GT Deep Learning: 01

Vincent Perrollaz

Institut Denis Poisson

27 Janvier 2026

Artifical Intelligence

Machine Learning

Deep Learning

Supervised Learning

Unsupervised Learning

Reinforcement learning

A.I. Definition

50 Turing "I propose to consider the question 'can machines think'?"

56 Dartmouth "the ability to solve hard problems"

60' Dijkstra[1] "The question of whether a computer can think is no more interesting than the question of whether a submarine can swim."

00' Russel Norvig "Acting Humanly" vs "Thinking Humanly" vs "Thinking Rationally" vs "Acting Rationally"

[1] Arrogance in computer science is measured in nano-dijkstra. - Alan Kay

Artifical Intelligence

Machine Learning

Deep Learning

Supervised Learning

Unsupervised Learning

Reinforcement learning

The bitter lesson[2] by `Rich Sutton`

The biggest lesson that can be read from 70 years of AI research is that general methods that leverage computation are ultimately the most effective, and by a large margin...

[2] clickable hyperlink

ML =

1. Statistical Modeling $\Longrightarrow$ Functions + Loss
2. Optimization

Examples:

- Linear Regression, Ridge Regression, Lasso
- Logistic Regression, Naive Bayes
- Nearest neighbors
- Support Vector Machines
- Decision Trees
- Random Forest
- Gradient Boosting
- ...

Artifical Intelligence

Machine Learning

Deep Learning

Supervised Learning

Unsupervised Learning

Reinforcement learning

Neural Networks

Class of functions

- Easy to picture
- Flexible architecture
- Great analytic properties
- Great algorithmic properties

Building block:

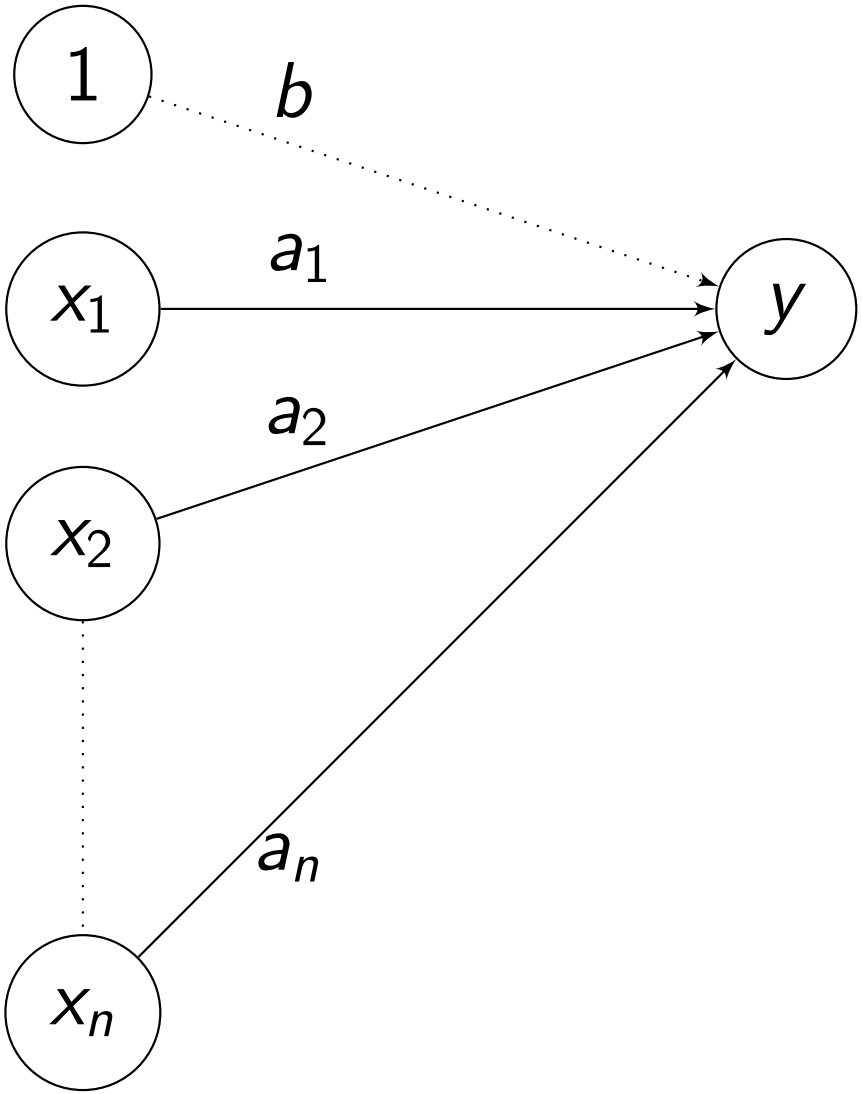

$$\sigma\left(\sum_{i=1}^{n} a_i x_i + b\right) = y$$

Multi Layer Perceptron:

input $\mathbb{R}^6$, hidden layers, output $\mathbb{R}^4$

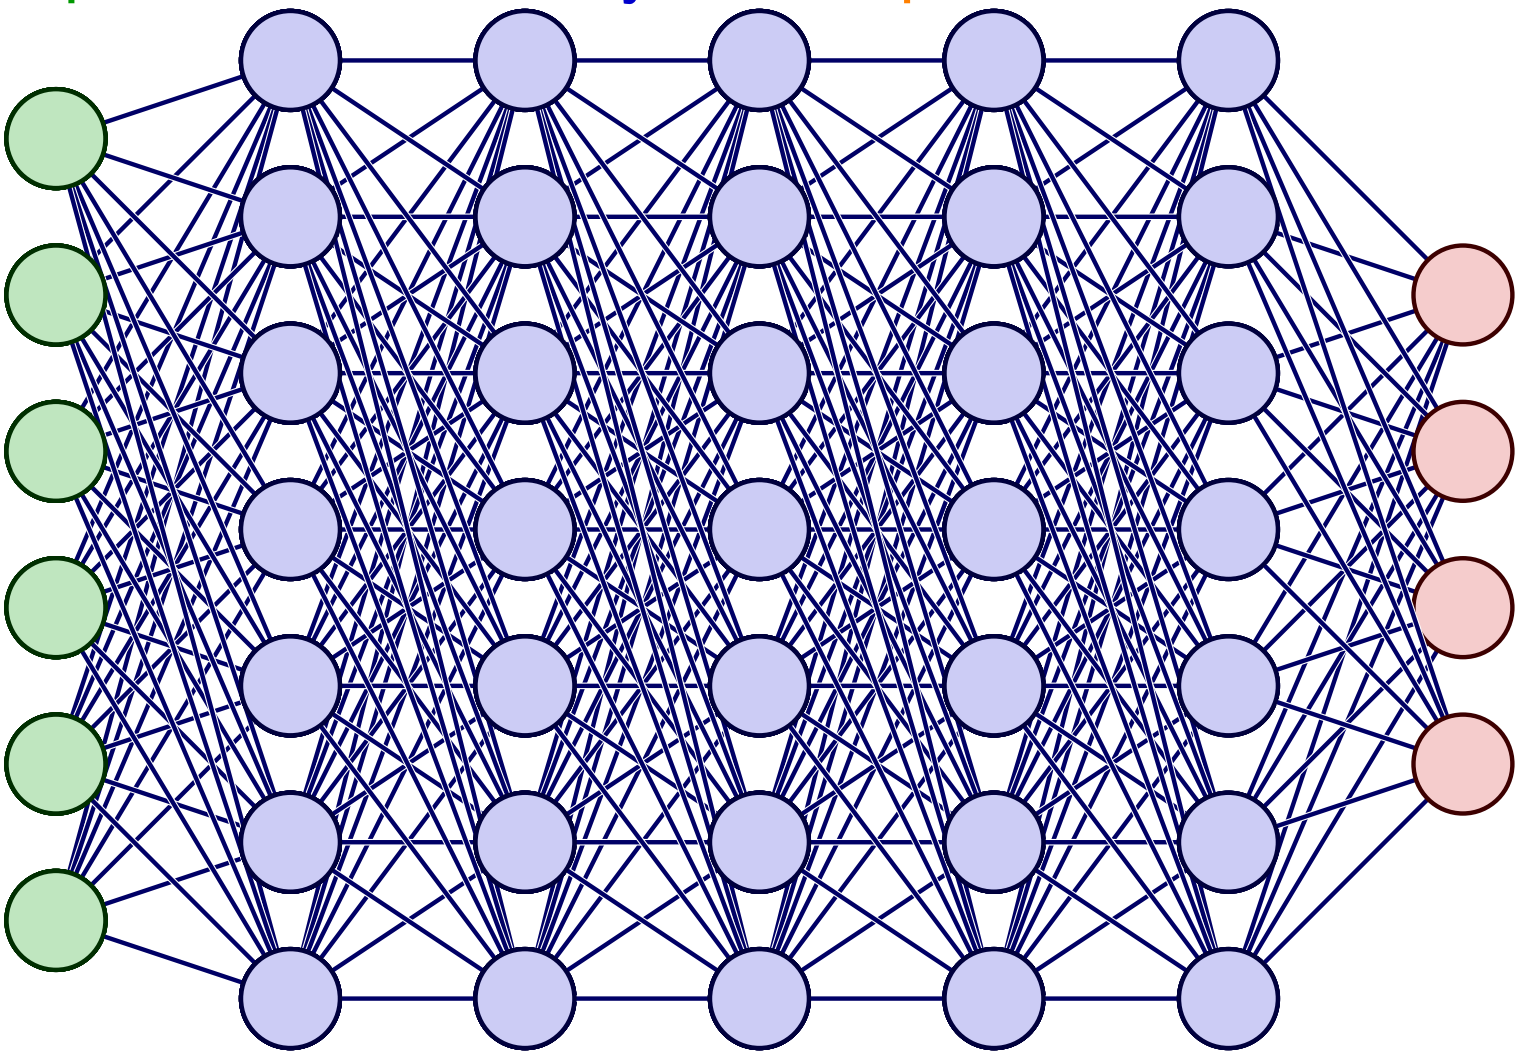

Deep Convolutional Network

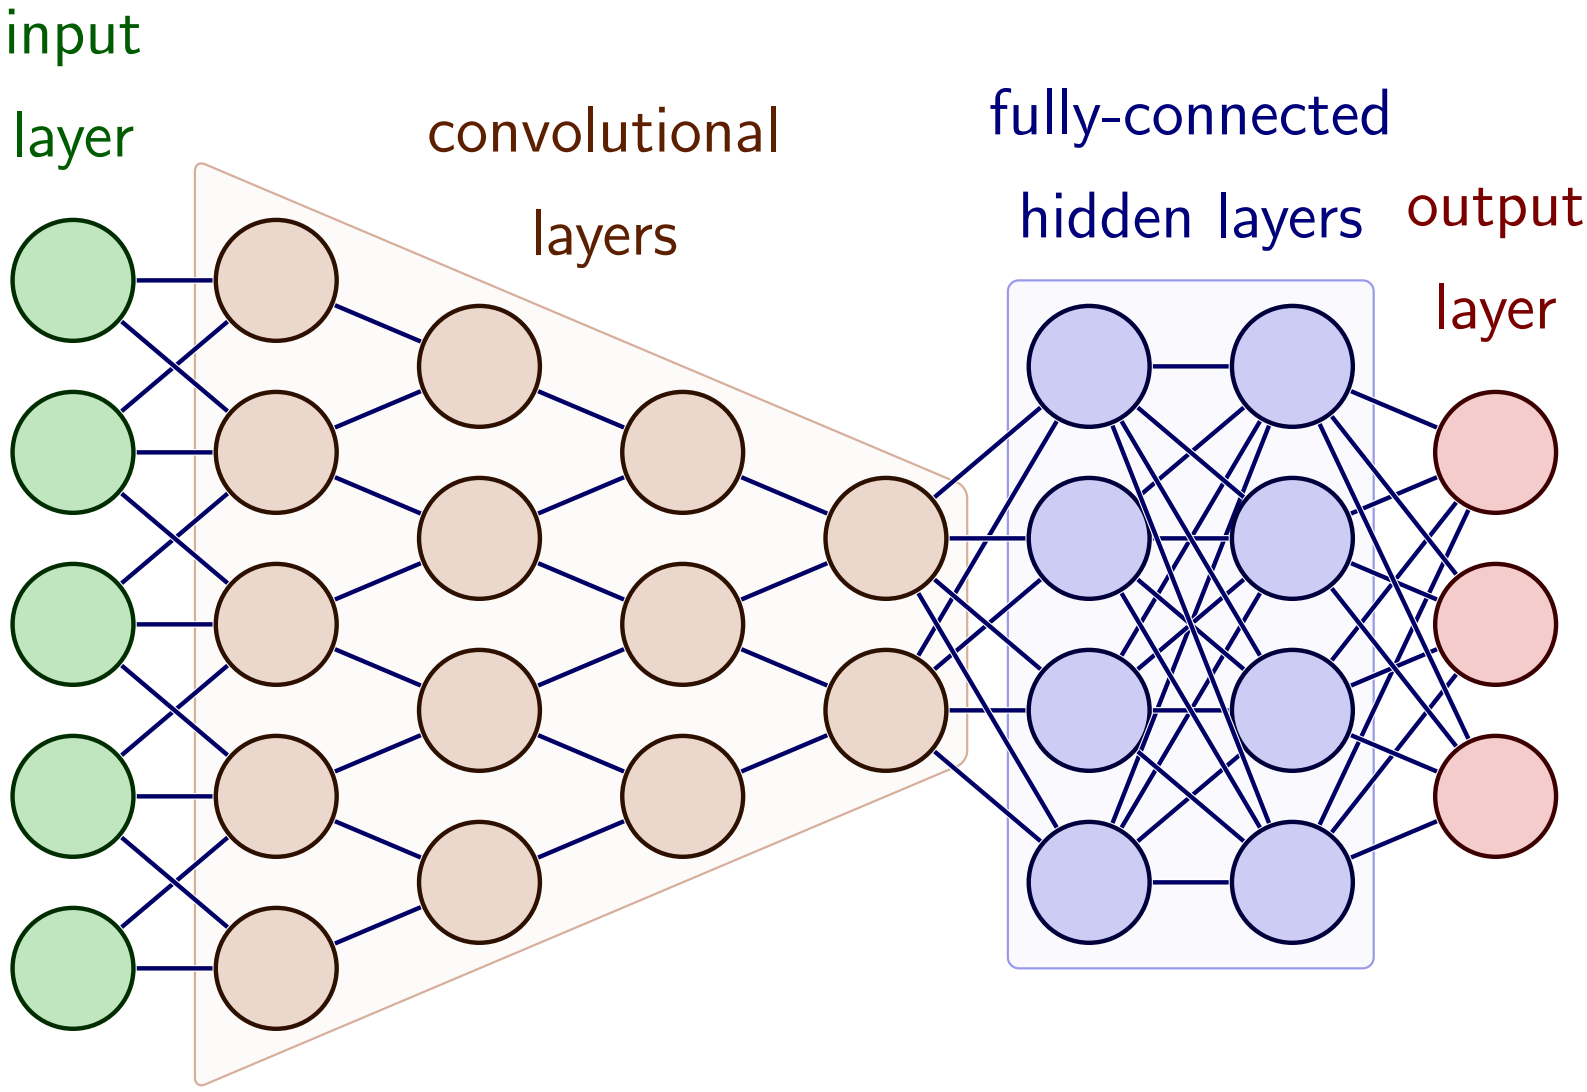

Autoencoder

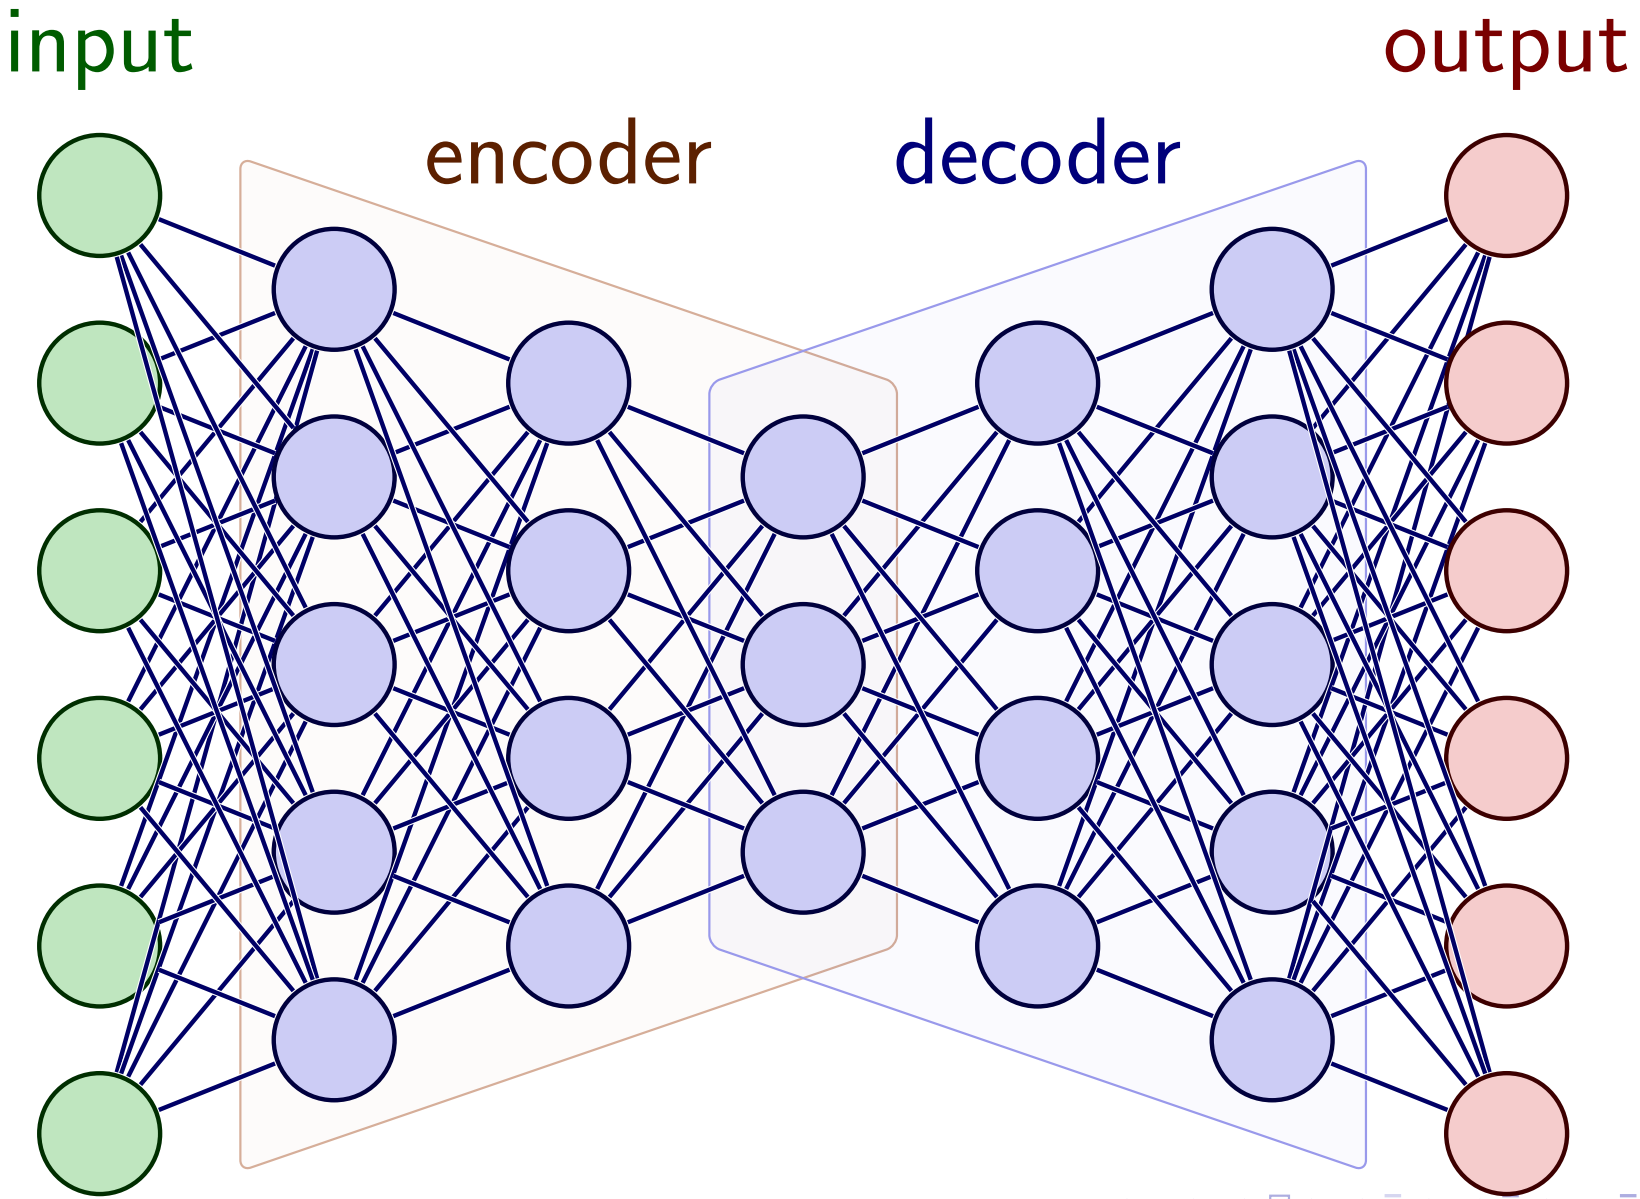

Artifical Intelligence

Machine Learning

Deep Learning

Supervised Learning

Unsupervised Learning

Reinforcement learning

- Supervised learning = "learn by examples"
- e.g. Linear Regression[3]

