



AgileAI

PAPER

Rev. 1.5.0 –April 11, 2024

AgileAI

an AgileConstellation Star



Summary

1. Overview.....	4
2. Artificial Intelligence	5
2.1 Main Areas of Application	7
2.2 The Issue of Data Privacy	8
2.3 Passive Agents and Intelligent Agents	9
3. AgileAI	12
3.1AgileAI Architecture	12
2.3.1 The Analysis Engine.....	12
2.3.2 The Learning Engine.....	13
2.3.3 The Strategic Engine.....	13
2.3.4 The Optimization Engine.....	14
2.3.5 The Interaction Engine	14
3.2 Implementation Toolit.....	15
3.3 Mindset	16
3.4 Specialization of Principles.....	17
3.5 Specialization of Practices.....	17
4. Conclusions	20



1. Overview

Like any self-respecting great change, *Artificial Intelligence* also has countless nuances and areas of application, some already consolidated and others yet to be discovered. For this reason, it is essential to start approaching this new technology with the right mindset (if we can limit ourselves to calling it such) which will certainly have enormous personal and professional impacts.

In particular, AI (or AI, *Artificial Intelligence*) will change the way organizations make strategic decisions, obtaining increasingly contextualized and detailed "suggestions" to support the People who will then have to make the choices. It is precisely in this perspective that **AgileAI** specializes: not a technology-oriented approach, but an ecosystem that aims to create artificial intelligence solutions at the service of organizations' ability to be increasingly protagonists in their markets. All with the awareness that AI will not replace the experience and creativity of professionals, but will be a valuable aid to extricate oneself from the ever-increasing density of information (or raw data) available.



2. Artificial Intelligence

The push to renew organizational models can "intelligently" exploit innovations in the field of artificial intelligence for a further implementation boost. All this passes through an intense automation of processes (and activities), particularly effective in all areas where decisions are closely linked to the data possessed.

The trend, also according to a study¹ by the *MIT Sloan Management Review*² (in collaboration with *The Boston Consulting Group*), is that of *cooperation between man and machine*, enhancing the latter's ability to collect and analyze large amounts of data in order to provide both operational and strategic "advice" and "suggestions".

The main advantages that can be obtained from this partnership concern: *better positioning of the company in the market, development of personalized marketing strategies, customer enhancement, loyalty, cost reduction and employee enhancement*. The level of achievable results depends on the organization's ability to innovate and renew itself, as well as the quality of the data possessed. This leads to distinguish four possible behavioral models of companies, in relation to the combination of artificial intelligence and market positioning strategies:

- **pioneers**, who have understood and adopted artificial intelligence to support flexible and adaptive organizational models;
- **experimenters**, who are investing in pilot projects to assess their impacts;
- **investigators**, who have understood the importance of this and are considering the possibility of activating specific trials;
- **who** have not taken any action in this regard.

Pioneers are obviously at an advantage in this new challenge, thanks also to a better understanding of the opportunities and technologies involved, as opposed to passive ones who struggle to understand their usefulness, also due to the absence of structured historical data, which is essential for training specific algorithms.

One of the underlying problems is the speed with which the related technologies evolve, putting at risk the investments of pioneers to the advantage of pursuers, or startups, which could turn the tables with a new innovative solution.

At a macro level, it is now possible to refer to artificial intelligence, starting from two macro-categories:

¹ <https://sloanreview.mit.edu/article/the-key-to-success-with-ai-is-human-machine-collaboration/>

² <https://sloanreview.mit.edu>



- **Weak Artificial Intelligence**, which identifies those systems capable of simulating some cognitive functions of humans without reaching the real intellectual abilities typical of the latter.
- **Strong Artificial Intelligence**, which refers to so-called "sentient systems", i.e. capable of developing real intelligence without emulating human-like thought processes or cognitive abilities.

Generally, we tend to talk about *Machine Learning* when we go to address the issue of how to make machines learn in relation to the execution of tasks, thanks to a series of algorithms capable of building *predictive models* and reducing the weight of *errors* at the end of each learning process.

