(2 x 2)
(with
dropout)
Flattened
0
1
2
9
INPUT
(28 x 28 x 1)
n1 channels
(24 x 24 x n1)
n1 channels
(12 x 12 x n1)
n2 channels
(8 x 8 x n2)
n2 channels
(4 x 4 x n2)
n3 units
OUTPUT

Convolutional Layers

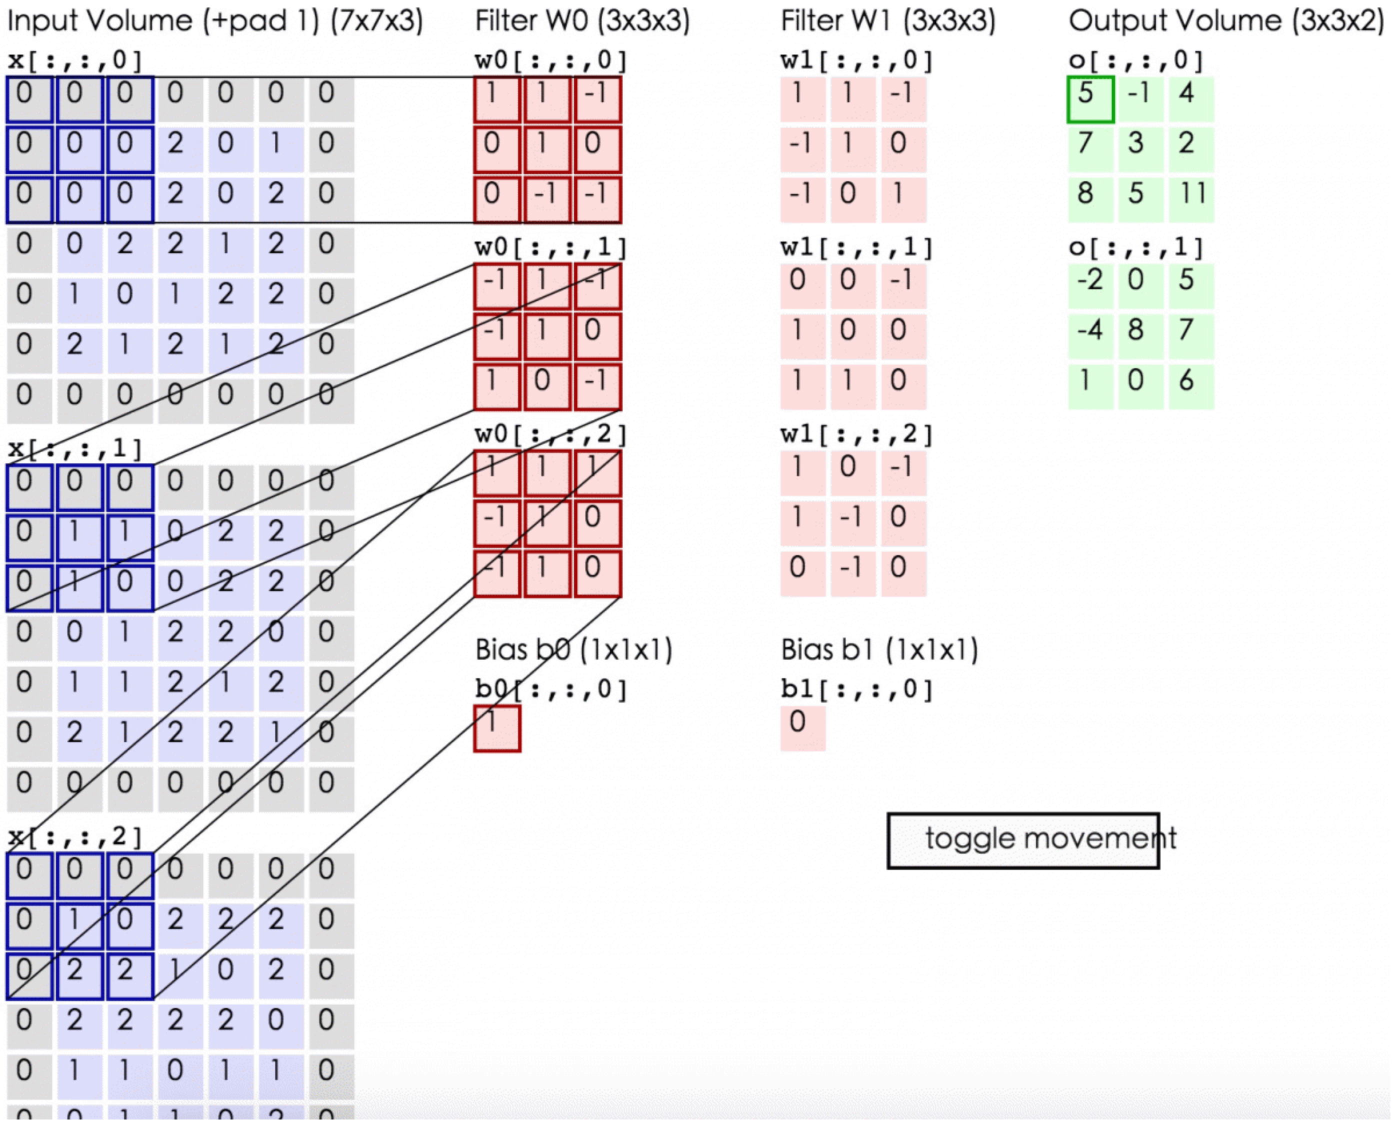

Convolution Filters

-1	-1	-1
1	1	1
0	0	0

-1	1	0
-1	1	0
-1	1	0

0	0	0
1	1	1
-1	-1	-1

0	1	-1
0	1	-1
0	1	-1

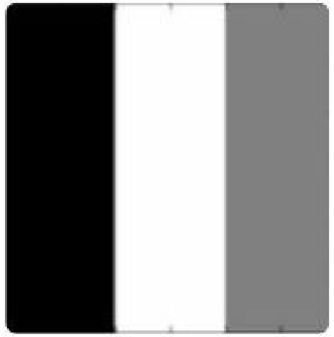

