Neural Networks 1 - Introduction

18NES1 - Lecture 1, Summer semester 2024/25

Zuzana Petříčková

February 18, 2025

Neural Networks 1 - Introduction

Introduction to Artificial Intelligence

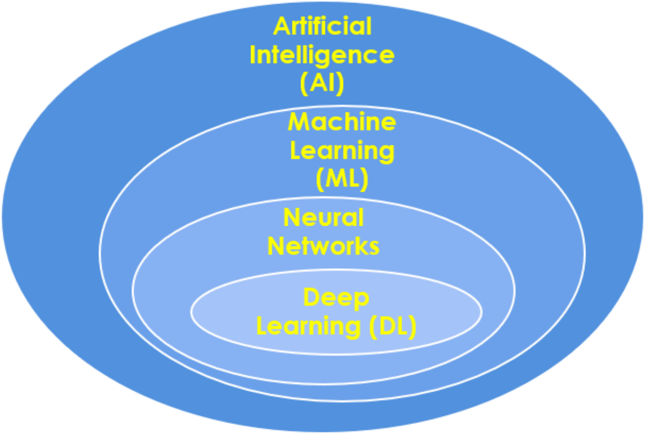

- What is AI (Artificial Intelligence)?

Introduction to Artificial Intelligence

- Turing Test (Alan Turing, 1950)
  - A computer/program passes the test if a human, during a five-minute conversation, cannot recognize that they are not talking to another human.
- First chatbot: Eliza (1966)

Introduction to Artificial Intelligence

- Turing Test (Alan Turing, 1950)
  - A computer/program passes the test if a human, during a five-minute conversation, cannot recognize that they are not talking to another human.
- The classical definition is no longer valid:
  - The first chatbot to surpass the Turing Test was Google BERT, 2022.

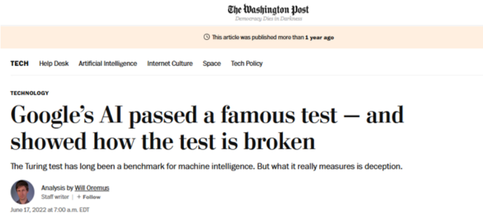
The Washington Post
Democracy Dies in Darkness

This article was published more than **1 year ago**

TECH Help Desk Artificial Intelligence Internet Culture Space Tech Policy

TECHNOLOGY

Google's AI passed a famous test — and showed how the test is broken

The Turing test has long been a benchmark for machine intelligence. But what it really measures is deception.

Analysis by Will Oremus
Staff writer | + Follow
June 17, 2022 at 7:00 a.m. EDT

Introduction to Artificial Intelligence

- **Classical definition:** The ability of machines/computer programs to replicate human capabilities that we consider **intelligent**:
  - The ability to reason and solve problems, and plan
  - The ability to adapt to new environments and learn
  - Creativity, innovation, and decision-making
- **Modern definition:** A scientific field focused on developing **advanced systems** to solve **complex problems**
  - Image recognition, language translation, playing chess, medical diagnostics, autonomous vehicles, . . .
  - **AI is integrated into everyday life:**
    - Content personalization on social media, personalized advertisements, spam detection, search engines, chatbots, and more

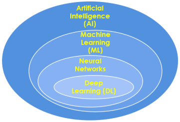

Machine Learning

- Models and techniques that allow a computer system to **learn** from data or past experiences.
- Often inspired by biological principles, the model "builds itself."

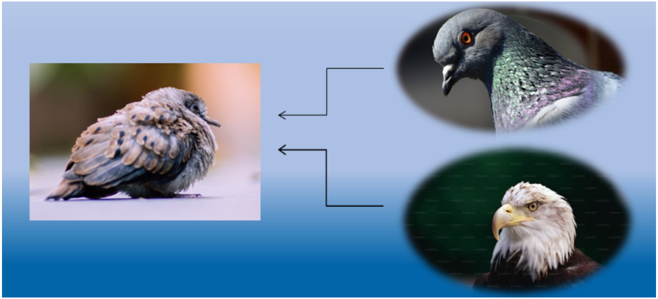

Machine Learning

Principle

- Based on biological principles, the model "builds itself."
- The computational model learns from data (training set) or past experiences.

