Applications of Deep Learning

Jeff Calder

Math 5467: Intro to Math of Image and Data Analysis
University of Minnesota

Deep Learning

Deep Learning is a field within machine learning based on artificial neural networks with many layers.

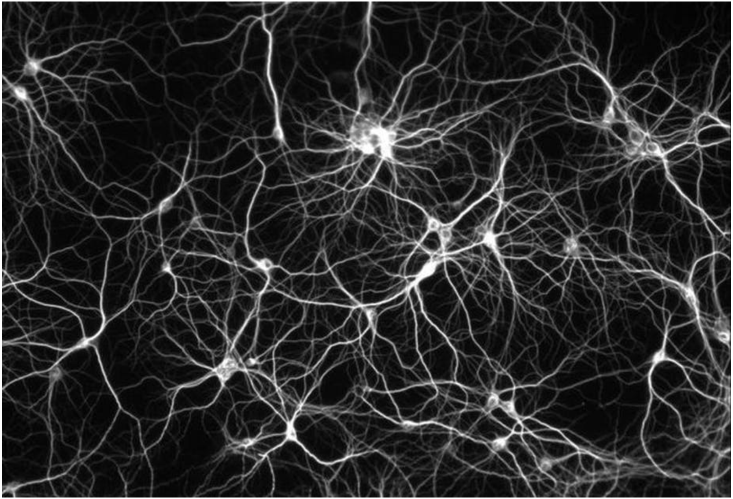

History of deep learning

Pre 2012 history:

- The original ideas date back to the 1940's (McCulloch and Pitts).

History of deep learning

Pre 2012 history:

- The original ideas date back to the 1940's (McCulloch and Pitts).
- The modern perceptron was invented by Frank Rosenblatt in 1962.
  - First working algorithms around 1967.

History of deep learning

Pre 2012 history:

- The original ideas date back to the 1940's (McCulloch and Pitts).
- The modern perceptron was invented by Frank Rosenblatt in 1962.
  - First working algorithms around 1967.
- In 1970, Linnainmaa developed automatic differentiation and backpropagation for training neural networks.

History of deep learning

Pre 2012 history:

- The original ideas date back to the 1940's (McCulloch and Pitts).
- The modern perceptron was invented by Frank Rosenblatt in 1962.
  - First working algorithms around 1967.
- In 1970, Linnainmaa developed automatic differentiation and backpropagation for training neural networks.
- Convolutional neural networks for handwritten digit recognition pioneered by LeCun in 1989.

History of deep learning

Modern history (post 2012) of deep learning:

- Before the 2000's neural networks were considered a dead-end in machine learning.

History of deep learning

Modern history (post 2012) of deep learning:

- Before the 2000's neural networks were considered a dead-end in machine learning.
- 2012 Breakthrough: Krizhevsky, Sutskever, and Hinton won the ImageNet competition with Convolutional Neural Networks (CNN).

ImageNet

Image classification problem with 14 million images and 20,000 classes.

History of deep learning

Modern history (post 2012) of deep learning:

- Before the 2000's neural networks were considered a dead-end in machine learning.
- 2012 Breakthrough: Krizhevsky, Sutskever, and Hinton won the ImageNet competition with Convolutional Neural Networks (CNN).

History of deep learning

Modern history (post 2012) of deep learning:

- Before the 2000's neural networks were considered a dead-end in machine learning.
- 2012 Breakthrough: Krizhevsky, Sutskever, and Hinton won the ImageNet competition with Convolutional Neural Networks (CNN).
  - First deep learning model to compete in ImageNet since the competition was founded in 2010.

History of deep learning

Modern history (post 2012) of deep learning:

- Before the 2000's neural networks were considered a dead-end in machine learning.
- 2012 Breakthrough: Krizhevsky, Sutskever, and Hinton won the ImageNet competition with Convolutional Neural Networks (CNN).
  - First deep learning model to compete in ImageNet since the competition was founded in 2010.
  - Their model, AlexNet, got 75% accuracy, was over 10% better than the next best team, and is a model still used today.