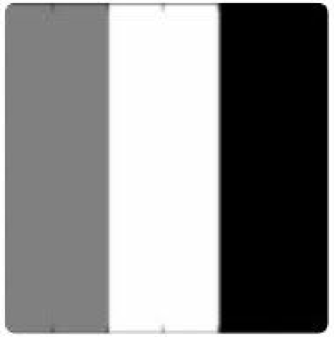

Convolution Filters

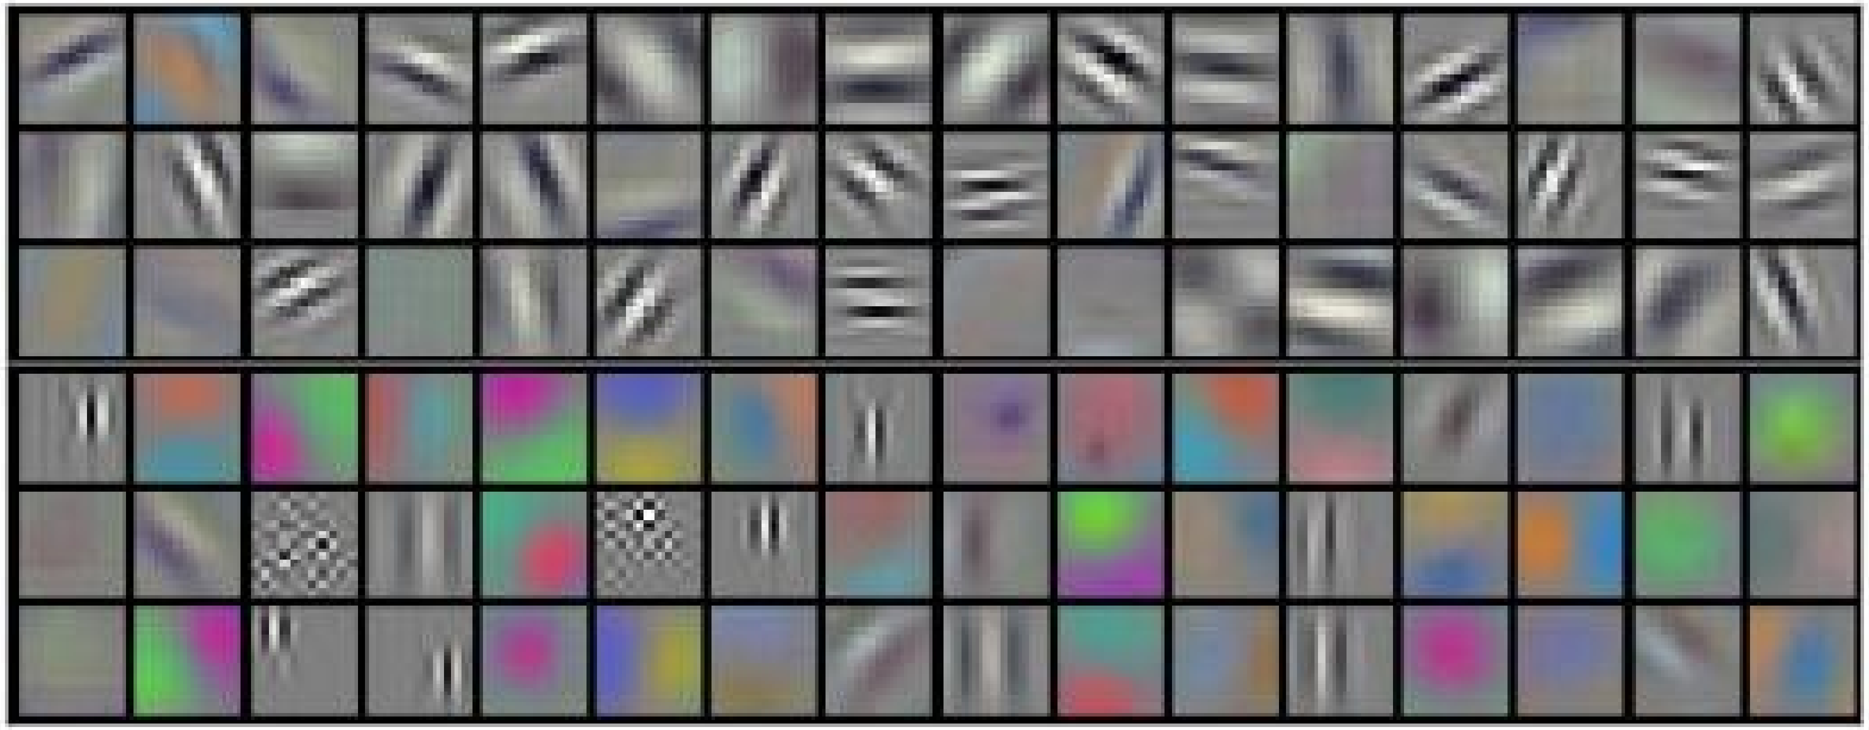

Convolution Filters

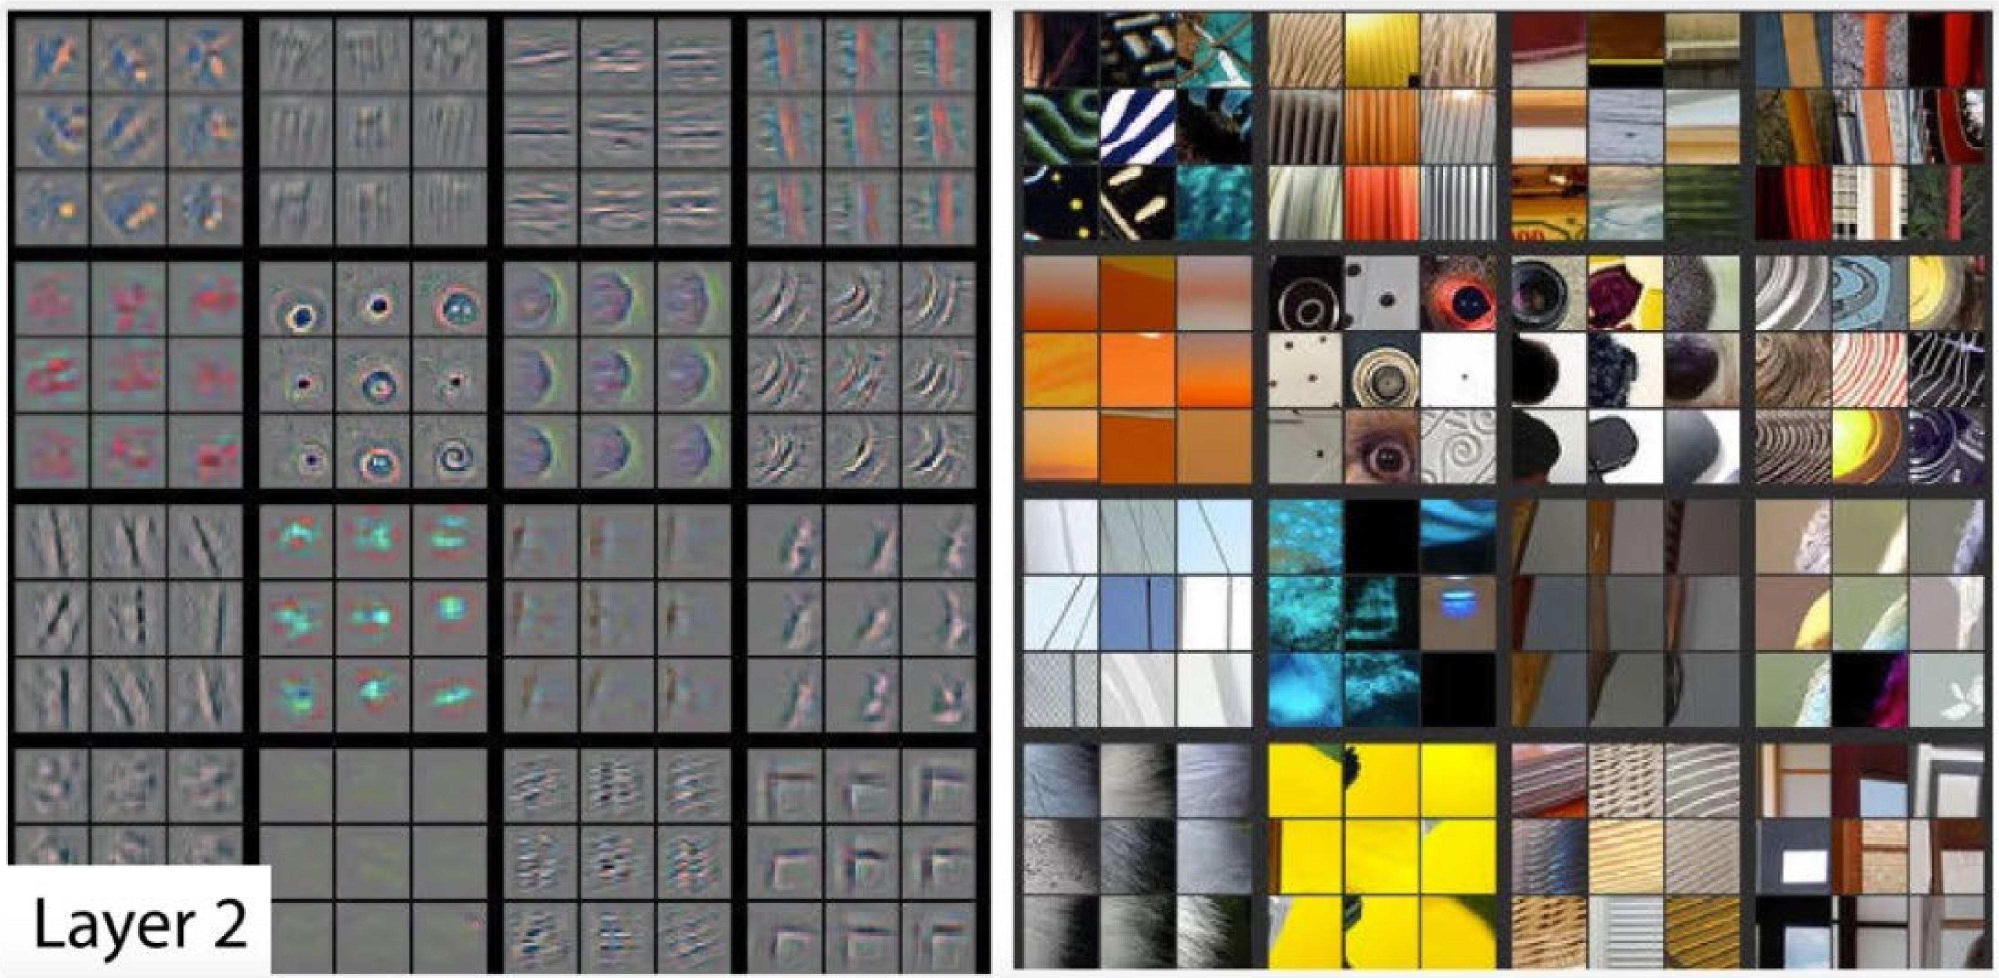

Convolution Filters

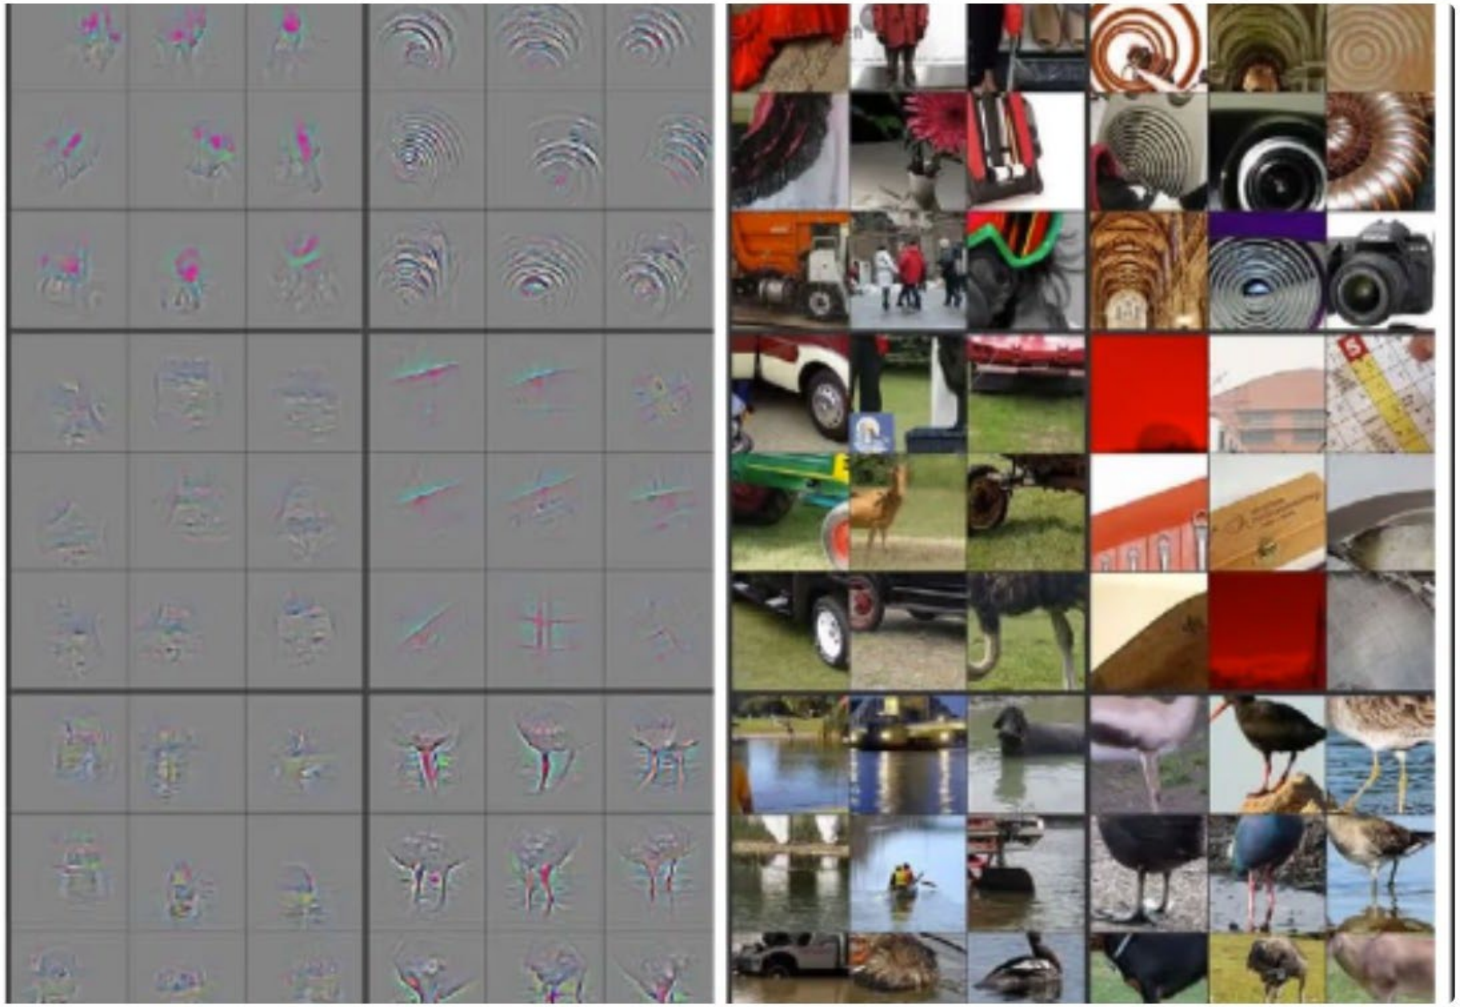

RNN

- Typical application: sequential data, natural language processing
- Intrinsic characteristic: (historical) context-related features