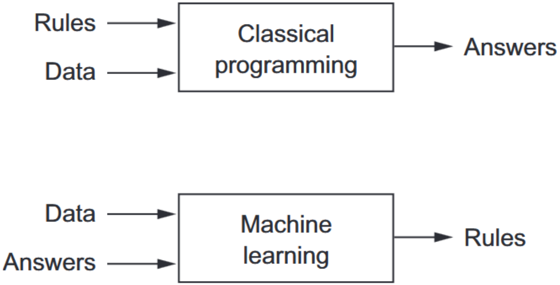

F. Chollet: Deep Learning with Python, Fig. 1.2

Machine Learning - Three Fundamental Types

- **Supervised Learning**
  - The model learns from labeled data (e.g., image classification)
  - Training set in the form of [*input, expected output*]
- **Unsupervised Learning (Self-Organization, Self-Supervised Learning)**
  - The model identifies patterns in unlabeled data (e.g., clustering images)
  - Training set in the form of [*input*]
- **Reinforcement Learning**
  - The model learns an optimal strategy based on past experiences, often using rewards and penalties (e.g., robotic soccer)

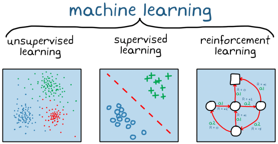

https://www.mathworks.com/discovery/reinforcement-learning.html

Machine Learning - Three Fundamental Types

Supervised Learning

- **Training set format:** [input, expected output]
- **Learning objective:** The model should approximate an unknown function as accurately as possible $\rightarrow$ predict the correct output for any given input.
- **Generalization:** The model should produce accurate outputs even for data not included in the training set.
- **Typical applications:**
  - Medical diagnosis, fraud detection in banking
  - Time series forecasting
  - Image classification and segmentation
  - Natural language processing, speech recognition

Machine Learning - Three Fundamental Types

Supervised Learning - Task Types

- **Classification:** Predicting a class/category

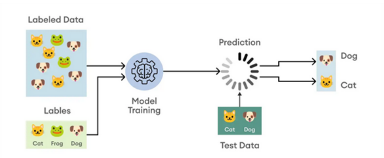

- **Regression:** Predicting a numerical value (price, temperature, handwriting slant, etc.)

- **Structured Data Learning** (e.g., sentences in natural language)

Machine Learning - Three Fundamental Types

Unsupervised Learning (Self-Organization)

- **Training set format:** [input]
- **Learning objective:** Identify inner structure or patterns in data
- **Applications: Dimensionality reduction** (data compression, visualization),
  Anomaly detection (e.g., in banking transactions),
  Clustering (e.g., grouping customers based on behavior, plagiarism detection),
  e-commerce (recommendation systems)

Types of tasks:

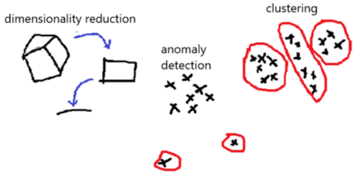

https://towardsdatascience.com/unsupervised-learning-algorithms-cheat-sheet-d391a39de44a

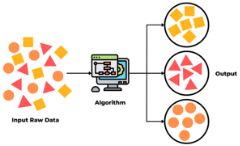

https://eastgate-software.com/what-is-unsupervised-learning/#ftoc-heading-7

Machine Learning - Three Fundamental Types

Unsupervised Learning (Self-Organization)

- Methods: Clustering, association rules, autoencoders, generative models

Reinforcement Learning

- The program learns an optimal strategy based on past experiences.
- Often involves a system of rewards and penalties.
- Methods: Q-learning, Deep Q-Networks, etc.
- Applications: Gaming industry, robotics, resource management, machine translation, ...

Typical Workflow for Solving a Machine Learning Task

- **Problem Definition**
  - What is the task? What kind of data do we have? What type of problem are we solving? What is the exact goal?
  - What machine learning models are suitable for this task?
  - Conduct a literature review of existing solutions, limitations, and challenges.
- **Data Collection**
  - Gather relevant data and try to understand it thoroughly.
  - Define a success metric (various evaluation metrics such as accuracy, precision, recall, etc.).