History of deep learning

Modern history (post 2012) of deep learning:

- Before the 2000's neural networks were considered a dead-end in machine learning.
- 2012 Breakthrough: Krizhevsky, Sutskever, and Hinton won the ImageNet competition with Convolutional Neural Networks (CNN).
  - First deep learning model to compete in ImageNet since the competition was founded in 2010.
  - Their model, AlexNet, got 75% accuracy, was over 10% better than the next best team, and is a model still used today.
- Now every competitor in the ImageNet competition uses deep learning.

History of deep learning

Modern history (post 2012) of deep learning:

- Before the 2000's neural networks were considered a dead-end in machine learning.
- 2012 Breakthrough: Krizhevsky, Sutskever, and Hinton won the ImageNet competition with Convolutional Neural Networks (CNN).
  - First deep learning model to compete in ImageNet since the competition was founded in 2010.
  - Their model, AlexNet, got 75% accuracy, was over 10% better than the next best team, and is a model still used today.
- Now every competitor in the ImageNet competition uses deep learning.
  - Best results from 2021 are 90.45% (Top 1) and 98.8% (Top 5).

History of deep learning

Modern history (post 2012) of deep learning:

- Before the 2000's neural networks were considered a dead-end in machine learning.
- 2012 Breakthrough: Krizhevsky, Sutskever, and Hinton won the ImageNet competition with Convolutional Neural Networks (CNN).
  - First deep learning model to compete in ImageNet since the competition was founded in 2010.
  - Their model, AlexNet, got 75% accuracy, was over 10% better than the next best team, and is a model still used today.
- Now every competitor in the ImageNet competition uses deep learning.
  - Best results from 2021 are 90.45% (Top 1) and 98.8% (Top 5).
  - Super-human performance: Humans get around 95% (Top 5).