“Vanilla” Neural Network

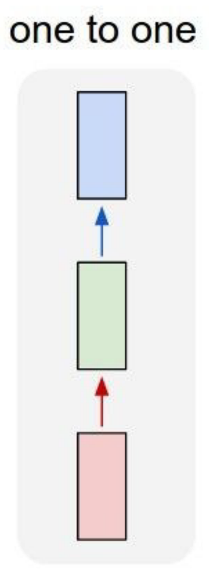

Vanilla Neural Networks

Recurrent Neural Networks: Process Sequences

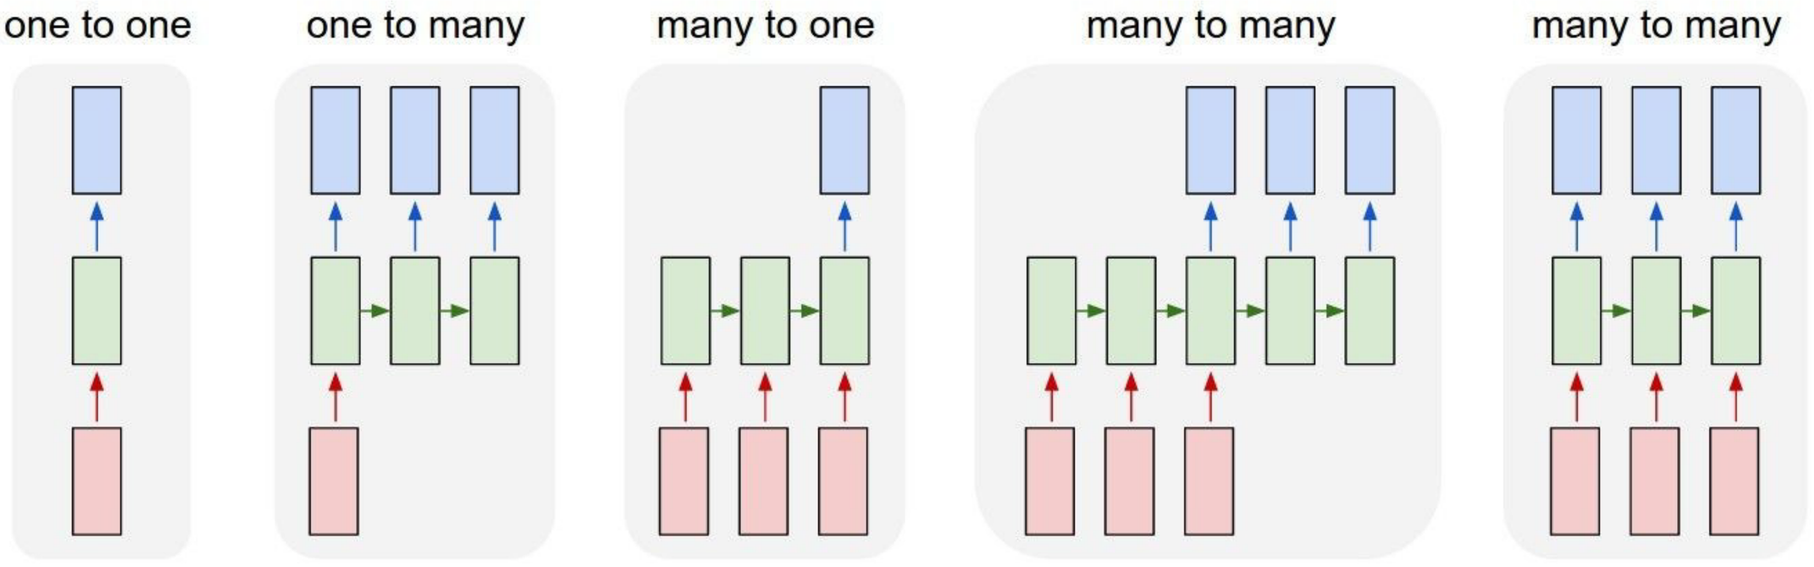

e.g. **Image Captioning**
image -> sequence of words

Recurrent Neural Networks: Process Sequences

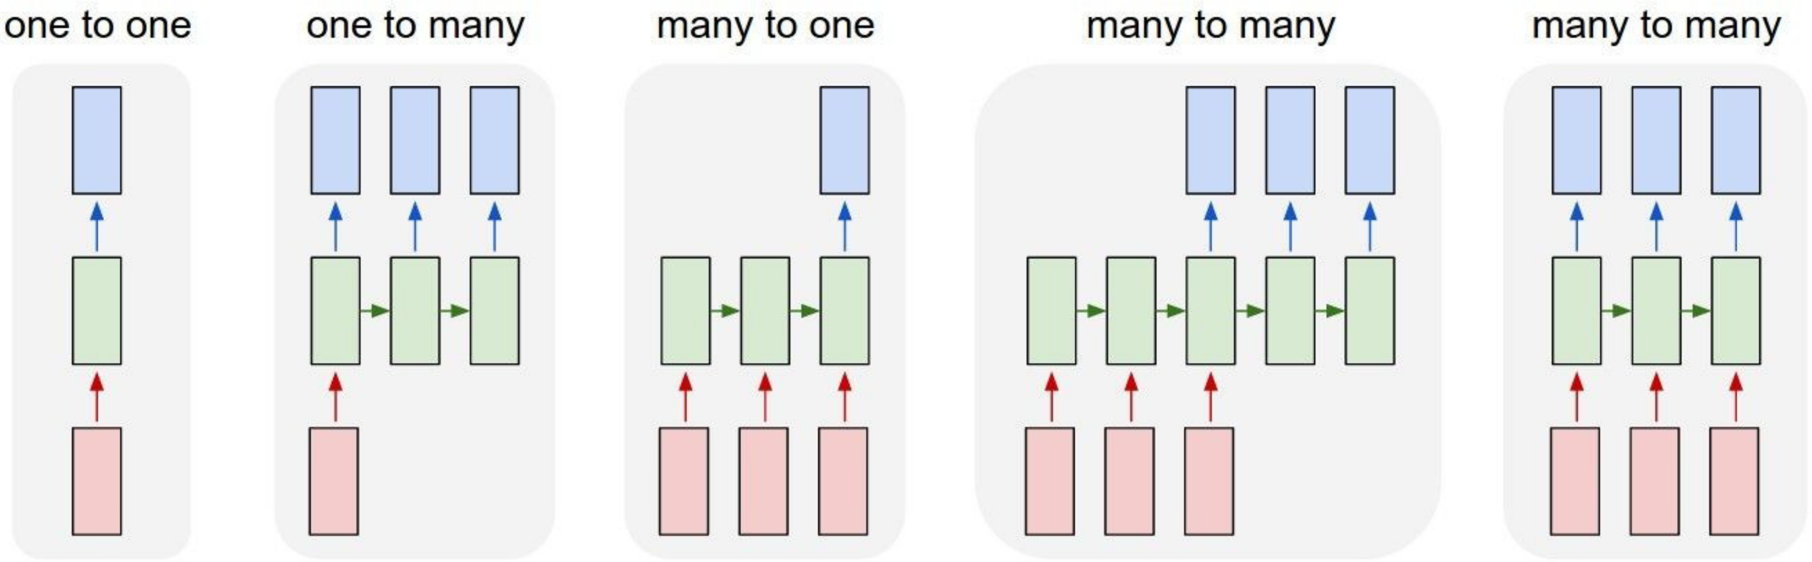

e.g. **Sentiment Classification**
sequence of words -> sentiment

Recurrent Neural Networks: Process Sequences

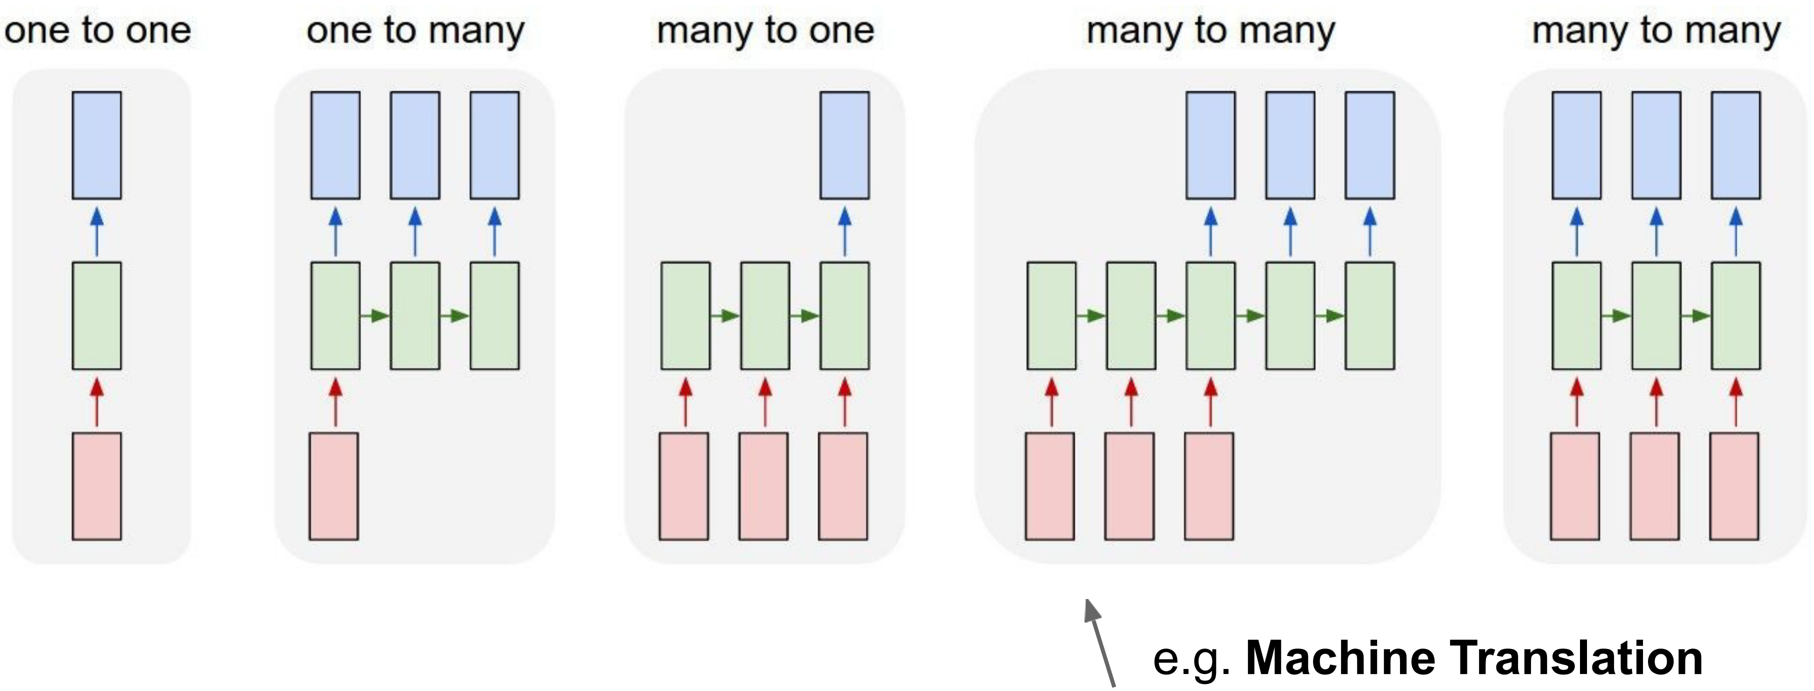

e.g. **Machine Translation**
seq of words -> seq of words

Recurrent Neural Networks: Process Sequences

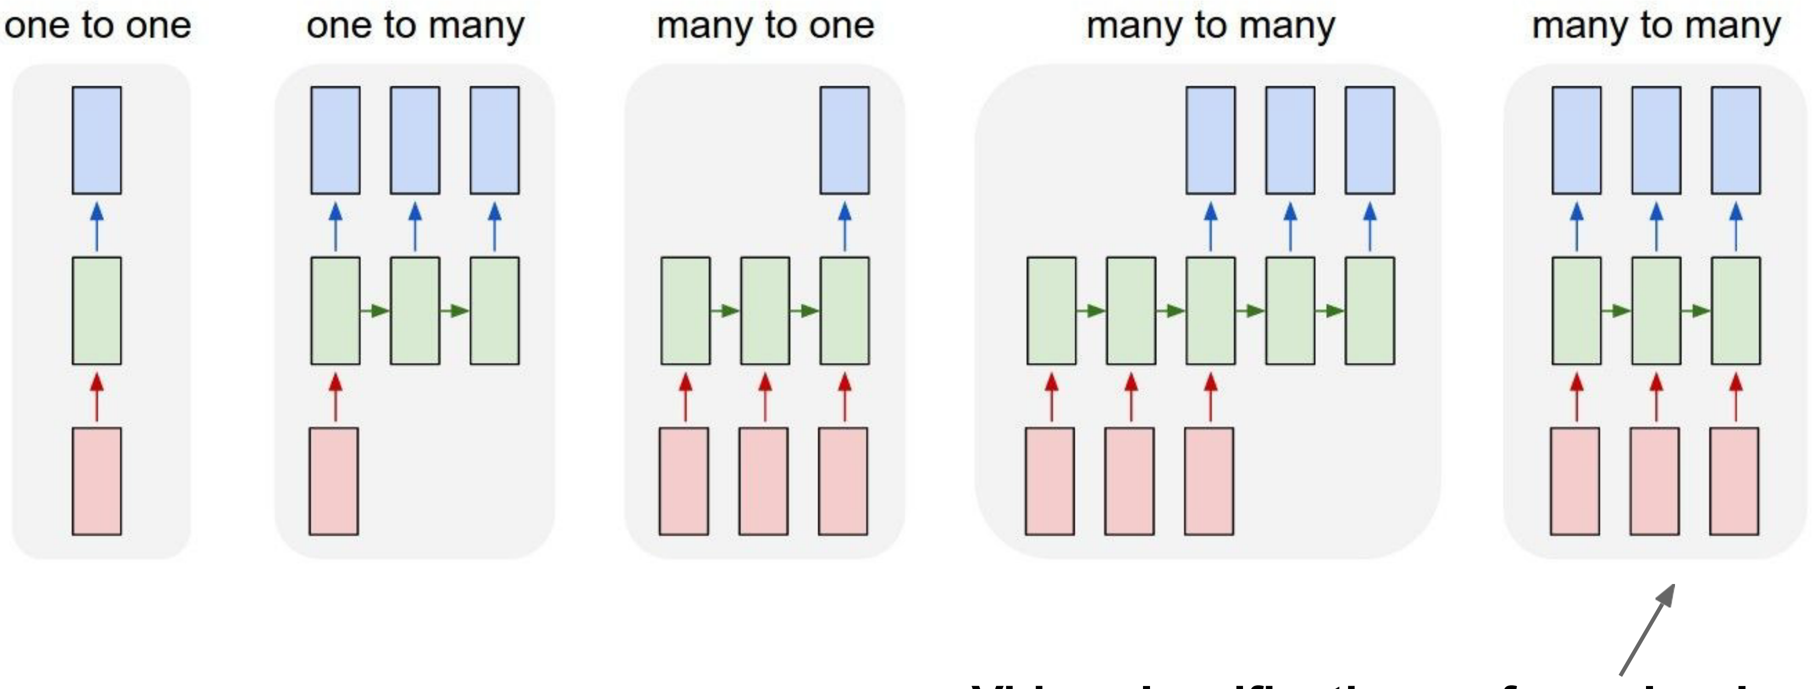

e.g. **Video classification on frame level**

Recurrent Neural Network

We can process a sequence of vectors **x** by applying a **recurrence formula** at every time step:

$$\boxed{h_t} = \boxed{f_W}(\boxed{h_{t-1}}, \boxed{x_t})$$

new state

old state

input vector at some time step

some function with parameters W

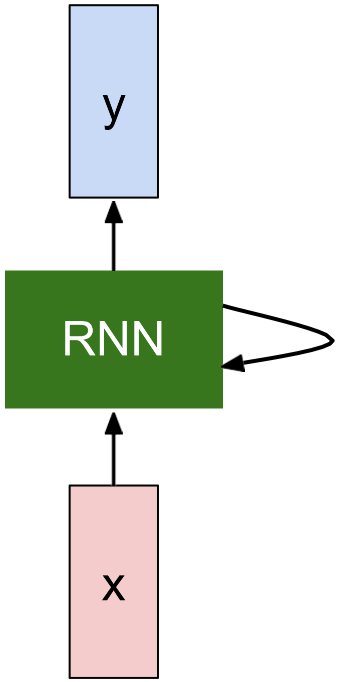

Recurrent Neural Network

We can process a sequence of vectors **x** by applying a **recurrence formula** at every time step:

$$h_t = f_W(h_{t-1}, x_t)$$

Notice: the same function and the same set of parameters are used at every time step.

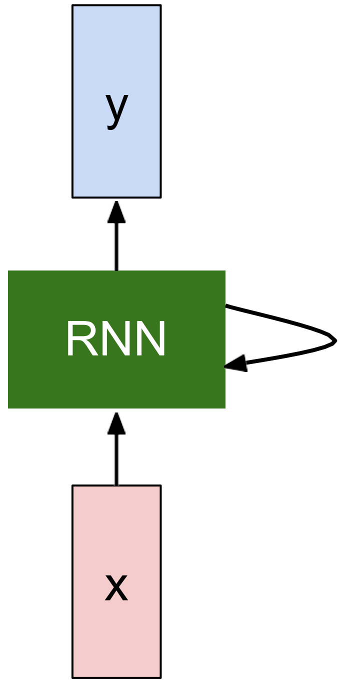

(Vanilla) Recurrent Neural Network

The state consists of a single *"hidden"* vector **h**:

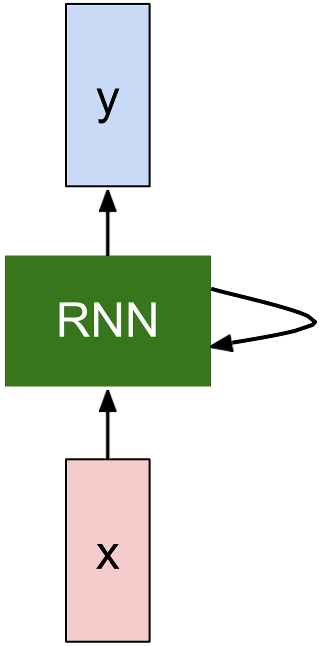

$$h_t = f_W(h_{t-1}, x_t)$$

$$h_t = \tanh(W_{hh}h_{t-1} + W_{xh}x_t)$$

$$y_t = W_{hy}h_t$$

RNN: Computational Graph

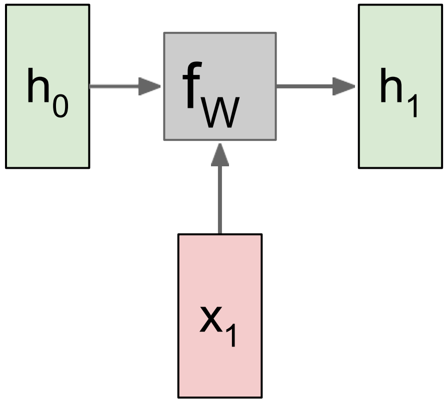

RNN: Computational Graph

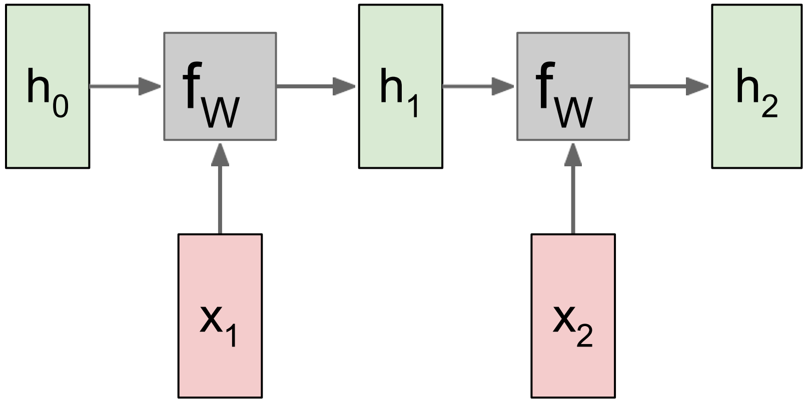

RNN: Computational Graph

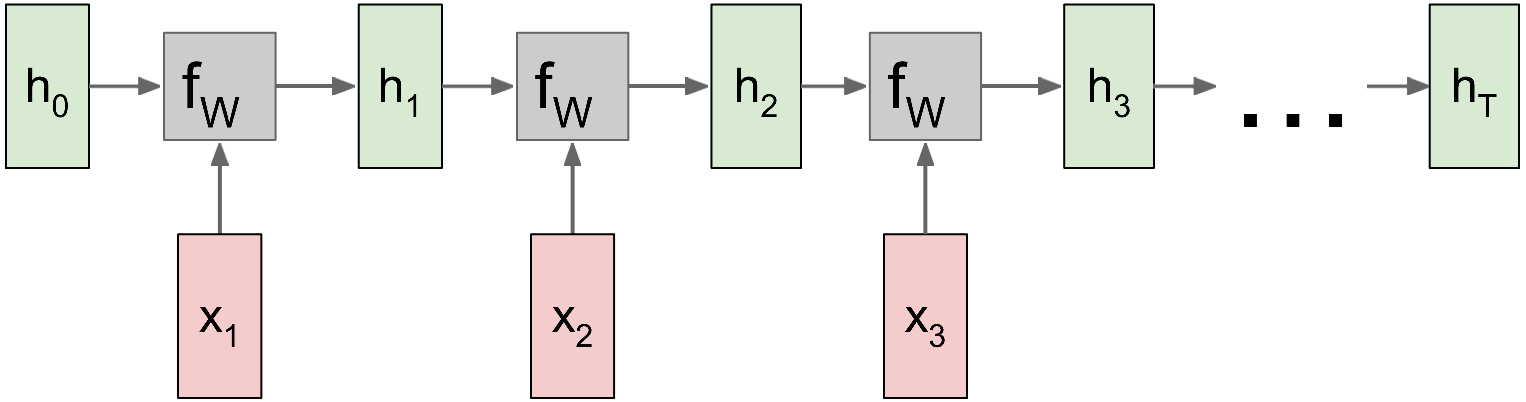

RNN: Computational Graph

Re-use the same weight matrix at every time-step

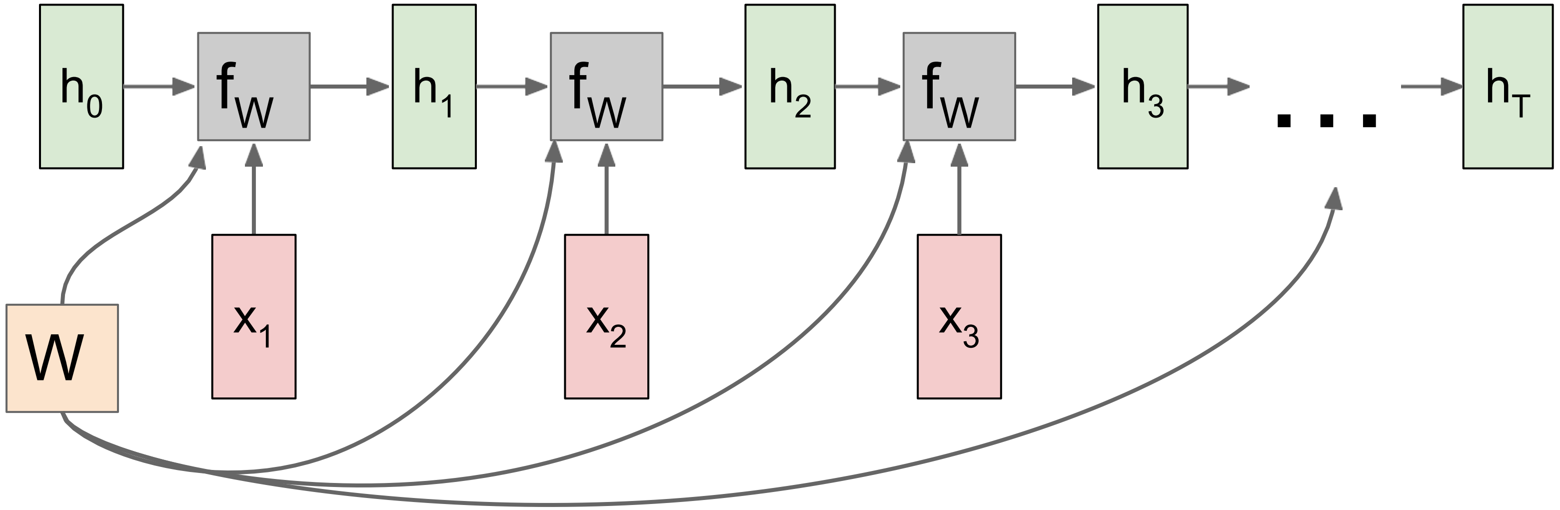

RNN: Computational Graph: Many to Many

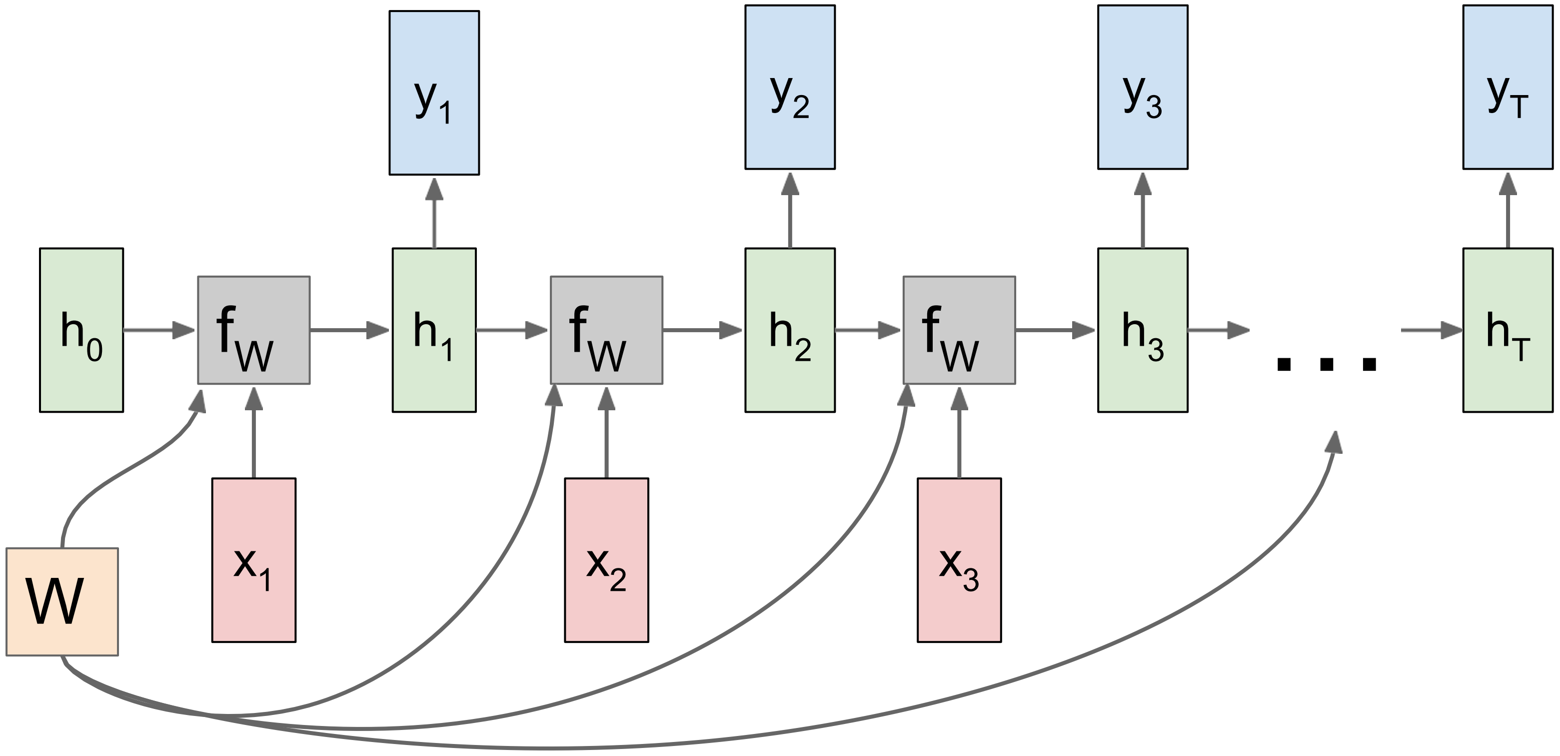

RNN: Computational Graph: Many to Many

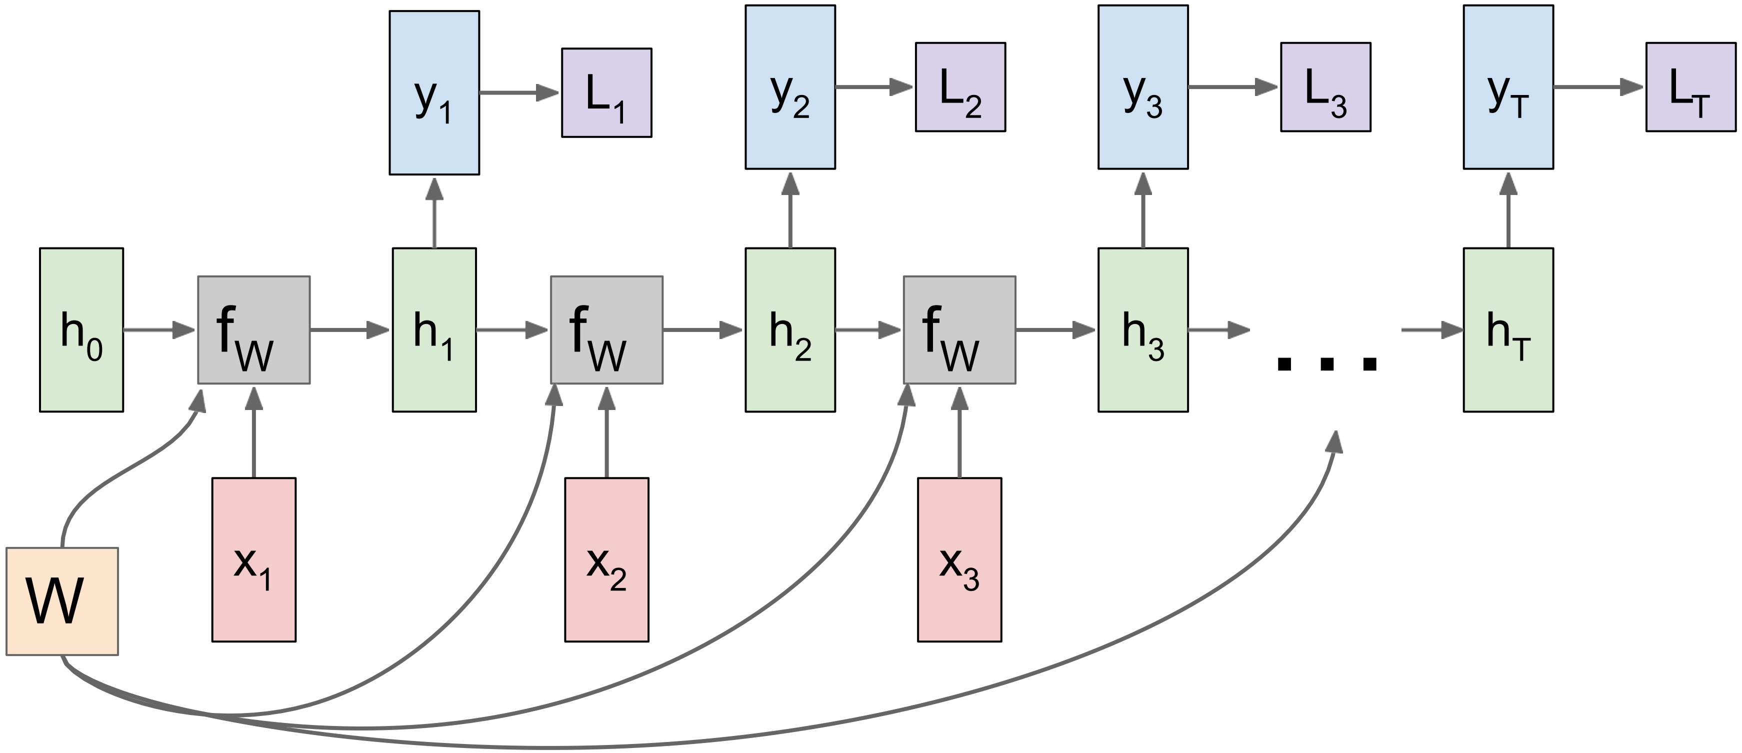

RNN: Computational Graph: Many to Many

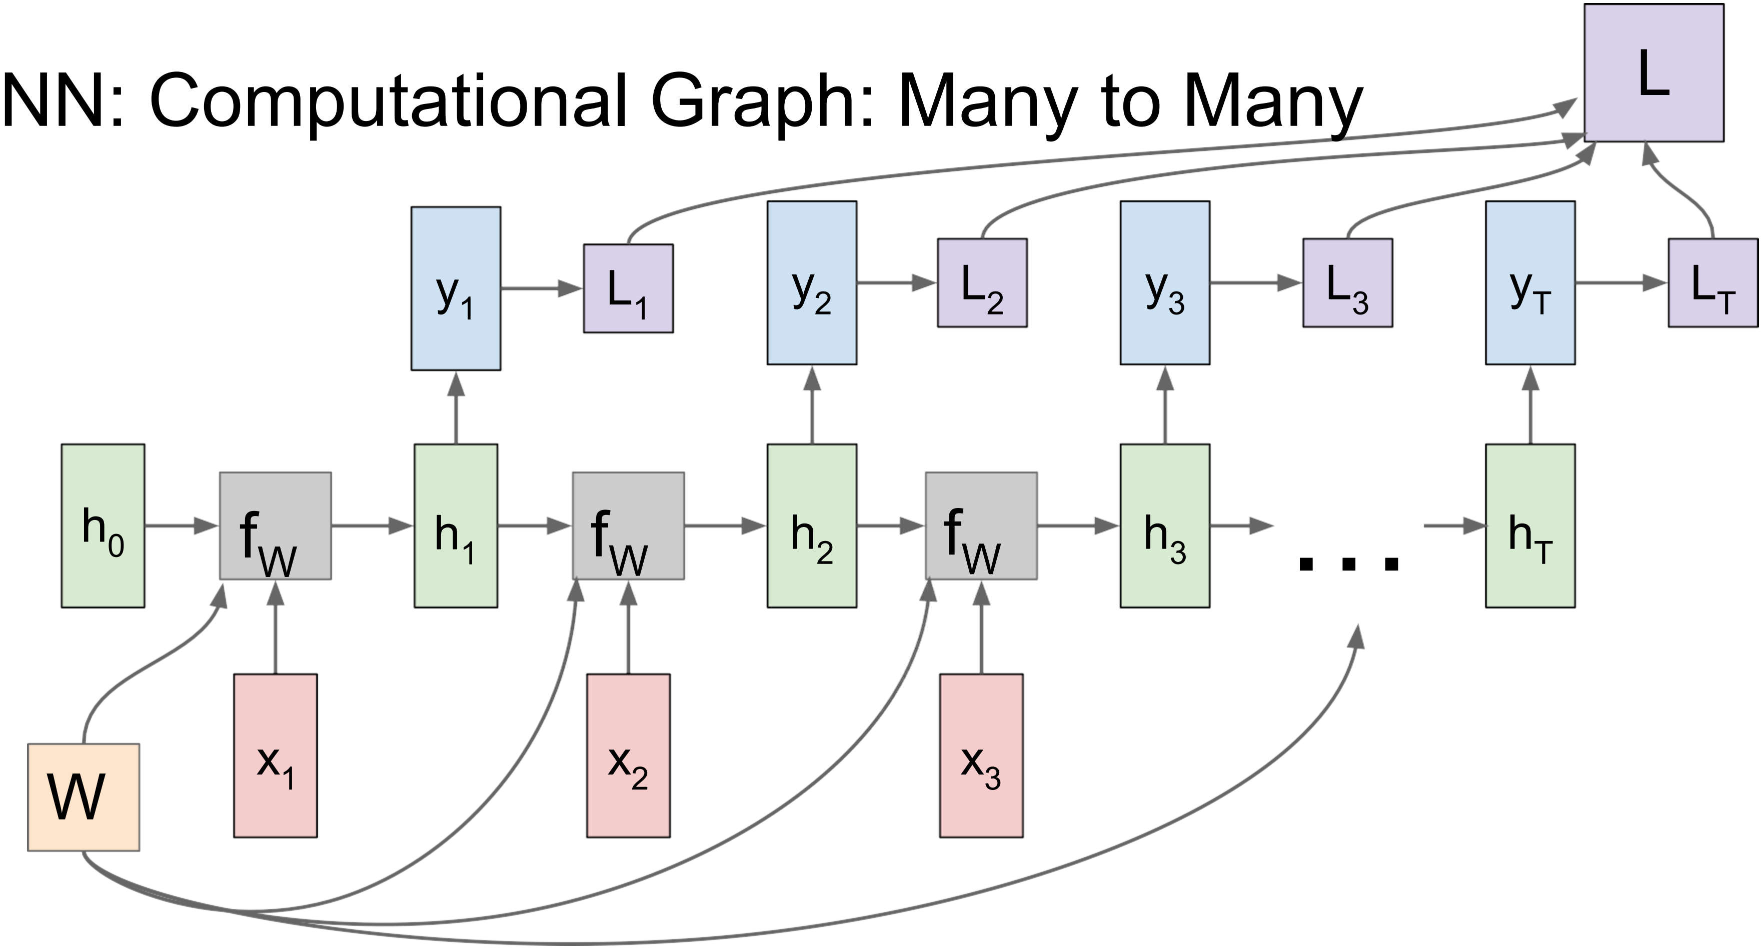

RNN: Computational Graph: Many to One

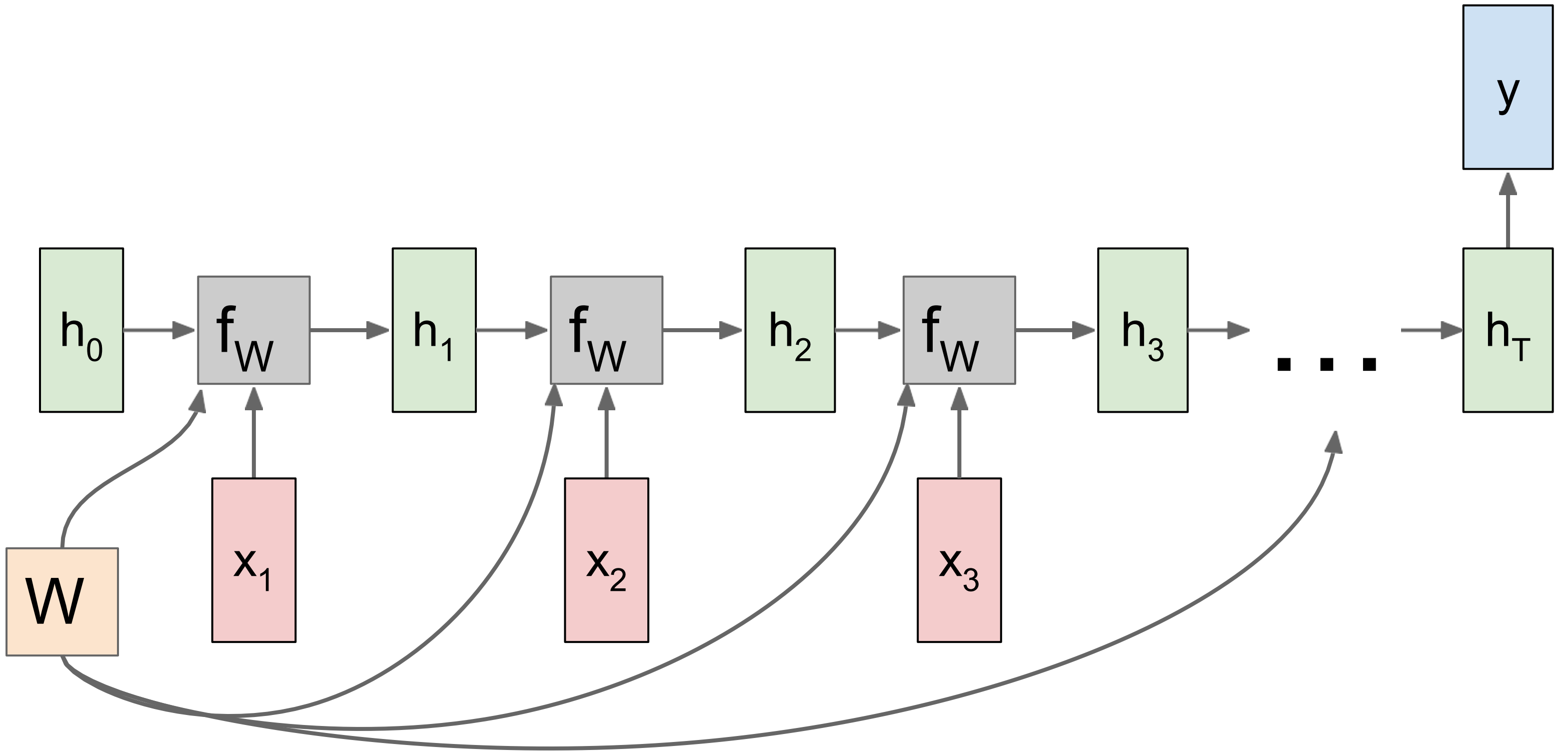

RNN: Computational Graph: One to Many