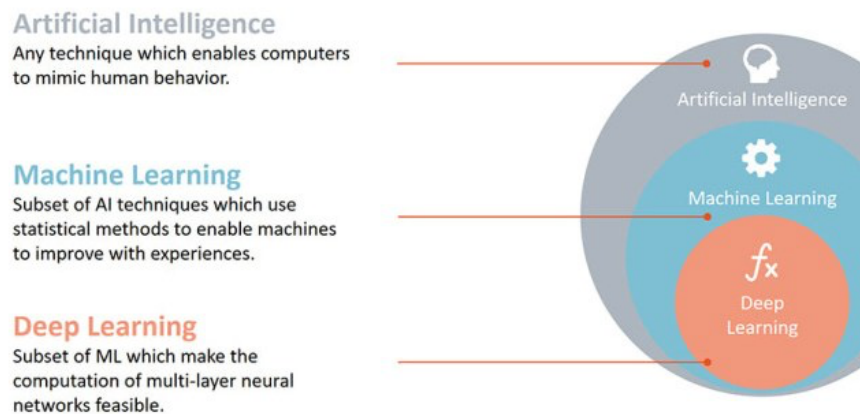


Figure 1 - Relationship between Artificial Intelligence, Machine Learning and Deep Learning

Based on the learning techniques, it is then possible to distinguish between different types of Machine Learning, which, using neural networks, make it possible to deal with specific tasks. The best known of these is certainly *Deep Learning*, but the faces of Machine Learning are much more multifaceted:

- *Deep Learning*: it is one of the most important areas of Machine Learning in which the brain's learning processes are simulated through *layered artificial neural networks*, therefore equipped with multiple H layers (2 or more), which allow information to be processed in a non-linear way in relation to the particular activation function chosen.
- *Model Prediction*: Includes a variety of techniques that can collect data and learn patterns to apply to it. Particularly used in the field of fraud prevention and qualitative analysis.
- *Online Learning*: focuses on *sequential* data (one after the other), where, as a result, decisions must be made as data becomes available.
- *Explainable regression & classification*: solves regression (prediction of ordinal quantities) and classification problems, allowing you to give an explanation of the results obtained.



- *Information Retrieval*: is the set of techniques used to manage the representation, storage, organization and access to objects containing information such as: documents, web pages, online catalogs and multimedia objects.
- *Reinforcement Learning*, its focus is on solving sequential decision-making problems through interaction with the environment in which it operates.

In particular, if Machine Learning can be defined as the method that "trains" A.I., Deep Learning is the one that allows you to "emulate the human mind", although due to the high computational cost attached, only recently have neural networks become reasonably approachable. This, in particular, thanks to new hardware architectures (especially the evolution of Graphics Processing Units - GPUs), which allow the processing of very large training sets³, typical, as mentioned, of Big Data.

It should be emphasized that Deep Learning is very sensitive to so-called "biases": in a supervised model, if the *labels* (categorizations) are created incorrectly, the model will learn incorrectly, and so will its inferences.

In general, when the objectives to be achieved are very specific and rigorous, deep learning techniques must necessarily be used, while for sufficiently elementary tasks, manageable by neural networks with a single hidden layer, the other declinations presented can be used.

2.1 Main Areas of Application

The main areas of application of artificial intelligence, according to a recent study by *osservatori.net*⁴ are:

- *Intelligent Data Processing, in which **specific data is analyzed to extract information and take actions accordingly***. This category includes several uses, such as predictive analytics and fraud detection.
- *Virtual Assistants/Chatbots, i.e . software agents capable of performing actions or providing services based on commands received in voice or text*. These systems, **increasingly used in corporate Customer Care** as the first level of customer care, stand out for their ability to understand the tone of the dialogue and memorize the information collected.
- *Recommendation System, which directs users' choices based on information provided by them (directly or indirectly)*. Among the most popular solutions are systems **that suggest**

³ In machine learning, a **training set** is a set of data that is used to train a supervised system (such as a neural network or probabilistic classifier).

⁴ https://blog.osservatori.net/it_it/applicazioni-intelligenza-artificiale



a purchase based on previous ones, thus influencing the *customer journey* and, more generally, the user's decision-making process.