Typical Workflow for Solving a Machine Learning Task

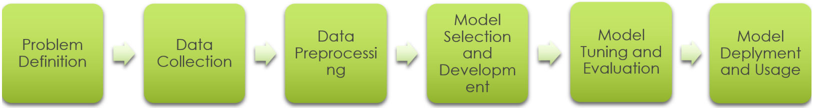

- **Data Preprocessing**
  - Transform data into a format suitable for machine learning models.
  - Techniques include vectorization, normalization, handling missing values, and data augmentation.
- **Model Selection and Development**
  - Choose an evaluation protocol (e.g., k-fold cross-validation).
  - Build a simple baseline model for reference.
  - Determine the model type — depends on the problem domain.
  - Select a specific model within the chosen type (tuning hyperparameters, defining architecture, using pre-existing templates).

Typical Workflow for Solving a Machine Learning Task

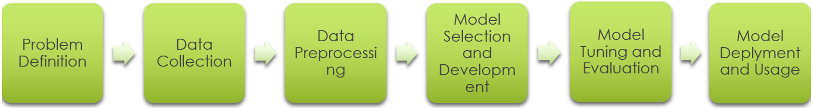

- **Model Tuning and Evaluation**
  - Experiment with different architectures.
  - Apply optimization techniques such as dropout, regularization, and hyperparameter tuning.
  - Evaluate the model—preferably using unseen (test) data.
- **Deployment and Maintenance**
  - Deploy the model across different platforms and devices.
  - Monitor model performance over time and update as necessary.

Machine Learning

Examples of Classical Machine Learning Models

- Stored data models
- Linear or non-linear functions (e.g., linear regression, logistic regression)
- Decision trees, rule-based systems
- Bayesian networks, fuzzy systems, evolutionary algorithms
- Artificial neural networks
  - Inspired by the structure and function of the human brain:
    - Information processing (speed, parallelism)
    - Information storage mechanisms
    - Redundancy and control mechanisms

Modern Machine Learning . . . Deep Learning

- Models based on deep neural networks

Machine Learning — Deep Learning

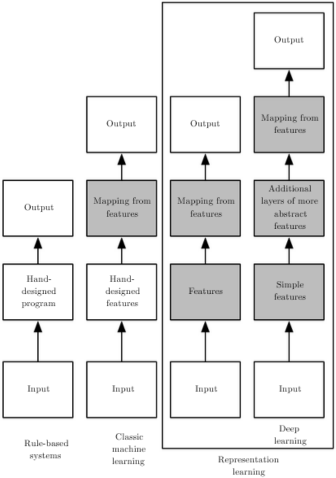

I. Goodfellow, Y. Bengio, and A. Courville: Deep Learning, 2016, Figure 1.5

- Utilizes artificial neural networks with multiple layers (so-called deep networks).
- The model automatically extracts features from data, reducing the need for extensive preprocessing.

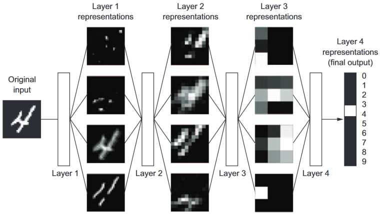

F. Chollet: Deep Learning with Python, Fig. 1.6

Machine Learning — Deep Learning

Advantages:

- Capable of processing vast amounts of data and identifying complex patterns.
- Flexible across various tasks (classification, recognition, content generation, etc.).

Disadvantages:

- Requires large amounts of data and significant computational power for training and inference.
- Functions as a "black box" — limited interpretability of learned representations.