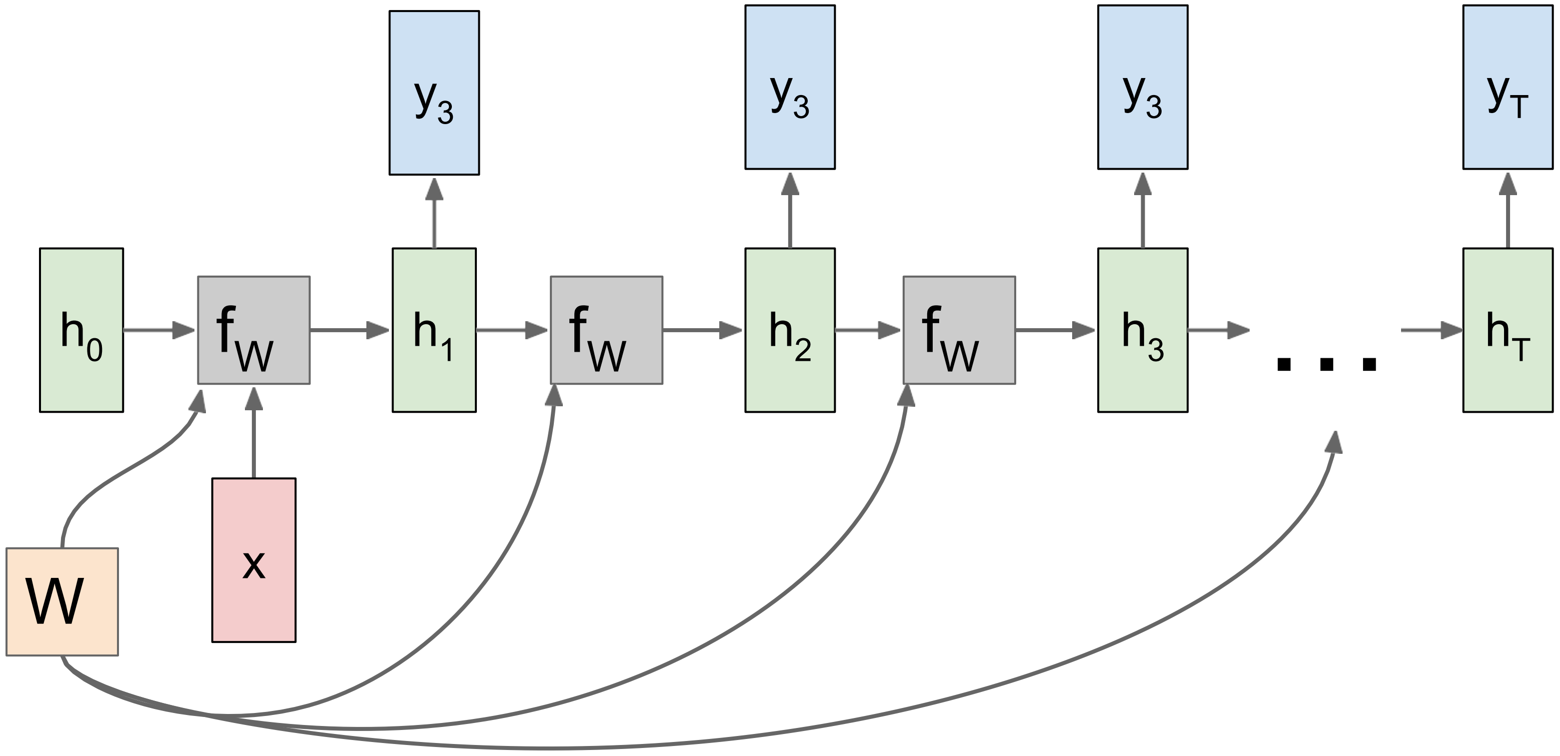

AutoEncoder

- Typical application: embedding, learning hidden representation

Some background first: Autoencoders

Unsupervised approach for learning a lower-dimensional feature representation from unlabeled training data

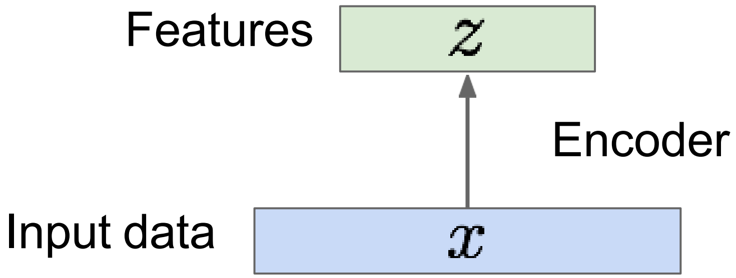

Some background first: Autoencoders

Unsupervised approach for learning a lower-dimensional feature representation from unlabeled training data

Originally: Linear + nonlinearity (sigmoid)
Later: Deep, fully-connected
Later: ReLU CNN

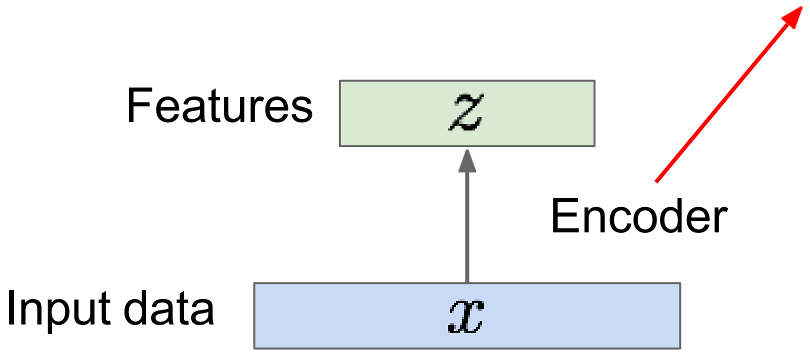

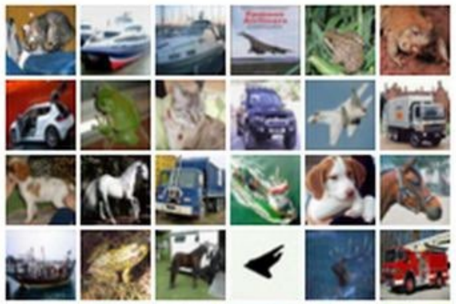

Some background first: Autoencoders

Unsupervised approach for learning a lower-dimensional feature representation from unlabeled training data

z usually smaller than **x** (dimensionality reduction)

Q: Why dimensionality reduction?