- *Natural Language Processing (NLP)*, focused on the processing of information expressed in natural language, with purposes that can vary from **content comprehension**, to **translation**, up to the **production of text autonomously** starting from data or documents provided as input.
- *Computer Vision*, which focuses on the analysis of images, single or in sequence (video), for the **recognition of people, animals and things** present within the image itself, **biometric recognition** (e.g. face, iris) and, in general, the **extraction of information**. These systems are becoming more widespread in the field of video surveillance, where image analysis is essential to identify any anomalous or dangerous situations.
- *Physical solutions, which include:*
 - **Autonomous Vehicles, self-guided means of transport** for the transport of people, animals or things, which can be used for navigation not only by road, but also by sea, river and air.
 - **Intelligent Objects**, i.e. **objects capable of performing actions without human intervention** and making decisions based on the conditions of the surrounding environment (e.g. the smart suitcase which, through a Bluetooth connection, is able to identify the owner's position through a smartphone);
 - **Autonomous Robot**, a robot capable of moving without human intervention, based on information collected from the surrounding environment. This category includes both industrial solutions, such as robots **designed for the automation of production and logistics processes**, and robots intended for the civilian market, such as sales assistants in stores to provide information to customers.

2.2 The Issue of Data Privacy

The amount and diversity of data used in artificial intelligence means that specific *privacy and security requirements must be met*.

Different contexts may require specific approaches, and organizations try to adopt the least "stringent" solutions to achieve wide-ranging results, often generating "gray areas": think, for example, of video surveillance cameras and facial recognition.

Privacy laws are continually updated in order to keep up with new technologies and to ensure a high level of transparency for the end user. In particular, *personal data* is increasingly perceived as a "bargaining chip": consumers are willing to give it up in exchange for specific services.



Depending on the different behaviors, there are two distinct categories of users:

- *Registrars*, who try to be as careful as possible about the disclosure of their personal data.
- *libertines* (typical of younger people), who share their personal data very easily, without particular attention.

In addition to the privacy aspect, there is also an increasing focus on security, so much so that terms such as *cyber attack* and *cyber security* have now entered our daily jargon, as well as the related protection solutions to support them.

In particular, in 2016 the European Commission, in agreement with the national supervisory authorities, approved the reform of the rules related to the protection of personal data which goes by the name of *General Data Protection Regulation (GDPR)*. The aim of the reform is to harmonize the regulations on the protection of personal data within the European Union, protection that with the *Lisbon Treaty of 2007* has become a fundamental right of citizens, to be guaranteed and protected like any other right. With the new European regulation, there is a shift from a *proprietary view* of the data to a vision of *the control of the data*, which favors the free circulation of the same, provided that the interested party knows how it will be used.

This regulation, therefore, is the reason for a twofold change. On the one hand, it has a strong cultural value: in the digital age we are living in, the defense of data means defending the integrity and identity of the person. On the other hand, data protection also has a strong significance at the organizational level within a company. In fact, when you decide to launch a new product, a new service or even a new procedure, you must first evaluate the problems related to the security of personal data.

Another directive approved at European level in 2016 is the *NIS* (acronym for *Network and Information Security*), which requires European Union member states to adopt a series of common measures for the security of networks and information systems. In Italy, it entered into force in 2018. The main objective of the directive is to define a homogeneous strategic line in the various European States in order to prevent and combat the risk of accidents caused by computer networks and information systems. In concrete terms, the directive deals with the management of attacks by *cybercriminals*, protection against hackers, the detection of risks and the reduction of such incidents.

2.3 Passive Agents and Intelligent Agents

An *agile* organization is an organization that focuses on "*sociality*", with the operational dimension characterized by the many uncertainties deriving precisely from the variability of "human relationships".



All this requires appropriate tools capable of collecting process data, covering different areas characterizing the operations of a company: from the strategic to the detailed.

The tools available today are mostly characterized as "*passive agents*" because, while facilitating real-time information exchange and collaboration through centralized information, they are limited to supporting the creation, management and monitoring of project/program *artifacts* and the visualization of *historical* data through specific diagrams and reports. Therefore, there is a lack of advanced analytical methods, equipped with "intelligence" that allow the tool to make "reasoning" and provide specific "recommendations", all fundamental characteristics to be able to talk about "*intelligent agents*" (or even "*active agents*"). Intelligent agents are necessarily based on A.I. and allow you to support the organization in its primary aspects, thanks to the ability to analyze many of the management activities.

The efficiency and effectiveness of active agents is clearly linked to data and the ability to update and manage it appropriately: the more up-to-date it is, the more the agent will be able to provide reliable predictions and avoid costly errors during the implementation of different initiatives. In the same way, active agents are able to operationally support the implementation of projects, improving the use of related resources and developing a whole series of predictive actions.