ImageNet Top 1 Accuracy

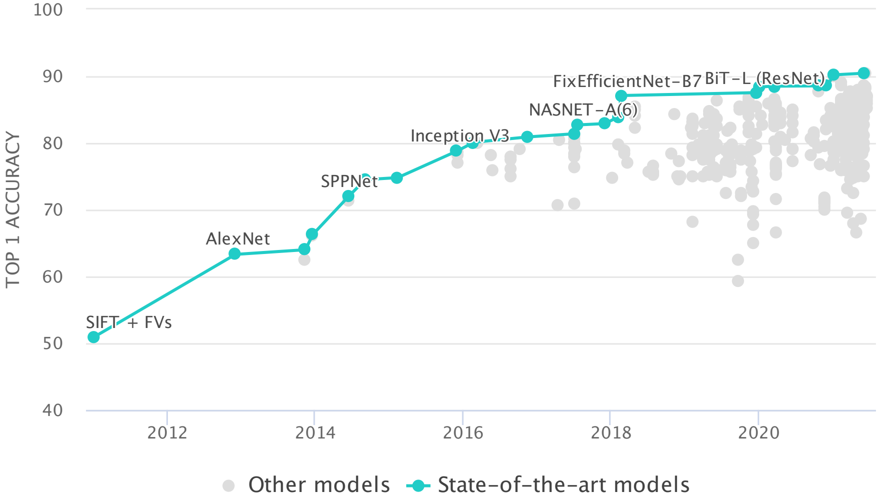

ImageNet Top 5 Accuracy

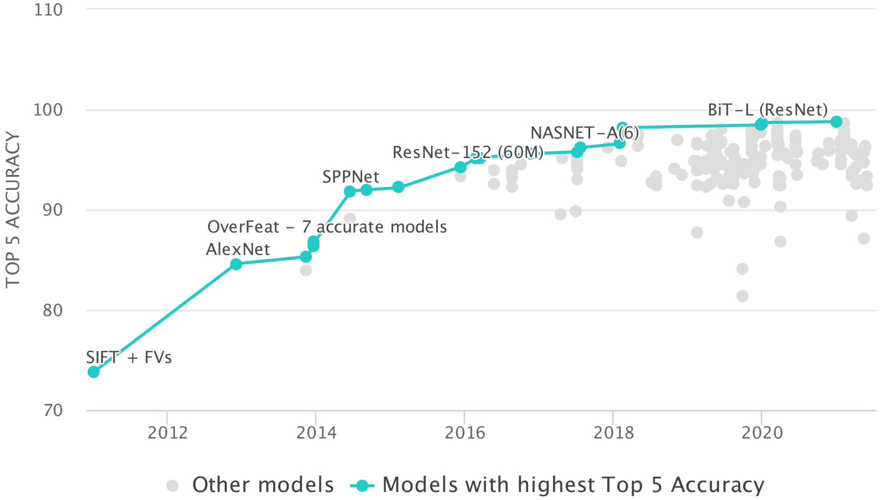

ImageNet: Number of Parameters

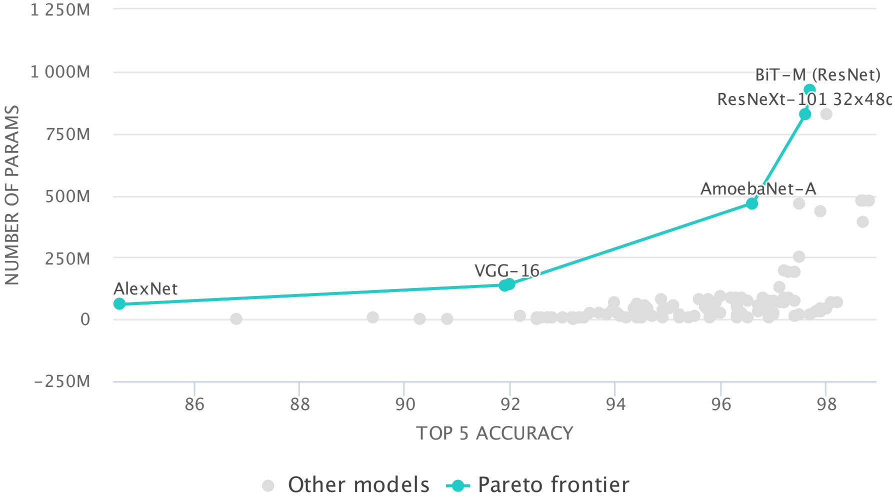

History of deep learning

Due to the success of deep learning at the ImageNet competition, deep learning quickly spread to other research problems.

- Computer vision (self-driving cars, medical image analysis, artistic style transfer)

History of deep learning

Due to the success of deep learning at the ImageNet competition, deep learning quickly spread to other research problems.

- Computer vision (self-driving cars, medical image analysis, artistic style transfer)
- Strategy game playing (Chess, Go, video games, etc.)

History of deep learning

Due to the success of deep learning at the ImageNet competition, deep learning quickly spread to other research problems.

- Computer vision (self-driving cars, medical image analysis, artistic style transfer)
- Strategy game playing (Chess, Go, video games, etc.)
- Recommender systems (Netflix, Amazon, etc.)

History of deep learning

Due to the success of deep learning at the ImageNet competition, deep learning quickly spread to other research problems.

- Computer vision (self-driving cars, medical image analysis, artistic style transfer)
- Strategy game playing (Chess, Go, video games, etc.)
- Recommender systems (Netflix, Amazon, etc.)
- Natural Language processing (Google translate, virtual assistants (e.g., Siri))

History of deep learning

Due to the success of deep learning at the ImageNet competition, deep learning quickly spread to other research problems.

- Computer vision (self-driving cars, medical image analysis, artistic style transfer)
- Strategy game playing (Chess, Go, video games, etc.)
- Recommender systems (Netflix, Amazon, etc.)
- Natural Language processing (Google translate, virtual assistants (e.g., Siri))
- Bioinformatics, drug discovery, and toxicology

History of deep learning

Due to the success of deep learning at the ImageNet competition, deep learning quickly spread to other research problems.

- Computer vision (self-driving cars, medical image analysis, artistic style transfer)
- Strategy game playing (Chess, Go, video games, etc.)
- Recommender systems (Netflix, Amazon, etc.)
- Natural Language processing (Google translate, virtual assistants (e.g., Siri))
- Bioinformatics, drug discovery, and toxicology
- Fraud detection

History of deep learning

Due to the success of deep learning at the ImageNet competition, deep learning quickly spread to other research problems.