Originally: Linear + nonlinearity (sigmoid)
Later: Deep, fully-connected
Later: ReLU CNN

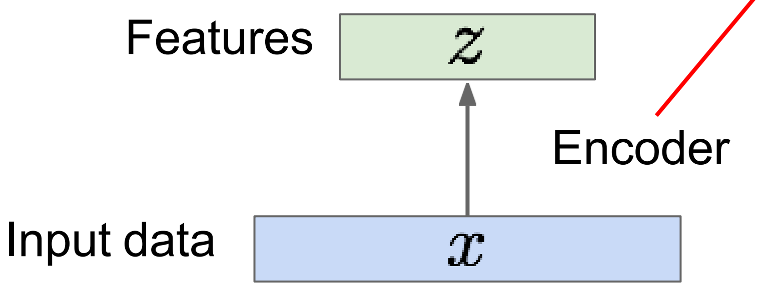

Some background first: Autoencoders

Unsupervised approach for learning a lower-dimensional feature representation from unlabeled training data

z usually smaller than **x** (dimensionality reduction)

Q: Why dimensionality reduction?

A: Want features to capture meaningful factors of variation in data

Originally: Linear + nonlinearity (sigmoid)
Later: Deep, fully-connected
Later: ReLU CNN

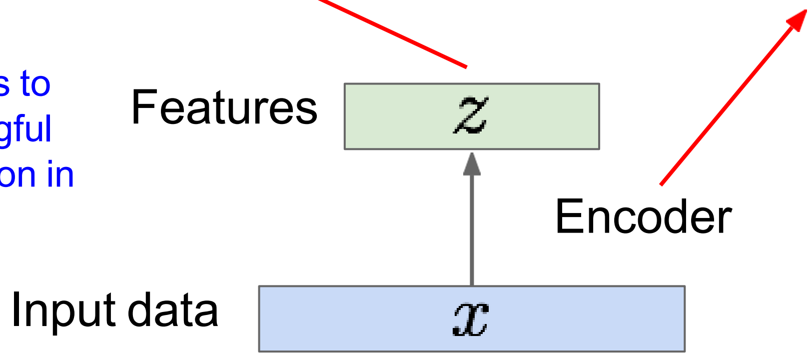

Some background first: Autoencoders

How to learn this feature representation?

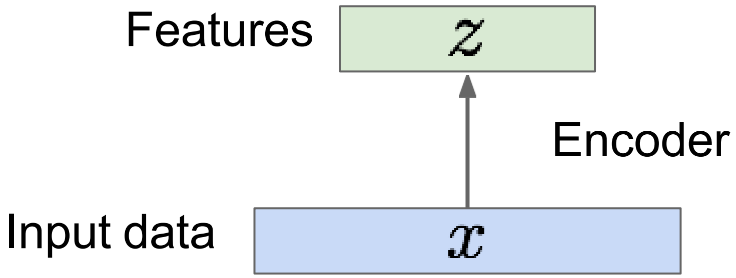

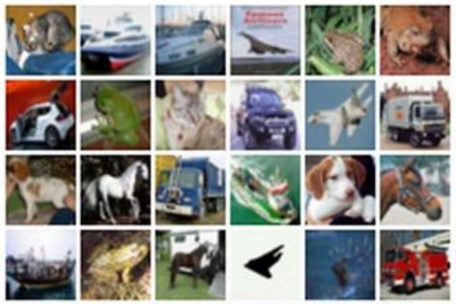

Some background first: Autoencoders

How to learn this feature representation?
Train such that features can be used to reconstruct original data
“Autoencoding” - encoding itself

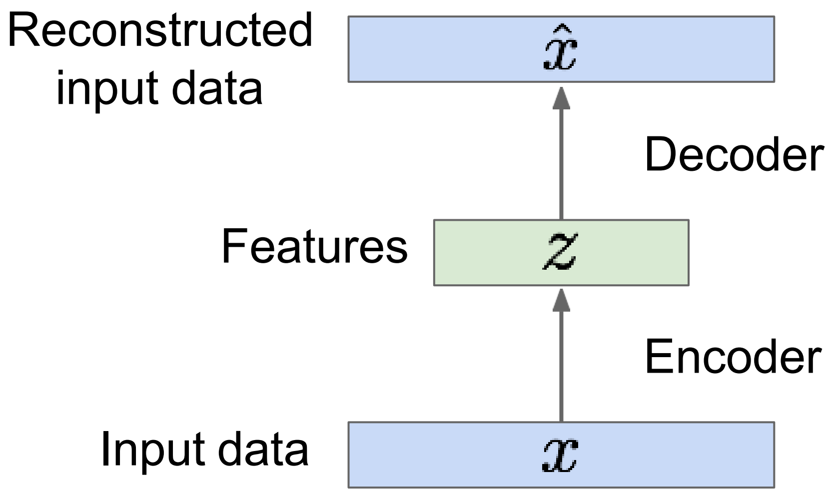

Some background first: Autoencoders

How to learn this feature representation?
Train such that features can be used to reconstruct original data
“Autoencoding” - encoding itself

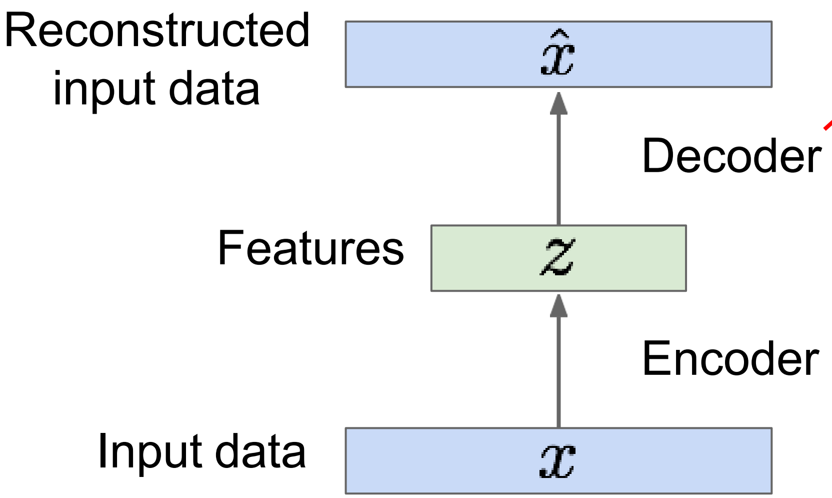

Originally: Linear + nonlinearity (sigmoid)
Later: Deep, fully-connected
Later: ReLU CNN (upconv)

Some background first: Autoencoders

How to learn this feature representation?
Train such that features can be used to reconstruct original data
"Autoencoding" - encoding itself

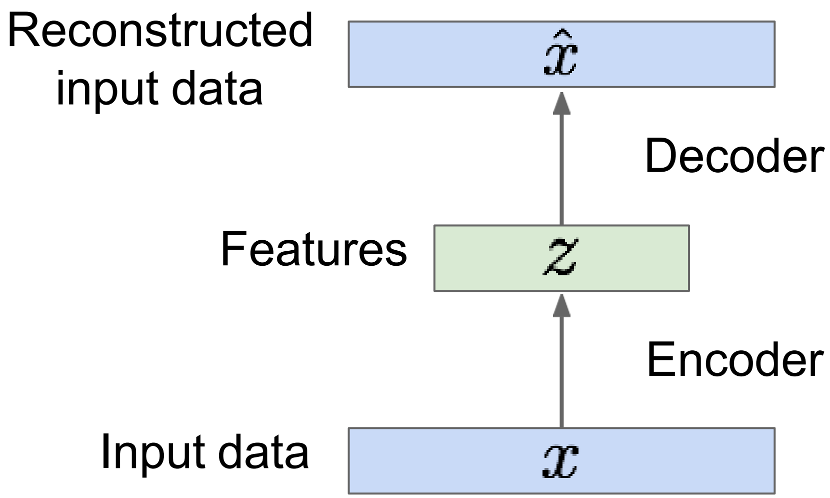

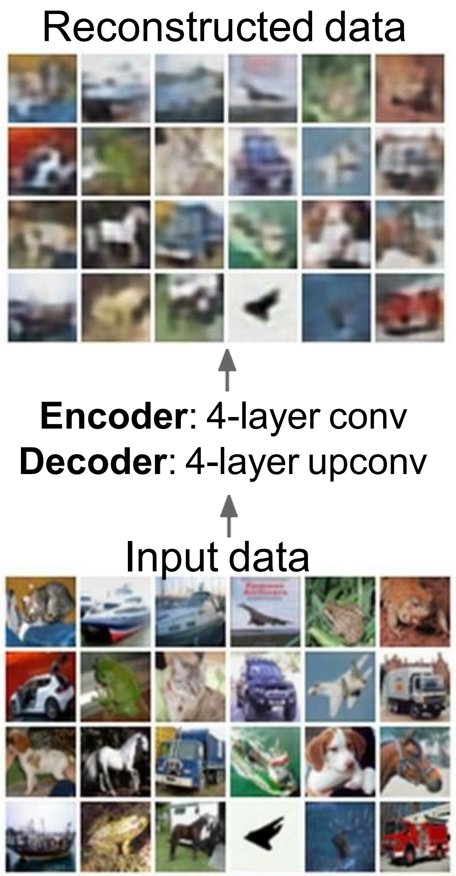

Some background first: Autoencoders

Train such that features can be used to reconstruct original data

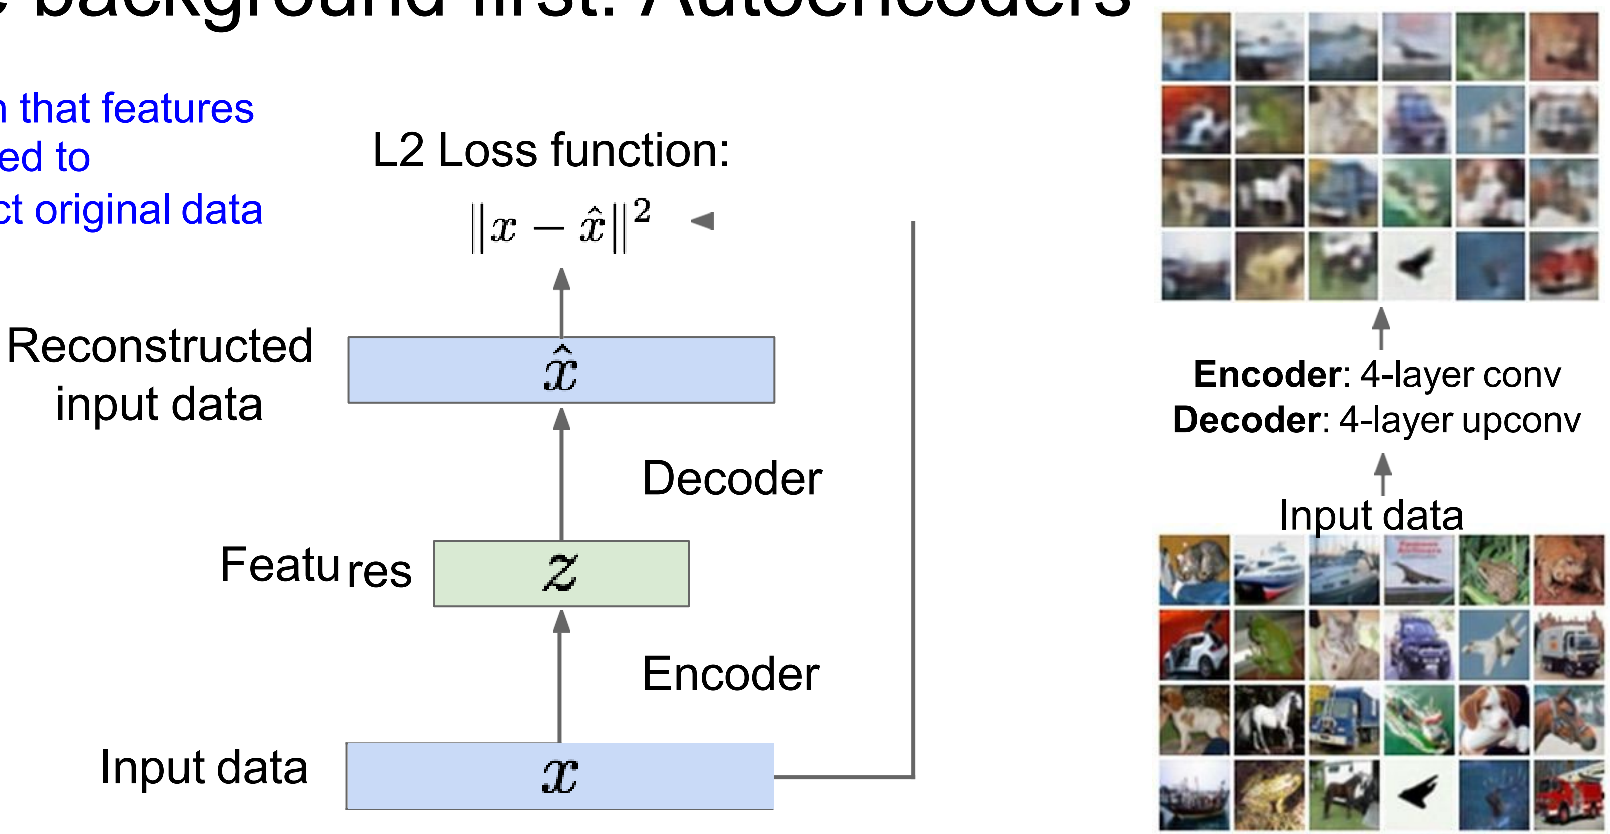

Some background first: Autoencoders

Train such that features can be used to reconstruct original data

Doesn't use labels!