$$\text{Given} : (x_k, y_k)_{1 \leq k \leq N} \in (\mathbb{R}^2)^N$$

$$\text{Find} : \arg\min_{(a,b) \in \mathbb{R}^2} \frac{1}{N} \sum_{k=1}^{N} (a x_k + b - y_k)^2$$

- Regression = "continuous output"
- Classification = "discrete output"
- Digits classification in accompanying notebook

[3] Legendre-Gauss™

Artifical Intelligence

Machine Learning

Deep Learning

Supervised Learning

Unsupervised Learning

Reinforcement learning

- Raw data, no output
- Clustering $=$ "Find groups"
- Embedding $\approx$ "Change of variables"
- Simple example in accompaying notebook
- Manifold Hypothesis
- NLP[4] $\approx \mathbb{R}^{800}$
- KING - MAN $+$ WOMEN $\approx$ QUEEN

[4] Natural Language Processing

Artifical Intelligence

Machine Learning

Deep Learning

Supervised Learning

Unsupervised Learning

Reinforcement learning

- RL = “Teach with stick and carrot”
- e.g. Play Tic Tac Toe
- Table: (State, Possible Action) $\rightarrow$ Value
- Play Explore/Exploit
- Win : Increase value previous actions
- Loss : Decrease value previous actions
- Bitter Lesson: `alphago` vs `alphazero`

- OLLAMA demo Qwen3:4b
- Local on my machine
- Qwen3 model type (open source by Alibaba)
- 4b number of parameters (edges) in neural net