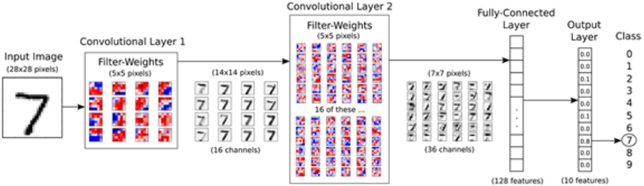

Neural Networks 1 - Introduction

Artificial Neural Networks — A Brief History

https://www.codesofinterest.com/p/build-deeper.html

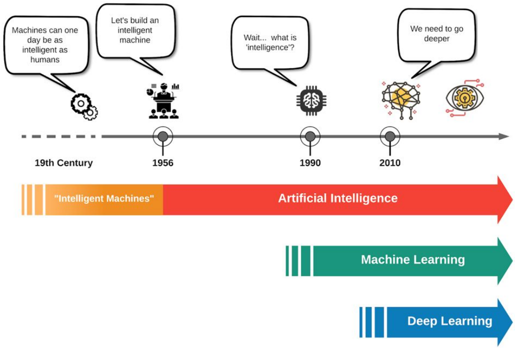

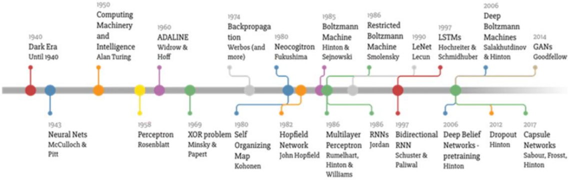

Mourtzis, Dimitris & Angelopoulos, John. (2020). An intelligent framework for modelling and simulation of artificial neural networks (ANNs) based on augmented reality. International Journal of Advanced Manufacturing Technology. 111. 10.1007/s00170-020-06192-y.

Artificial Neural Networks — A Brief History

Development Progressed in Waves:

- Periods of rapid advancement and high expectations were followed by phases of disappointment and stagnation.

Key Time Periods:

- **1940 – 1960:** Theoretical foundations
- **1960 – 1970:** First boom — the "single neuron era"
- **1970 – 1980:** First "AI Winter" for neural networks
- **1980 – 1990:** Second boom — the era of "shallow" neural networks
- **1990 – 2000:** Gradual stabilization of the field
- **2000 – 2010:** Second "AI Winter" for neural networks
- **2010 – present:** Third boom — the era of "deep" neural networks

Artificial Neural Networks — A Brief History

Early Foundations (1940-1960)

- Attempts to model the functioning of biological neurons.
- **1943: First mathematical model of a neuron** (W. McCulloch, W. Pitts) — capable of representing logical and arithmetic functions.

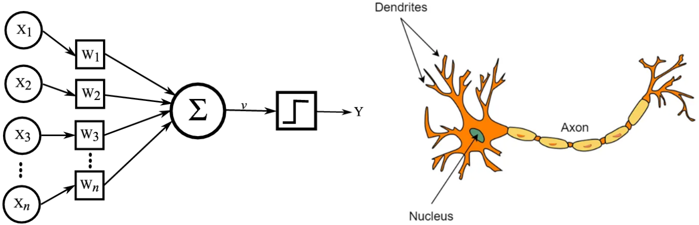

- **1949: Mathematical concept of learning** (D. Hebb) — introduced the first learning algorithm for artificial neurons, modeling conditioned reflexes.
- **1951: First neurocomputer, SNARC** (M. Minsky).

Artificial Neural Networks — A Brief History

First Boom (1960–1970): The "Single Neuron Era"

- **1957: Perceptron** (F. Rosenblatt) — first practical artificial neuron model with real-valued parameters and a functional learning algorithm. Sparked enormous enthusiasm.
- **1958: First successful neurocomputer,** *Mark I Perceptron* (F. Rosenblatt, C. Wightman).
- **1962: Adaline** and the **sigmoid activation function** (B. Widrow, M. Hoff).
- Rapid progress in neurocomputing in the 1960s, but soon faced fundamental limitations.

Artificial Neural Networks — A Brief History

- **1969: Perceptrons** (M. Minsky, S. Papert) — demonstrated the limitations of perceptrons, showing they cannot model even basic logical functions.
- What if multiple neurons were connected into a network?
  $\rightarrow$ A promising idea, but how would such a model be trained? Existing learning algorithms for perceptrons and linear neurons were insufficient.

1970 – 1980: First "AI Winter"

- Disillusionment due to unfulfilled expectations.
- Declining credibility of the field.
- Decreased funding and reduced interest in AI research.

Artificial Neural Networks — A Brief History

Second Boom (1980 - 1990): The Era of "Shallow" Neural Networks

- Renewed interest in the field (John Hopfield, DARPA, etc.); introduction of the multilayer perceptron (MLP).
- **1986: Backpropagation Algorithm** for training MLPs (P. Werbos and D. Rumelhart, with earlier work by G. Hinton and Y. LeCun).
    - A fundamental concept that remains in use today.