L2 Loss function:

$$\|x - \hat{x}\|^2$$

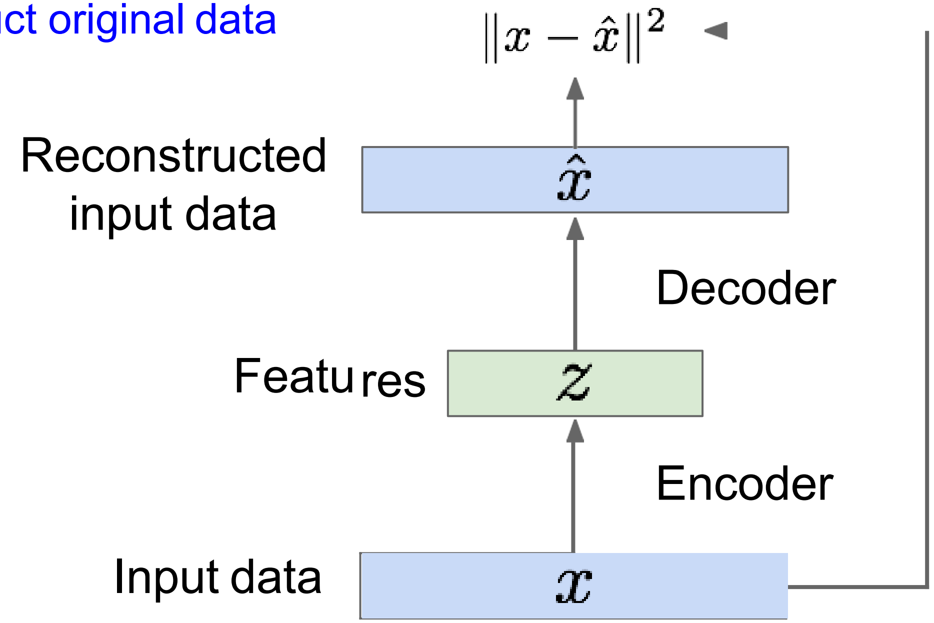

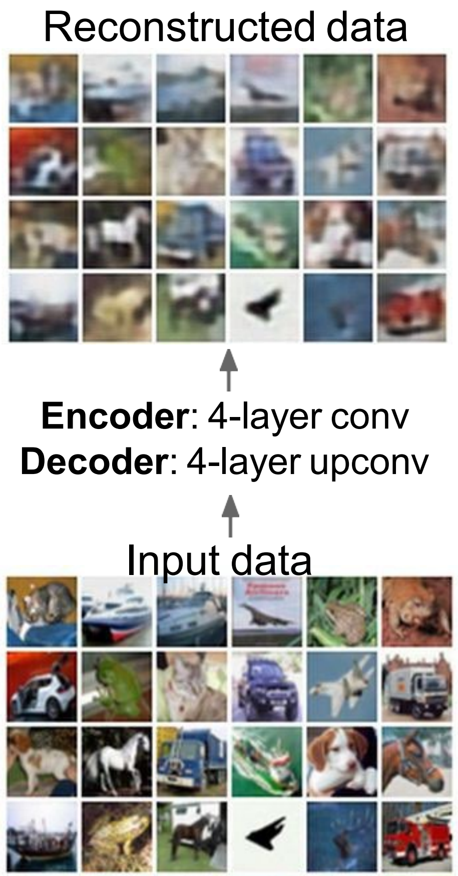

Some background first: Autoencoders

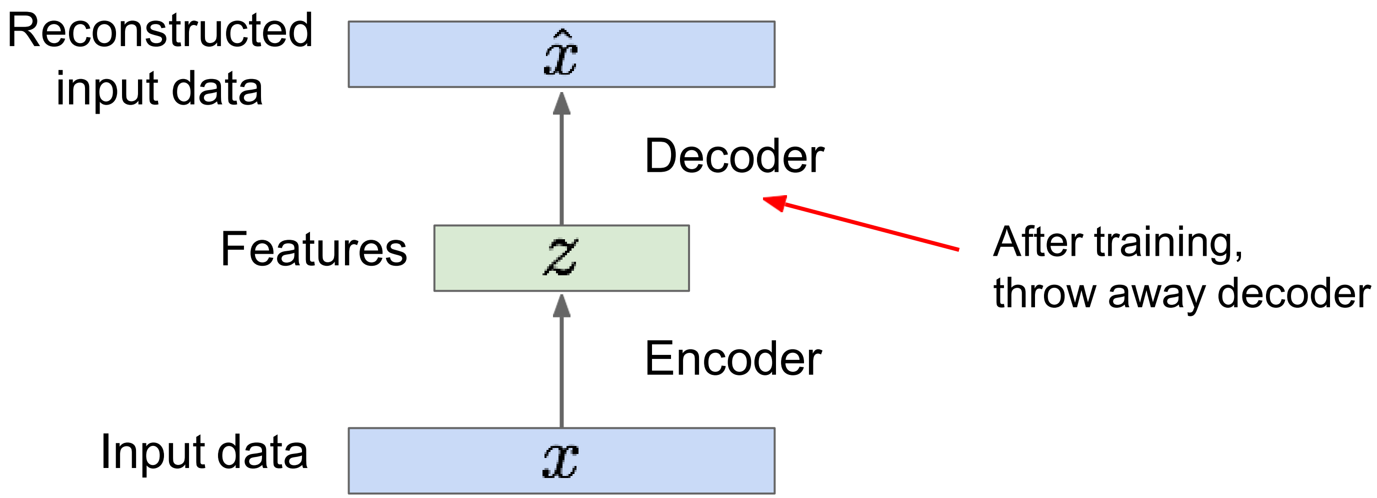

Some background first: Autoencoders

Encoder can be used to initialize a **supervised** model

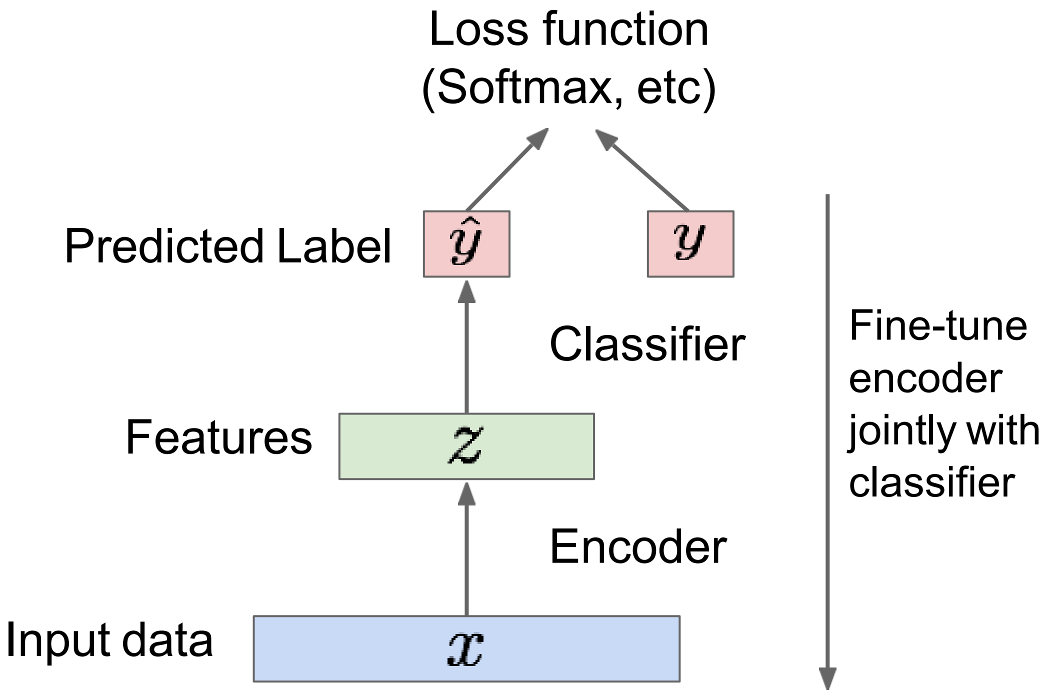

bird plane

dog deer truck

Train for final task (sometimes with small data)

Generative networks

- Typical application: Realistic samples for artwork
  - Actually it depends on your creativity

Generative Models

Given training data, generate new samples from same distribution

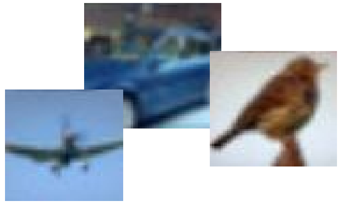

Training data ~ $p_{data}(x)$

Generated samples ~ $p_{model}(x)$

Want to learn $p_{model}(x)$ similar to $p_{data}(x)$

Generative Models

Given training data, generate new samples from same distribution

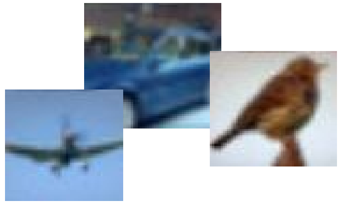

Training data ~ $p_{data}(x)$

Generated samples ~ $p_{model}(x)$

Want to learn $p_{model}(x)$ similar to $p_{data}(x)$

Addresses density estimation, a core problem in unsupervised learning

Several flavors:

- Explicit density estimation: explicitly define and solve for $p_{model}(x)$
- Implicit density estimation: learn model that can sample from $p_{model}(x)$ w/o explicitly defining it

Why Generative Models?

- Realistic samples for artwork, super-resolution, colorization, etc.

- Generative models of time-series data can be used for simulation and planning (reinforcement learning applications!)
- Training generative models can also enable inference of latent representations that can be useful as general features