A particularly interesting thing is that intelligent agents save a considerable part of the time spent by experts in formulating hypotheses and analyzing supporting information, intervening in the typical steps of defining a strategy:

- *Analysis*: The initiatives to be implemented originate from different sources, such as, for example, new explicit requests from customers, market changes, insights, etc. Processing all this data manually, especially with different perspectives, is very expensive and it is often not possible to consider all the related factors.
- *Review*: the initiatives to be pursued require continuous reviews and re-evaluations to be aligned with what is really happening and to consolidate them over time. All this requires being ready for the current conditions of internal operation of the company and the evolution of the market, which is obviously complex and difficult to achieve.
- *Planning*: Identifying which initiatives to implement, and especially which not to implement, is almost always a sort of "alchemy" that can put the company's very future at risk. This is a very challenging choice process since many factors need to be considered, often in dissonance with each other but with strong dependencies and not easily prioritized.
- *Monitoring*: As the initiative is developed, its progress must be monitored. This includes *risk management*, typically based on high-level guidance and often subjective judgments. This is a particularly difficult and costly task due to the inherent uncertainty, time dependencies and, above all, the dynamic nature of today's products and services.



It should be clear that a strategic approach to the integration of artificial intelligence into business processes cannot be separated from having an adequate set of intelligent agents, first and foremost *chatbots* to break down the barriers of human-machine interaction.



3. AgileAI

3.1 AgileAI Architecture

The **AgileAI Architecture** (*A-Cube*) represents a robust architectural proposal for the development of an integrated human-machine system, with particular reference to the aforementioned strategic choice steps: *analysis, review, planning* and *monitoring*.

A-Cube is divided into five main engines: *the analytics engine, the learning engine, the strategy engine, the optimization engine and the interaction engine*.

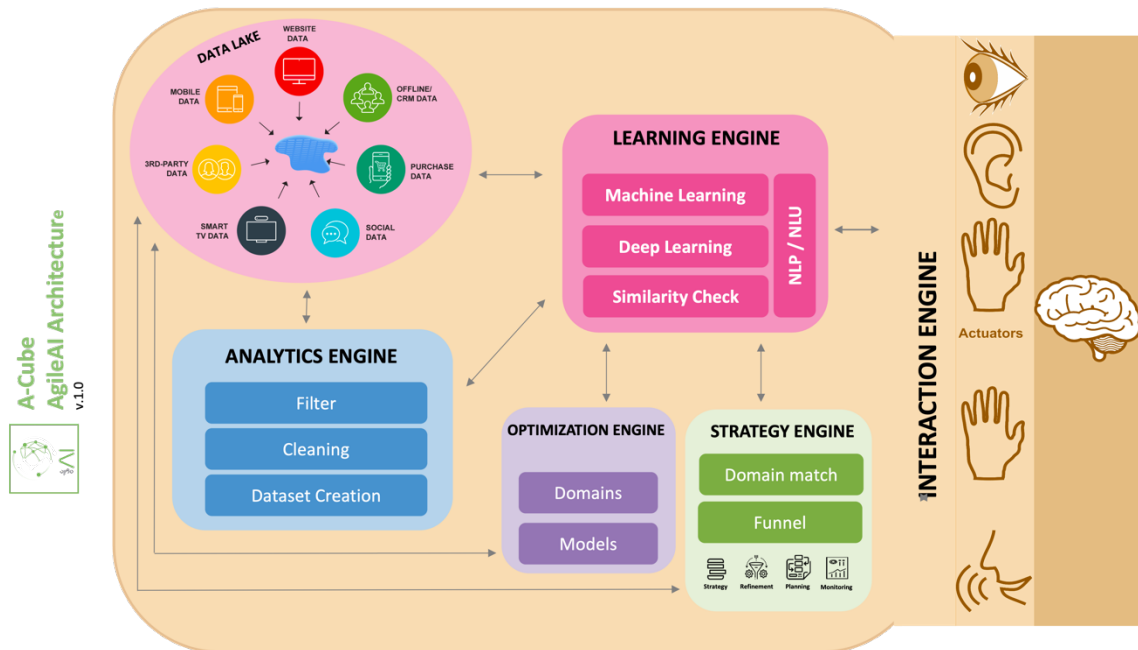


Figure 2 - A-Cube

The different engines are connected to a *Data Lake*⁵ that acts as a storage for Big Data, internal or external, that can be used by the organization for its decision-making processes.