- Computer vision (self-driving cars, medical image analysis, artistic style transfer)
- Strategy game playing (Chess, Go, video games, etc.)
- Recommender systems (Netflix, Amazon, etc.)
- Natural Language processing (Google translate, virtual assistants (e.g., Siri))
- Bioinformatics, drug discovery, and toxicology
- Fraud detection
- Many, many, many, others.

Automated image captioning

A woman is throwing a **frisbee** in a park.

A **dog** is standing on a hardwood floor.

A **stop** sign is on a road with a mountain in the background

A little **girl** sitting on a bed with a teddy bear.

A group of **people** sitting on a boat in the water.

A giraffe standing in a forest with **trees** in the background.

[Yann LeCun, Yoshua Bengio, Geoffrey Hinton. Deep learning. Nature, 2015.]

https://deepai.org/machine-learning-model/neuraltalk

Word2Vec

Word2Vec maps words into vectors in $\mathbb{R}^m$ in such a way that relationships between words are encoded into the vector geometry.

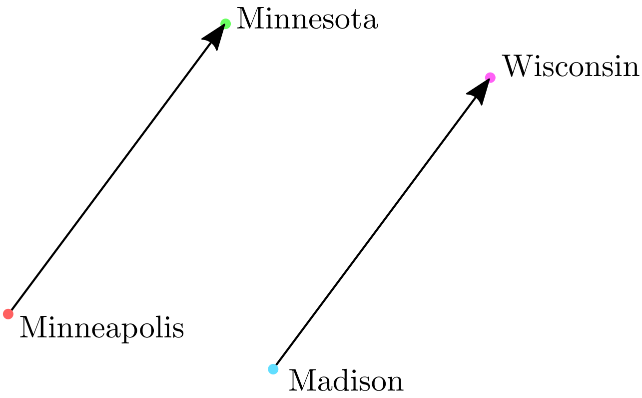

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man $+$ Woman $=$ Queen (most famous example)

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings =

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane
- Bird - Wings + Legs =

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane
- Bird - Wings + Legs = Turtle

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane
- Bird - Wings + Legs = Turtle
- Donut - Icing + Cheese =

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane
- Bird - Wings + Legs = Turtle
- Donut - Icing + Cheese = Sandwich

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane
- Bird - Wings + Legs = Turtle
- Donut - Icing + Cheese = Sandwich
- France - Paris + Tokyo =

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane
- Bird - Wings + Legs = Turtle
- Donut - Icing + Cheese = Sandwich
- France - Paris + Tokyo = Japan

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane
- Bird - Wings + Legs = Turtle
- Donut - Icing + Cheese = Sandwich
- France - Paris + Tokyo = Japan
- Minnesota - Minneapolis + Madison = Wisconsin

Word2Vec examples

Using Word2Vec we can do arithmetic with words.

- King - Man + Woman = Queen (most famous example)
- Car - Wheels + Wings = Airplane
- Bird - Wings + Legs = Turtle
- Donut - Icing + Cheese = Sandwich
- France - Paris + Tokyo = Japan
- Minnesota - Minneapolis + Madison = Wisconsin
- Choose your own example...

Good Old-Fashioned AI

Machine learning (and deep learning) learn by examples and are not explicitly programmed to complete tasks.

Good Old-Fashioned AI

Machine learning (and deep learning) learn by examples and are not explicitly programmed to complete tasks.

Good old-fashioned AI refers to the AI that is explicitly programmed by a human expert to complete tasks in the way a human would.

Good Old-Fashioned AI

Machine learning (and deep learning) learn by examples and are not explicitly programmed to complete tasks.

Good old-fashioned AI refers to the AI that is explicitly programmed by a human expert to complete tasks in the way a human would.

Deep Blue vs. Garry Kasparov, 1996-1997

Chess-playing AI (Deep Blue)

- Deep blue was designed by IBM with chess grandmaster Joel Benjamin
- It was programmed with a playbook of opening moves and the ability to recognize certain strategies.
- Deep blue ultimately won due to the ability to look ahead more moves than a human chess player can.

