Introduction to Artificial Intelligence

from data analysis to generative AI

Alberto Ciaramella, Marco Ciaramella

IntelliSemantic

13 February 2025

Index

1. AI definition (what?)
2. AI technologies (how?)
3. AI applications (why?)
4. AI opportunities and challenges
5. The AI book
6. Follow-up
7. Questions and answers

1. Artificial Intelligence: definition

- “Artificial Intelligence is the simulation of human intelligence process by machine”
- This definition is very broad, and includes different technologies originated independently, from *signal processing to pattern recognition, from logic programming to natural language processing*, hence we have to go deeper.

2. AI technologies

- AI technologies can be classified based on two independent features:
  - **models** (knowledge-based *(slide 5)*, classic learning *(slide 6)*, artificial neural networks *(slides 7-8-9)*)
  - **tasks** (supervised learning *(slide 10)*, unsupervised learning *(slide 11)*, reinforcement learning (*slide 12)*

Knowledge based (KB) systems

- KB systems rely on predefined knowledge
- Expert Systems are an example; they achieved commercial interest in the 1980s
- Expert Systems application examples:
  - Suggest a workstation configuration on the basis of users needs
  - Suggest whether to accept or reject a loan request
  - Suggest treatment of diseases based on patient characteristics
- Advantages
  - Results easy to explain
- Limitations
  - Generality. Custom solutions only.

Classic learning

- Rely on examples and on statistics / geometry-based algorithms
  - A simple algorithm: K-NN classifies a new sample on the basis of the majority among the K nearest classified samples.
- These systems were well consolidated since the ‘90s
- Advantages
  - Fair accuracy in many cases
  - Computational requirements typically not too high
  - Results easy to justify (as the K-NN algorithm)
- Limitations
  - Suitable for simpler tasks (e.g. classification)
  - Accuracy does not improve substantially by increasing examples

Artificial Neural Network (ANN) models

- Rely on **examples** and on **artificial neural networks** (ANN)
- ANN are made up of artificial neurons suitably arranged (typically by layers)
- An artificial neuron suitably combines different weighted inputs
- ANNs were initially proposed (in the 1950s) with the aim of simulating the functions of the animal brain. Constraints were then imposed for feasibility reasons
  - in any case today the functional difference between an ANN and a biological one is greater than the flight of a plane or the flight of a bird
- ANNs today can be considered **more advanced mathematical models** which:
  - capture deeper and more complex relationships than possible with classical methods
    - with an challenge to justify results *(explainable AI faces this challenge)*
  - improve their performance by increasing the number of examples
  - support *also* new functions, as generative AI, *if suitably configured*

Neural network generations (1)

- **Shallow**: up to 3 layers
  - Limited functionalities and limited abstraction capabilities
  - This was the status of the art until ‘90s, *not satisfactory due to training limitations*
  - Mostly mentioned for historical reasons, but still used in specific applications
- **Deep**: no more limitation about the number of layers
  - This breakthrough is due to a generalized training algorithm discovered in the 1980s and recognized in 2024 with the Nobel Prize (Hopfield, Hinton)
    - *Based on an clever application the Leibniz chain rule indeed!*
  - High abstraction capabilities demonstrated in the 2010s, typically with text and images, *including also best practices and heuristics identified by many researchers.*
  - Performances/features improve with the models size and with computational requirements
    - *The race for ever larger models begins*

Neural network generations (2)

- **Large Language Foundation models** (hot research topic in the 2020s !)
  - To solve the computational challenge associated to large models in ANN, these models are built in two stages:
    - The foundation model is trained to provide the **general knowledge**
      - This phase typically implies high computational efforts and costs
    - The application model is trained to provide the **application specific knowledge**
  - All this implies a new and “longer” structure of the value chain, in which only some big players can sustain the role of foundation model providers
    - *but this “fear” can be partially alleviated by recent developments (DeepSeek)*
  - New applications enabled by large language foundation models are fashionable, *as typically more horizontal (as ChatGPT)*, but *cover only some of AI use cases.*