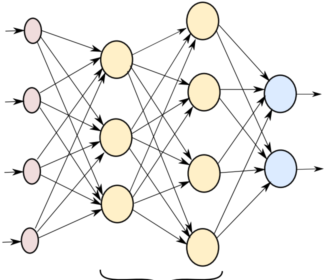

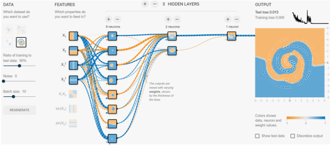

https://playground.tensorflow.org/

Artificial Neural Networks — A Brief History

Second Boom (1980 - 1990): The Era of "Shallow" Neural Networks

Emergence of new neural network models composed of multiple neurons:

- **Multilayer Perceptron (MLP)**
- **Kohonen Maps** (T. Kohonen)
- **Hopfield Networks** (J. Hopfield)
- **Radial Basis Function (RBF) Networks** (J. Moody, C. Darken)
- **Growing Neural Gas (GNG) Model** (B. Fritzke)
- **Support Vector Machines (SVMs)** (V. Vapnik)
- **Extreme Learning Machines (ELMs)** (G.-B. Huang)
- **Recurrent Neural Networks (RNNs)** (e.g., Jeffrey Elman)
- **Convolutional Neural Networks (CNNs)** (Y. LeCun, Y. Bengio, et al.)

Artificial Neural Networks — A Brief History

Second Boom (1980 - 1990): The Era of "Shallow" Neural Networks

- **1987: IEEE International Conference on Neural Networks**
- **1987: Establishment of INNS (International Neural Network Society)** and launch of the journal *Neural Networks*.
- International journals on neural networks:
  - *Neural Computation* (1989)
  - *IEEE Transactions on Neural Networks* (1990)
  - *Neural Network World* (1991)

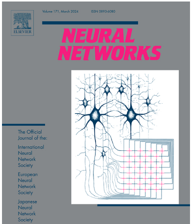

Artificial Neural Networks — A Brief History

Stabilization Period (1990-2000)

- MLP (Multilayer Perceptron), RNN (Recurrent Neural Network), Kohonen Maps, and SVM (Support Vector Machines) became dominant models.
- Efforts to address learning challenges in neural networks:
  - Application of advanced optimization techniques.
  - Enhancing robustness, generalization, and mitigating overfitting (where models perform perfectly on training data but fail on new data).
  - Learning strategies — parallelization and computational efficiency.
- Attempts to train deeper neural networks:
  - Beyond MLP, CNNs (Convolutional Neural Networks) were also developed.
  - However, computational limitations (processing power and memory) imposed significant barriers.
  - The vanishing and exploding gradient problem emerged as a major challenge.

Artificial Neural Networks — A Brief History

2010 – Present: The Third Boom — The Era of Deep Neural Networks

- **What triggered the boom?**
  - Advances in GPUs, cloud computing, and easy access to large datasets.
  - Improvements in learning algorithms (e.g., Deep Learning, Y. LeCun, Y. Bengio, G. Hinton, 2015).
  - **Breakthrough: CNNs won the ILSVRC competition (2012)**
    - **ImageNet Dataset:** 16 million color images across 20,000 categories.

Artificial Neural Networks — A Brief History

2010 – Present: The Third Boom — The Era of Deep Neural Networks

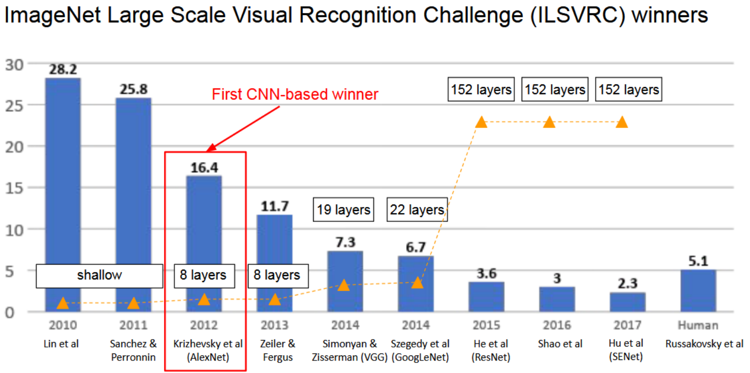

https://cs231n.stanford.edu/2021/slides/2021/lecture_9.pdf (Fei-Fei Li, Ranjay Krishna, Danfei Xu)