Chess-playing AI (Deep Blue)

- Deep blue was designed by IBM with chess grandmaster Joel Benjamin
- It was programmed with a playbook of opening moves and the ability to recognize certain strategies.
- Deep blue ultimately won due to the ability to look ahead more moves than a human chess player can.
- This approach fails for more complicated games, like the Chinese board game Go:

Deep reinforcement learning

Deep reinforcement learning deals with how an intelligent agent ought to take action in an environment to maximize a notion of cumulative reward (e.g., strategy game playing).

Deep reinforcement learning

Deep reinforcement learning deals with how an intelligent agent ought to take action in an environment to maximize a notion of cumulative reward (e.g., strategy game playing).

In a major breakthrough in 2015, Deep Mind's AlphaGo was the first computer to beat a professional Go player.

AlphaGo and AlphaGo Zero

- AlphaGo is a convolutional neural network, trained on data from human Go games.

AlphaGo and AlphaGo Zero

- AlphaGo is a convolutional neural network, trained on data from human Go games.
- In 2017, Deep Mind introduced AlphaGo Zero, which did not use any human data in training.

AlphaGo and AlphaGo Zero

- AlphaGo is a convolutional neural network, trained on data from human Go games.
- In 2017, Deep Mind introduced AlphaGo Zero, which did not use any human data in training.
  - AlphaGo Zero learned to play Go by playing against itself repeatedly in a deep reinforcement learning algorithm.

AlphaGo and AlphaGo Zero

- AlphaGo is a convolutional neural network, trained on data from human Go games.
- In 2017, Deep Mind introduced AlphaGo Zero, which did not use any human data in training.
  - AlphaGo Zero learned to play Go by playing against itself repeatedly in a deep reinforcement learning algorithm.
- Quotes from professional Go players during 2015-2017 matches:
  - "I misjudged the capabilities of AlphaGo and felt powerless."
  - "It knows everything and can play creatively."
  - "All but the very best Go players craft their style by imitating top players. AlphaGo seems to have totally original moves it creates itself."

Deep Reinforcement Learning

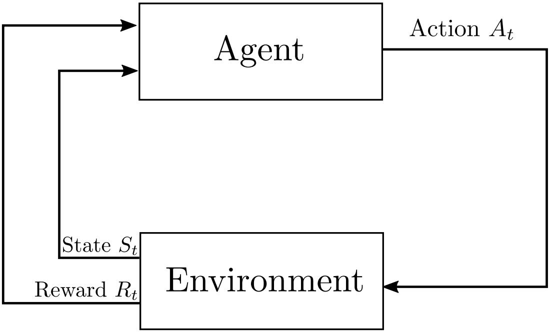

Generative Adversarial Networks (GAN)

Generate new samples from a given training dataset.

Generative Adversarial Networks (GAN)

Generate new samples from a given training dataset.

Uses 2 neural networks, a generator and a discriminator network.

Generative Adversarial Networks (GAN)

Generate new samples from a given training dataset.

Uses 2 neural networks, a generator and a discriminator network.

- Discriminator is trained to recognize images from the training set.

Generative Adversarial Networks (GAN)

Generate new samples from a given training dataset.

Uses 2 neural networks, a generator and a discriminator network.

- Discriminator is trained to recognize images from the training set.
- Generator's goal is to trick the discriminator into thinking its generated images belong to the training set

Generative Adversarial Networks (GAN)

Generate new samples from a given training dataset.

Uses 2 neural networks, a generator and a discriminator network.

- Discriminator is trained to recognize images from the training set.
- Generator's goal is to trick the discriminator into thinking its generated images belong to the training set
  - The input of the generator is random noise.

Generative Adversarial Networks (GAN)

Generate new samples from a given training dataset.

Uses 2 neural networks, a generator and a discriminator network.

- Discriminator is trained to recognize images from the training set.
- Generator's goal is to trick the discriminator into thinking its generated images belong to the training set
  - The input of the generator is random noise.
- Training pits the generator against the discriminator until the generator can consistently fool the discriminator.