2.3.1 The Analysis Engine

The purpose of the **Analytics Engine** is to query the data lake in search of the most significant data in relation to the specific objective. It has three primary components:

- *Filter*, which aims to filter the huge amount of data available, identifying those most relevant to the objective sought by those who are using the system.
- *Cleaning*, the selected data is cleaned and made consistent for aggregated use.
- *Dataset Creation*, aggregates are created that will be used by the learning engine.

⁵ https://it.wikipedia.org/wiki/Big_data



In particular, the creation of datasets takes into account the type of use for which they will be intended, structuring them according to the four most common types of analysis:

- *Descriptive Analysis*, which allows you to answer the question "*what happened?*". In practice, a summary of historical data is created for further analysis.
- *Diagnostic Analysis*, which allows you to answer the question "*why did this happen?*". The goal is to try to understand the causes of events and behaviors.
- *Predictive Analytics*, which allows you to answer the question "*what could happen in the future?*". In this case, we look for meaningful relationships between the variables and their representation in the available models.
- *Prescriptive analytics*, which allows you to answer the question "*how should we respond to potential future events?*". The goal is to support the decision-making process by putting the company in the best operating condition.

2.3.2 The Learning Engine

The **Learning Engine** is responsible for producing a series of potential options to be suggested as answers to questions. It uses algorithms and technologies related to *Machine Learning* and *Deep Learning* for its learning and deduction activities, in particular looking at the most modern NLP techniques, such as: *word2vec*, *paragraph2vec*, *Long Short-Term Memory* (used in Google Translate) or *Convolutional Neural Networks* (used in Facebook's DeepText engine).

This engine is also equipped with the *Similarity Check* module, which compares with known scenarios that have a high degree of similarity with the one under consideration. The goal is to provide *just-in-time* hypotheses that leverage what are the immediately identifiable and comparable "experiences".

2.3.3 The Strategic Engine

The **strategic engine** is responsible for refining the inferences made by the learning engine. This is done by making a comparison between the identified options and the related domain models. Its action starts from the options generated by the *learning engine*, with a strong correlation with respect to the information made available to the *optimization engine*. The inferred knowledge becomes, in turn, part of the information assets of the system itself, and therefore available for future queries and for the overall improvement of the overall analysis and prediction capabilities. The engine uses two primary components:

- *Domain match*, which calculates the optimal set of options most relevant to the specific domain in relation to a minimum evaluation threshold.
- *Funnel*, which classifies the candidate options by identifying the *Gold Option*, i.e. the option that most answers the question asked.



The strategic engine is trained to respond, particularly with respect to four strategic aspects: *strategy, refinement, planning and monitoring*.

2.3.4 The Optimization Engine

The **Optimization Engine** is responsible for managing the dynamics characterizing the context under analysis and the specific domains of relevance. This is a crucial aspect if we consider how the possible answers, in order to stop being "generic" and have a real operational resonance, must always be in tune with the specific context.

This engine makes use of two components:

- *Domains*, which is responsible for managing and providing information about the various domains that are significant to the organization.
- *Models*, which deals with modeling domains in relation to specific interest, narrowing the field of analysis (*models boundary*) and allowing to decrease the number of options that are actually significant from an informational point of view.

Models persisted in the data lake and updated whenever context evolves significantly.

2.3.5 The Interaction Engine

The **Interaction Engine** is responsible for managing the "dialogue" between the user and the artificial intelligence system, identifying the most appropriate form to allow the user to make their requests in the most natural way possible and receive, in an equally natural way, the identified option(s). As we have seen, one of the most common forms of interaction is that of *digital assistants* (chatbots in particular) that allow you to ask questions in natural language and receive answers just as fluidly.