Supervised learning

- the system is trained by examples. Any example include **parameters** (for example temperature, pressure) and **labels** (for example healthy, sick)
- when the system has been trained (learning phase), the system extracts its own model on the basis of which it can process new data presented to it.
- supervised learning requires that labels of the training database are defined by a human, and this represents a cost, also significant
- typical use cases solved by supervised learning are:
  - **classification** (to identify a class): based on the temperature and possibly other parameters detected, is the patient healthy or sick?
  - **regression** (to identify a value): based on the distance from the centre, pollution and other, what is the price of an apartment per square meter in a given area?

Unsupervised learning

- The system simply relies on **unlabeled dataset** (hence no more labeling costs)**,** and can **suggests new facts**
- In the simplest cases, the system, based on the proximity/distance of the parameters, identifies the most similar groups of data, anomalous data and like, as in:
  - **Clustering**: to identify similar groups of elements, as for example similar customers (targeted marketing) or similar patients (personalized medicine)
  - **Outliers identification** is a subclass of clustering, with 2 groups only: standard and anomalous behavior. E.g. credit card fraud identification.
- **Generating new content** is also a use case solved by unsupervised learning, if specific deep neural models are used
  - as for example to suggest new models of clothes compared to the old ones (it is up to the professional to select the most suitable one)

Reinforcement learning (RL)

- In RL the system does not use examples, but observes the environment and carries out actions, based on which it receives rewards or punishments.
- Based on rewards and punishments, the system learns the most appropriate strategies to use in a sequence of operations to achieve the intended result
- Reinforcement learning is used in autonomous systems, as:
  - Games (chess, go etc.)
  - Autonomous cars
  - Automatic trading
- More in general **planning** and **control** are typical use cases for reinforcement learning.

AI technologies: the proposed taxonomy

The proposed taxonomy

Includes models (rows)

and tasks (columns)

Tasks have been described

at the high level and as detail

(which is more precise).

Practical implementations

can also combine different

solutions

Task complexity increases →

Model complexity increases ↓

Task (High level)	**Supervised learning**	**Unsupervised learning**	**Reinforcement learning**
Task (detail)	**Classification, Regression**	**Clustering, Generation**	**Planning, Control**
Predefined knowledge			
Classic models			
Neural network models - shallow			
Neural network models - deep			
Neural network models - large			

3. AI applications: a different taxonomy

AI applications can be categorized according to two features:

A) target (rows in this table): *(Slide 15 for details)*

B) modality (columns in this table), i.e. kind of inputs

New horizontal use cases typically addressed by foundation models.

	Data	**Audio / Signal / Sequence**	**Text**	**Video/ images**
Vertical	Slide 16	Slide 16	Slide 17	Slide 18
Horizontal			Slide 19	

Target

- **Vertical applications**:
  - for health, manufacturing, finance, energy, chemistry, agro etc.
  - typically to address core business topics (process or product)
  - mostly relying on data/sequences and on established technologies
- **Horizontal applications**
  - for text (translation, summarization, question answering, content generation), video and images (generating)
  - mostly relying on Foundation Models
  - in most verticals they address specific department needs (e.g. marketing)
  - in some verticals they address the core business (e.g. publishing, ICT)

Use cases with data or sequences

- **Categorization** examples (categories known a-priori)
  - Is this patient healthy or not on the basis of the measured data?
    - Medical record, electrocardiograms (ECG), encephalograms (EEC)
  - Is this engine correctly working on the basis of parameters measured?
  - Is this credit card access or service access eventually suspect?
- **Clustering** examples (categories not known a-priori):
  - which are the clusters of my customers? To identify most suitable actions and campaigns for any cluster (personalized campaigns)
  - which are the clusters of patients in a specific disease? To identify the most appropriate actions for any cluster (personalized medicine)