These people do not exist (generated by a GAN)

https://thispersondoesnotexist.com/

Neural Style Transfer

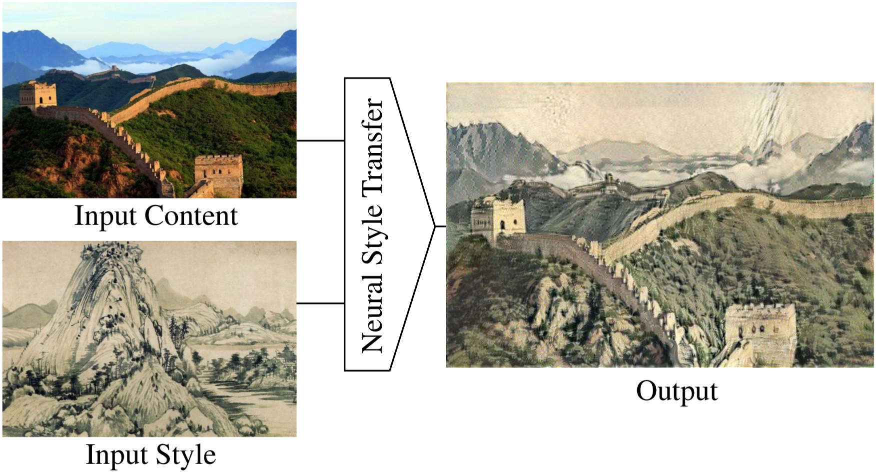

Neural Style Transfer

(a) Content

(b) Style

Neural Style Transfer

(c) Small Stroke Size

(d) Large Stroke Size

Neural Style Transfer

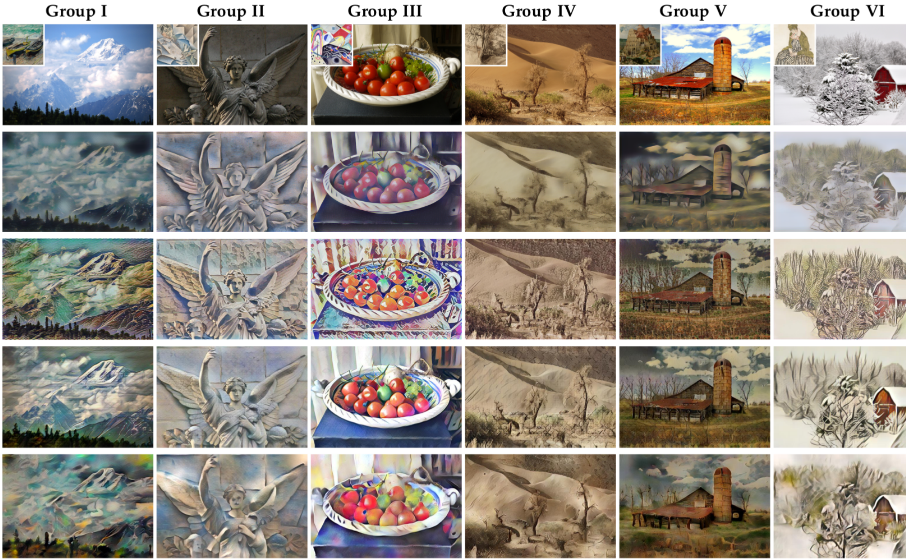

Neural Style Transfer

Neural Style Transfer works by minimizing the loss

$$E(I) = \sum_{i=1}^{L} \|F_i(I) - F_i(I_{\text{style}})\|^2 + \lambda \|I - I_{\text{content}}\|^2,$$

where

- I_{style} is the style image.
- I_{content} is the content image.
- F_i are the features extracted from the i^{th} layer of a pre-trained convolutional neural networks, usually pretrained on imagenet.

Autoencoders

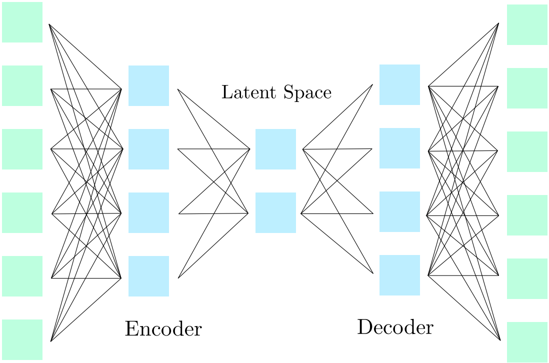

- Nonlinear version of PCA.
- Applications
  - Image or data compression.
  - Simplifying data or extracting features.