Artificial Neural Networks — A Brief History

2010 – Present: The Third Boom — The Era of Deep Neural Networks

- Rapid emergence of new deep learning models and architectures:
  - **Modern Recurrent Neural Networks (RNNs):** LSTM (Hochreiter, Schmidhuber, 1997).
  - **Generative Adversarial Networks (GANs)** (I. Goodfellow, 2015).
  - **First transformer-based generative language model:** GPT-3 (2020, OpenAI), later followed by ChatGPT (2022).
  - **First text-to-image model:** Stable Diffusion (2022, CompVis).

Modern Deep Neural Network Architectures

Key Deep Learning Architectures

- **Multilayer Perceptron (MLP, sometimes DNN)** — A standard, universal model.
- **Convolutional Neural Networks (CNNs)** —
  - Image and video processing.
  - Classification, recognition, and segmentation tasks.
- **Autoencoders** —
  - Unsupervised learning.
  - Feature extraction, data denoising, compression, and reconstruction.

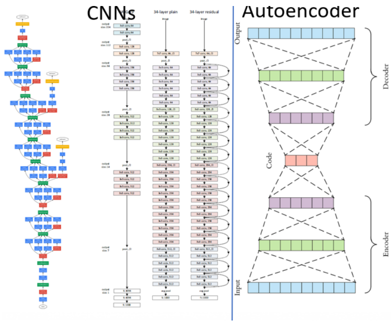

Modern Deep Neural Network Architectures

Recurrent Neural Networks (RNNs)

- Used for sequential data analysis (time series, speech, text, handwriting).
  - **Long Short-Term Memory Networks (LSTMs)**
    - Effective for processing sequential data (e.g., speech and handwriting recognition).
    - Captures long-range dependencies.
  - **Gated Recurrent Unit Networks (GRUs)**
    - Efficient for modeling sequential data.
    - Applied in speech recognition and machine translation.

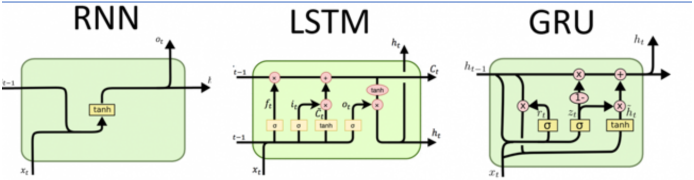

Modern Deep Neural Network Architectures

Key Advanced Deep Learning Architectures

- **Generative Adversarial Networks (GANs)** — Generate new data based on learned patterns.
- **Deep Belief Networks (DBNs)** — Unsupervised generative models.
- **Deep Q-Networks (DQNs)** — Used in reinforcement learning applications.
- **Siamese Networks** — Applied in image recognition and object tracking; measures similarity between two inputs.
- **Capsule Networks (CapsNets)** — Used for image recognition; models hierarchical relationships between object parts.
- **Transformer Networks (BERT, GPT)** — Used for natural language processing (text classification, translation, etc.).

Course Organization

What will we cover in the Neural Networks 1 course?

- This is an introductory course, so we will focus on fundamental concepts.
- We will explore foundational "shallow" neural network models, starting with the perceptron and extending to models with a single hidden layer (e.g., MLP, Kohonen maps).
- We will also introduce basic deep learning models, including the multilayer perceptron (MLP) and convolutional neural networks (CNNs).
- Our goal is to understand how artificial neural networks work internally, how they behave in different scenarios, and how to correctly apply them to various problem types.
- Practical exercises will be conducted in Python.

Course Organization

Interested in learning about more advanced, modern, and deeper neural network models?

- More in-depth knowledge about deep neural networks is covered in the following advanced courses:
  - **Neural Networks 2** — A practical, engineering-focused course on deep neural networks.
  - **Machine Learning 2**
  - **Applications of Optimization Methods**, and other related subjects.