3.2 Implementation Toolit

AgileAI proposes a structured implementation approach that is developed on 4 operational levels that define the *AgileAI Implementation Toolkit*:

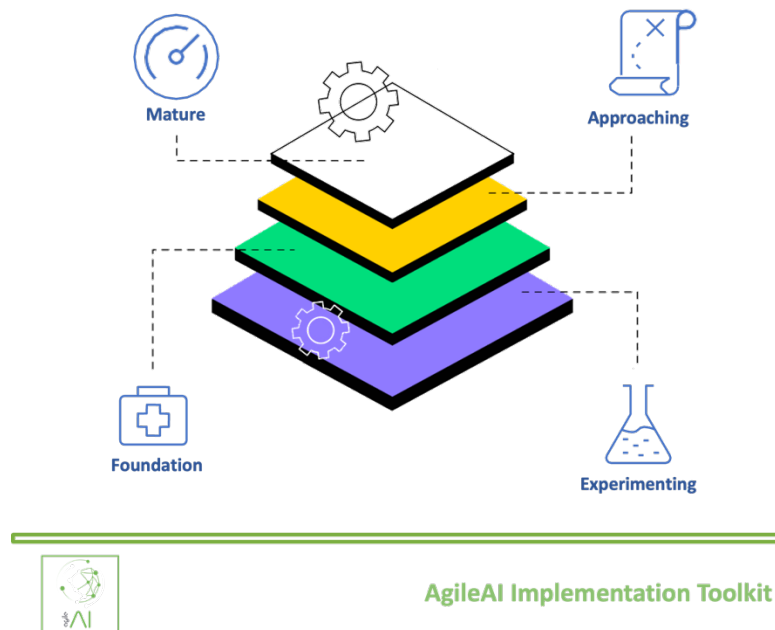


Figure 3 - AgileAI Implementation Toolkit

Specifically:

- **Foundation**, we are questioning the very meaning of A.I. in its context and how to take advantage of it:
 - difficulty in embracing the relative risks
 - Low digitization
 - Basic analytical skills
- **Approaching**, awareness and conviction of usefulness for the specific context:
 - Ongoing digitization
 - Automation and process optimization
 - Beware of indirect impacts
- **Experimenting**, experimentation and concrete application of the solutions to support:
 - Understanding the life cycles attached to predictive models
 - Implementation of a reference technology architecture
- **Mature**, the organization is an *Intelligent Learning Organization*:
 - High digitalization
 - Implementation of new "smart" business models

Each of these steps is a fundamental step in the path of structured integration of artificial intelligence, allowing you to develop the most suitable strategy for your company.



3.3 Mindset

At the heart of AgileAI we find the *Philosophy, Principles and Practices* of the **AgileConstellation Manifesto**⁶, of which AgileAI is a *Star*, defining specific practices and principles pertaining to the reference domain.