Use cases for text

- **Categorization** examples (categories known a-priori)
  - for categorizing documents by predefined topics (finance, sports etc).
  - for distinguishing positive and negative reviews (sentiment analysis)
- **Clustering** examples (categories not known a-priori):
  - for clustering more similar documents
- more advanced applications as **translation, summarization, question answering, content generation** more easily solved by Generative AI

Use cases for images and video

- **Categorization** examples (categories known a-priori)
  - Quality control in manufacturing
  - Radioscopy control in healthcare
- **Regression** examples
  - Crop prediction from satellite images for agriculture
- **Generative AI** examples
  - Propose new clothing models in the fashion industry
  - Facilitate the generation of new images for brochures and papers
  - *Warning*: advanced generative AI can be used also for building realistic whilst deceiving pictures of celebrities to fabricate fake news.

The new wave of foundation models solutions

- Hot today, typically for **advanced text related use cases**:
  - Large spectrum use cases
    - Question answering (chatbots)
    - Search engines with summarization capabilities
  - More topic specific, whilst also broad
    - ICT: to support software developments
      - a new wave Software Engineering Tools
    - Legal: to draft legal documents, including patents
- Results of new tools improve with the user skills
    - Topic-specific
    - Tool-specific
      - with the emergence a new kind of specific skills, named “prompt engineering”

4. Opportunities and challenges

- Computational resources
- Data
- Applications

Computational resources

- Computational resources requirements typically increase with more recent technologies, as large language models. How to mitigate?
  - Use more computational intensive technologies only when necessary for features provided and/or accuracy
    - In most cases, less computationally intensive technologies are well suited to solve core issues in many verticals
  - Reduce computation requirements also in new advanced solutions (**Green AI**), which is an active research topic
    - The recent DeepSeek announcement is just an example
  - Use more efficient hardware architectures.
    - Decentralized architectures (**edge computing**) are another active research topic

Data

- **Good quality data** is a prerequisite for developing **any** AI solution
- **Data Science** is a prerequisite for having good quality data
  - Hence Data Science plus AI (they are not alternatives)
- **Open repositories** of data are definitely useful for first demo and trials
  - they are continuously improving for the variety and of quality data
- Possible challenges, to be instantiated by different modalities and verticals:
  - **Not so good quality data** available (format, provenance), which could require also specific *standardization efforts*
  - **Too less open data,** also due to business fears *(blockchain can be a solution)*
- **Large data sets** are a prerequisite for new advanced AI, which can be delivered only if large data set exist for a specific modality
  - Today this is mostly materializing for text related applications

Applications

- Opportunities
  - Many core business use cases in many verticals can be solved by using today well established AI technologies: *do not wait further*
  - These solutions can be developed by relying on open source libraries and affordable efforts
- Challenges
  - Application owners could be not aware that some problems of his/her interest can be realistically solved by using AI established technologies
  - Most of the AI debate today is on AI advanced research, whose aim is more on horizontal solutions, which can trigger FUD (fear, uncertainty and doubt) and which *can delay the timely adoption of well established AI technologies*

5. The AI book

- What is this book about?
- Why this book?
- Who is this book for?
- Main topics
- Book structure
- Book details and availability
- To summarize

What is this book about?

This book:

- provides an **agile overview** of the different Artificial Intelligence technologies
- summarizes **application** opportunities and examples in AI
- exemplifies how to develop specific **AI applications** by using open source libraries
- includes an extensive **bibliography** of articles and programs for further study

Why this book?

This book answers to questions:

- which are in general the **opportunities** of AI ?
- is my **problem** solvable by using the AI? Which are the specific benefits of using the AI ?
- which are the suggested **AI methods** to solve my problem ?
- how to **validate** the results I found by using AI ? how to eventually improve them ?

Who is this book for ?

This book's intended audience includes:

- **students** and **classes** entering this field, typically at the undergraduate level, interested to have an overview as well as implementation examples
- **professionals** interested to broaden their skills in AI with a quick learning curve
- **managers** interested in identifying the potential benefits of AI for them and to evaluate solutions available
- **citizens** interested to deepen their understanding in this more and more debated area with a transparent and simple approach.