Taxonomy of Generative Models

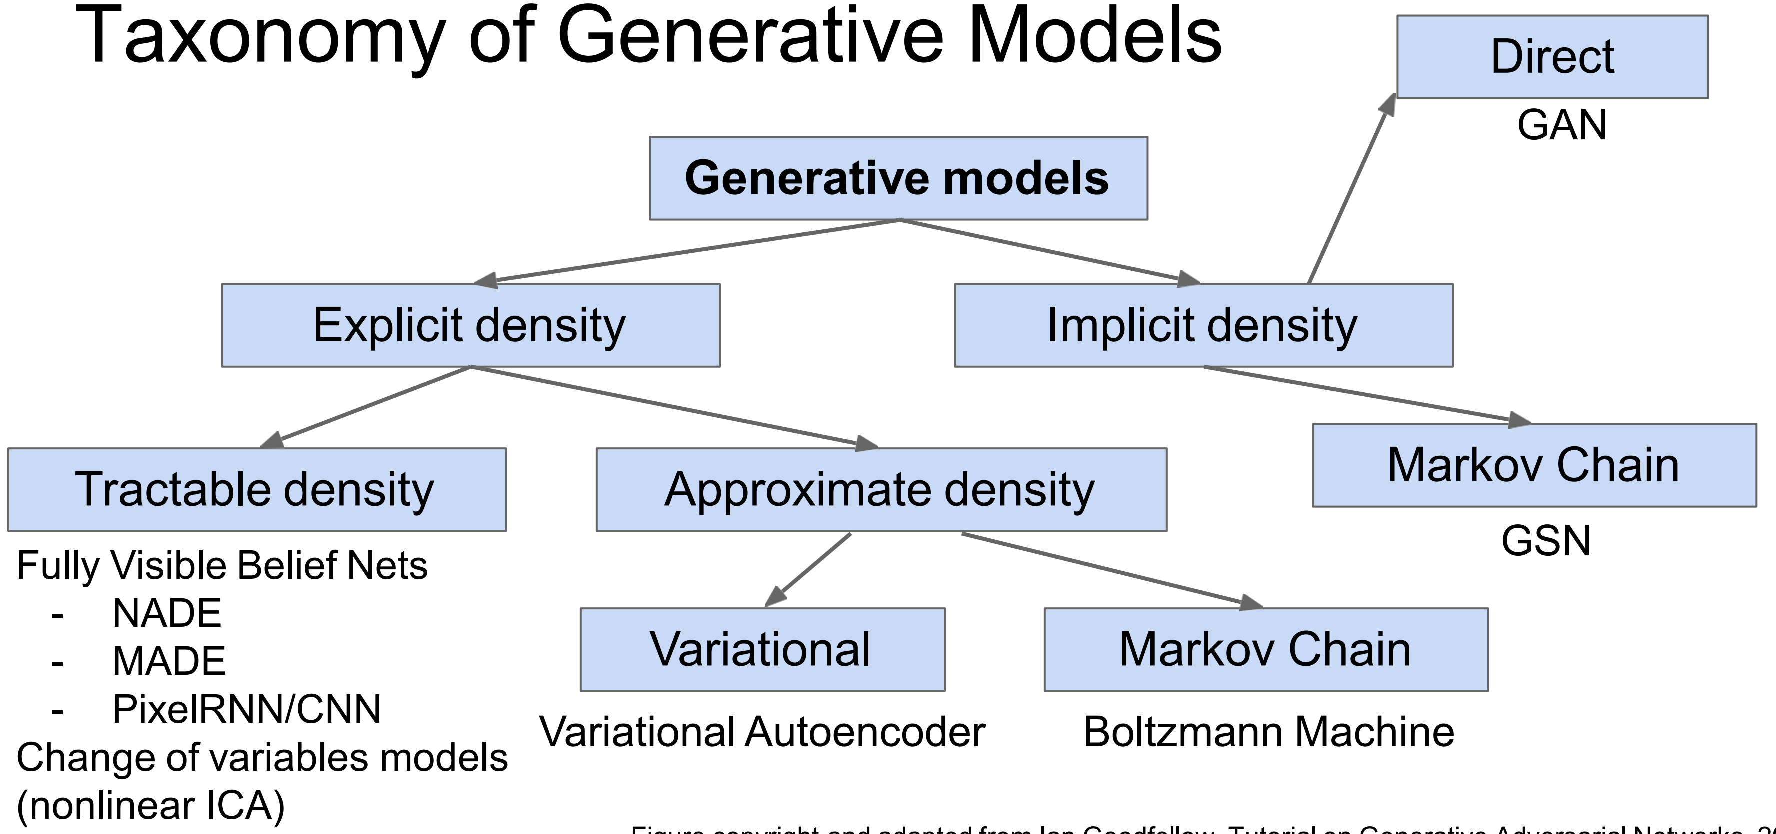

Generative Adversarial Networks

Ian Goodfellow et al., “Generative Adversarial Nets”, NIPS 2014

Problem: Want to sample from complex, high-dimensional training distribution. No direct way to do this!

Solution: Sample from a simple distribution, e.g. random noise. Learn transformation to training distribution.

Q: What can we use to represent this complex transformation?

A: A neural network!

Output: Sample from training distribution

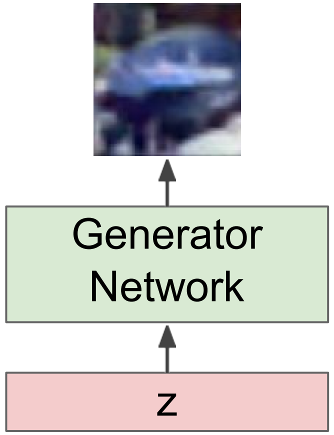

Input: Random noise

Training GANs: Two-player game

Ian Goodfellow et al., "Generative Adversarial Nets", NIPS 2014

Generator network: try to fool the discriminator by generating real-looking images
Discriminator network: try to distinguish between real and fake images

Training GANs: Two-player game

Ian Goodfellow et al., “Generative Adversarial Nets”, NIPS 2014

Generator network: try to fool the discriminator by generating real-looking images
Discriminator network: try to distinguish between real and fake images

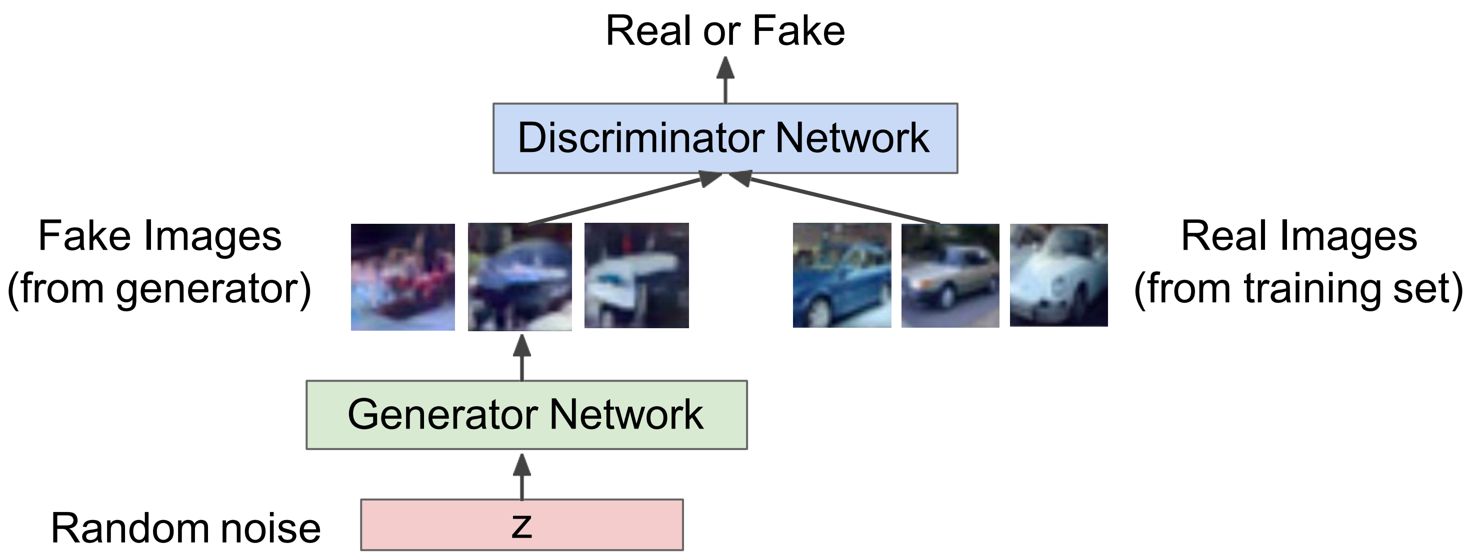

Training GANs: Two-player game

Ian Goodfellow et al., “Generative Adversarial Nets”, NIPS 2014

Generator network: try to fool the discriminator by generating real-looking images
Discriminator network: try to distinguish between real and fake images

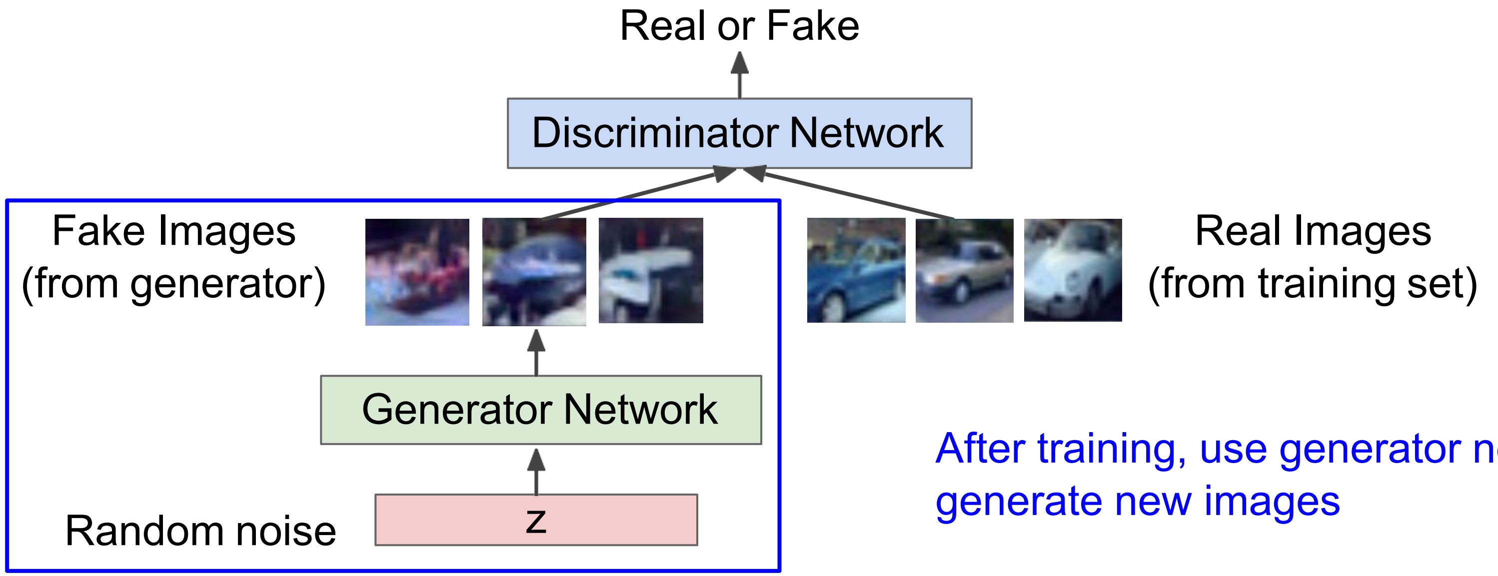

After training, use generator network to generate new images

Ian Goodfellow et al., “Generative Adversarial Nets”, NIPS 2014

Generative Adversarial Nets

Generated samples

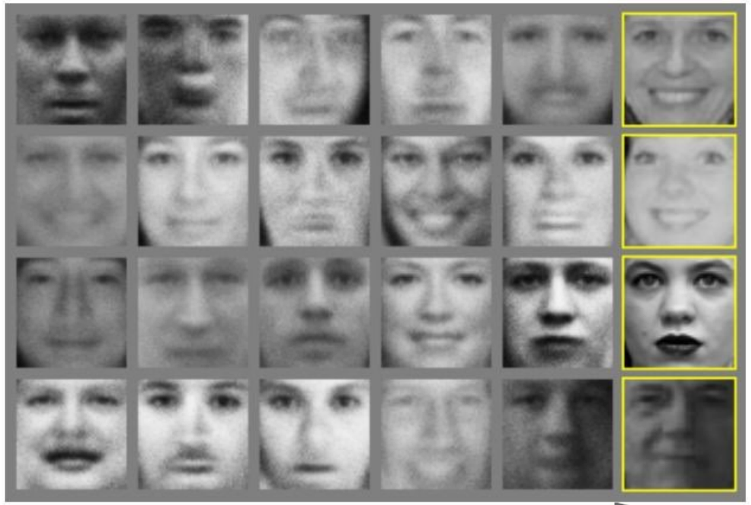

Nearest neighbor from training set

2017: Year of the GAN

Better training and generation

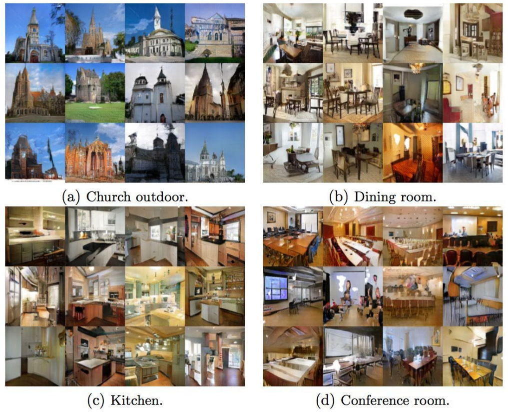

LSGAN. Mao et al. 2017.

BEGAN. Bertholet et al. 2017.

Source->Target domain transfer

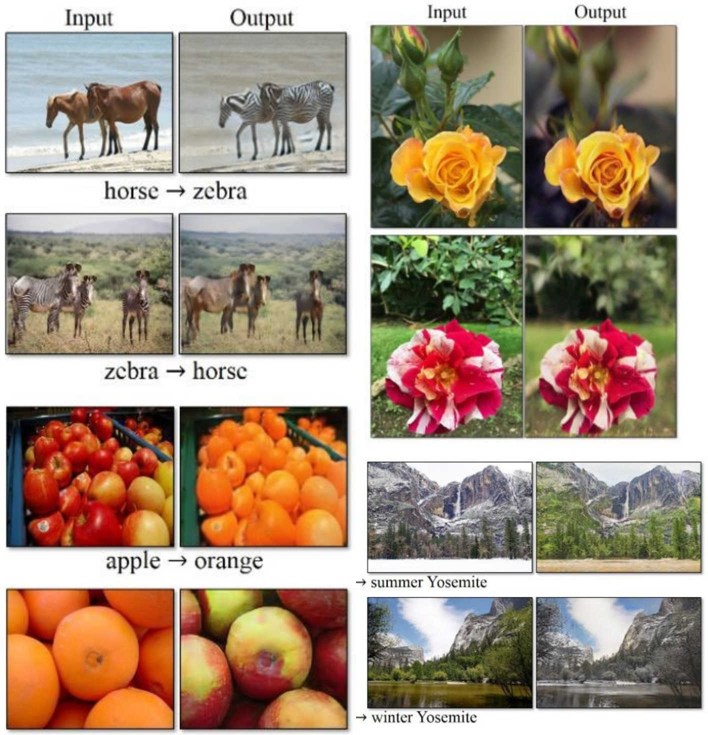

CycleGAN. Zhu et al. 2017.

Text -> Image Synthesis

Reed et al. 2017.

Many GAN applications

Pix2pix. Isola 2017. Many examples at https://phillipi.github.io/pix2pix/