Main topics

- **technologies** (supervised and unsupervised learning, classic and neural networks based methods)
  - overview in chap. 1, implementation examples from chap. 5 to 8
- **applications** (horizontal and vertical)
  - overview in chap. 1, details in chap. 9
- **data sets**
  - overview in chap. 3, data analysis in chap. 4 and appendix B
- **programming hints** in chap. 2 and appendix A

Implementation examples in detail

Reference language: Phyton

Open source libraries used:

a) scikit-learn (classic models)

b) Keras (neural network models)

Optimal coverage of most common use cases focused to simpler to explain and develop methods.

Task (High level)	**Supervised learning**	**Unsupervised learning**	**Reinforcement learning**
Task (detail)	**Classification, Regression**	**Clustering, Generation**	**Planning, Control**
Predefined knowledge			
Classic models	Chap.5	Chap.6	
Neural network models- shallow		(Cap.8)	
Neural network models - deep	Chap.7	Chap.8	
Neural network models - large	(Chap. 8)	(Chap.8)	

Book structure

This 289 pages book includes:

- many **text boxes** for summarizing key concepts
- more than 100 **figures** and tables to facilitate understanding
- more than 200 **references** for further study

Tables and figures: an example

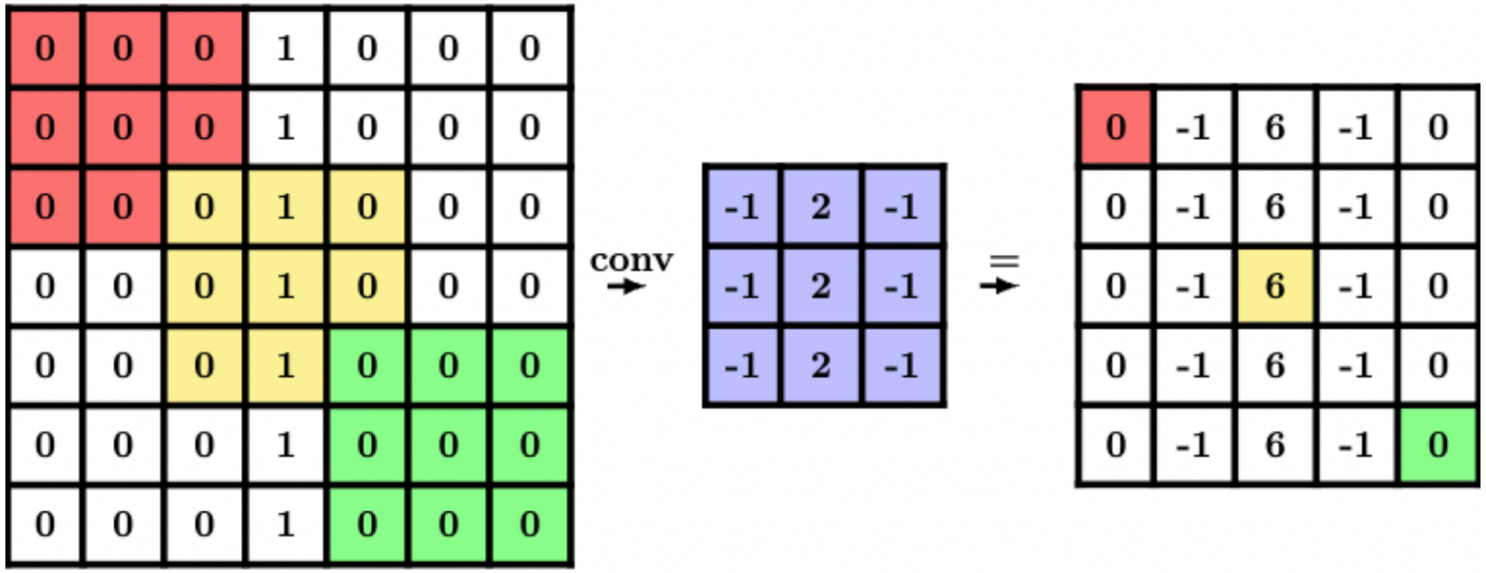

Figure 7.8: A full convolution example

Text boxes: an example

Important note

Generative AI is more and more attracting interest in corporate divisions, as:

- in customer operations, to improve today chat bot and IVR interactions
- in marketing, to facilitate the production of marketing documents
- in software engineering, to support developers in writing, testing and debugging sections of code
- in research and development, to facilitate the analysis of technical and scientific literature

Book details

Publisher: Intellisemantic Editions
(July 27, 2024)
Language: English Paperback: 289 pages
ISBN-10: 8894787605
ISBN-13: 978-8894787603
Item Weight: 1.59 pounds
Dimensions: 8 x 0.66 x 10 inches

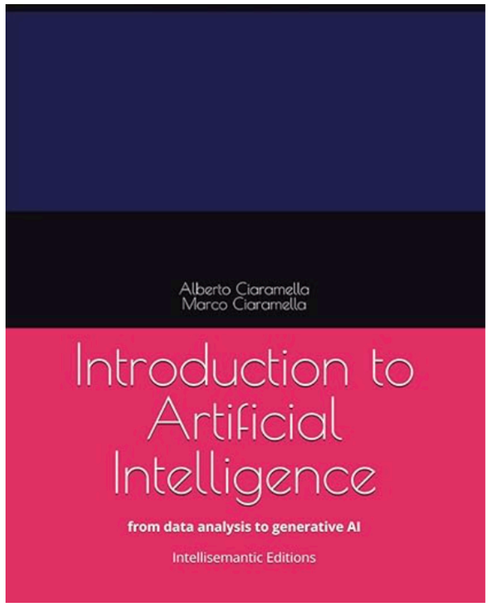

Availability (1)

- The book can be ordered worldwide on Amazon https://www.amazon.com/dp/8894787605
  - This site provides also access to sample free pages
  - Some Italian customers (teachers, students) in some cases are eligible for a price refund: please check (carta del docente, carta giovani, carta del merito)
- The book is distributed in different countries as:
  - Germany (Thalia, Osiander, Hugendubel)
  - USA/Canada (Barners & Noble, Magers & Quinn, Books Inc.)
  - UK (Blackwell's, Alibris)
  - Sweden/Finland/Norway (Adlibris, Akademibokhandeln, Bokus)
  - Poland (Allegro)
  - Perù (Buscalibre)
- Libraries (e.g. University libraries) can order this book at Casalini Libri (worldwide distributor)

Availability (2)

- The book has been presented at Bett UK (22-24 January 2025)
- The book will be presented at the London Book Fair (11-13 March 2025)
- For eventual other information contact contact directly ai_book@intellisemantic.com

To summarize

- The **AI design perspective:** this book provides a framework for different technologies and solutions to support better informed decisions about the technology to implement
- The **AI implementation perspective**: this book shortlists a good variety examples of not difficult to understand and implement with open source solutions in order to speed up the learning curve of a student or of a more seasoned programmer evolving to AI
  - The book presents only cursory examples of foundation models programming (open source of course), taking into account time limits of a first AI course
- This book does not detail foundation models from the user perspective (prompt engineering), which we consider out of the main focus of this book
  - we have been experimented indeed prompt engineering for Software Engineering, which is a different course we teach

6. Follow-up

Please contact IntelliSemantic at info@intellisemantic.com if you have questions about this presentation or if you want to know more about IntelliSemantic offers and activities, including:

- AI publishing
- AI and data analysis courses
- AI consultancy and feasibility studies
- AI projects, including Eu-funded research projects

See also https://www.intellisemantic.com

7. Questions and answers

10 minutes

Creative Commons license

This presentation is covered by CC BY-ND, hence:

- BY: credit must be given to the creator
- ND: No derivatives or adaptations of the work are permitted

More details https://creativecommons.org/licenses/by-nd/4.0/