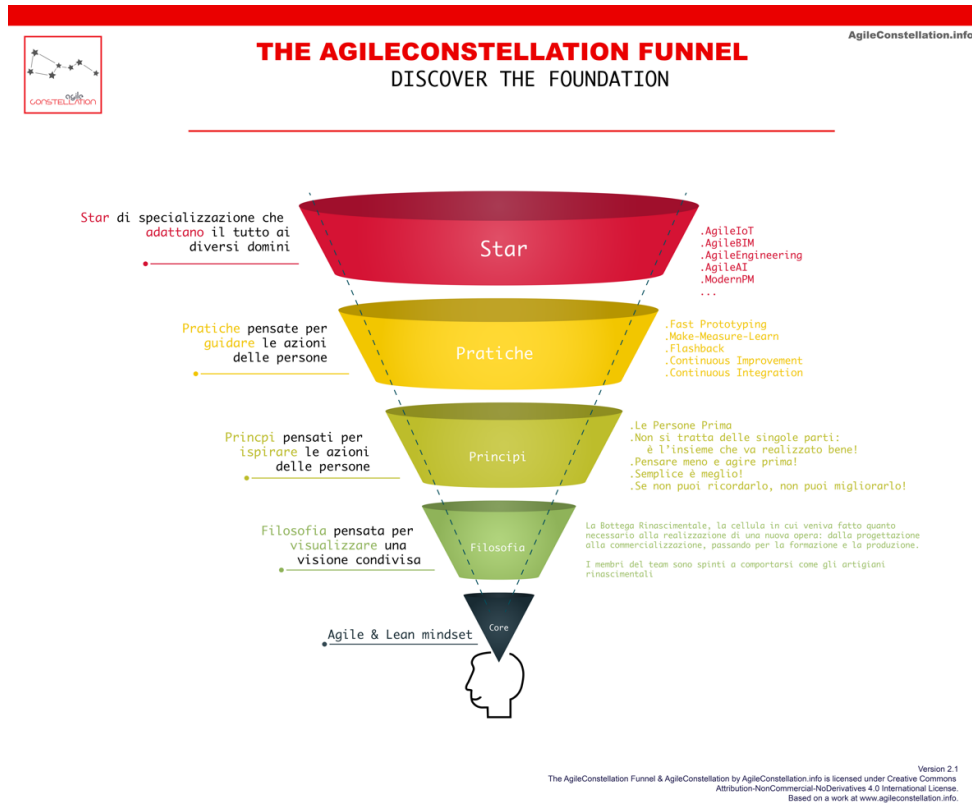


Figure 4 - AgileConstellation funnel

We have, therefore:

1. **Philosophy**, inspired by the **Renaissance Workshop**, or the cell that fulfills what is necessary for the creation of a new work: from design to construction and marketing.
2. **Principles (core)**:
 1. *It's not about the individual parts: it's the whole thing that needs to be done well!*
 2. *Think less and act sooner!*
 3. *Simple is better!*
 4. *If you can't remember it, you can't improve it!*
3. **Practices (core)**:
 1. *Fast Prototyping*, validating the sustainability of the solution
 2. *Make-Measure-Learn*, quickly experiment with different assumptions and assumptions
 3. *Flashback*, quick alignment in which the observer goes to the work desk

⁶ www.agileconstellation.info



4. *Continuous Improvement*, constantly improving every aspect
5. *Continuous Integration*, constantly integrating the different souls of the solution

3.4 Specialization of Principles

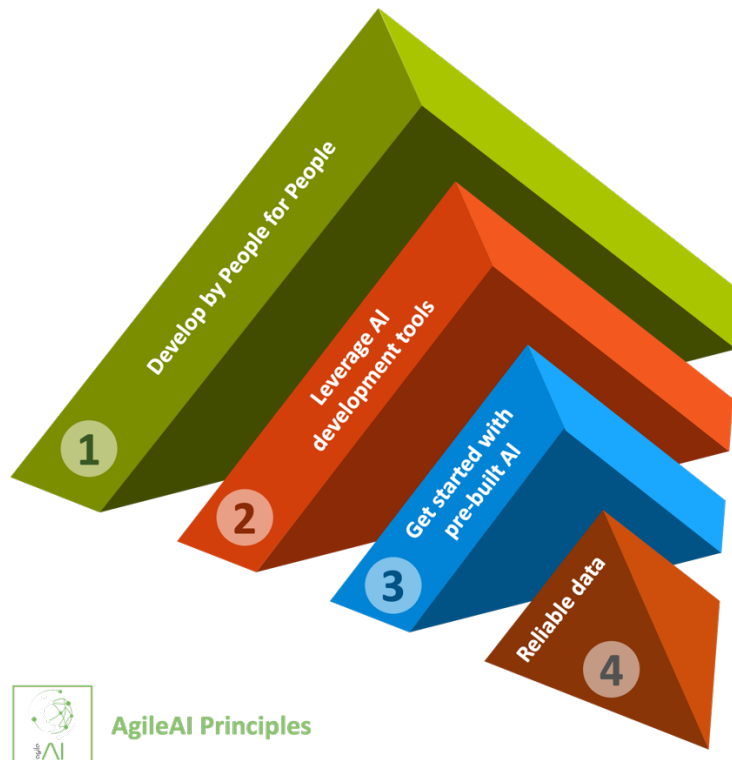


Figure 5 - AgileAI Principles

In addition to the principles "inherited" from the AgileConstellation funnel, AgileAI defines 4 additional principles that look specifically at aspects related to data and related analysis algorithms:

- **Develop by People for People**, develop any intelligence system for people, not to replace them in decisions.
- **Leverage AI development tools**, constantly improve what has been achieved, without fear of changing something that works.
- **Get started with pre-built AI**, before embarking on the creation of something new, analyze the market solutions, or the internal ones, available.
- **Reliable data**, only with reliable and updatable data can artificial intelligence systems provide valuable analysis.

3.5 Specialization of Practices

Compared to the inherited practices, the first specialization concerns *Fast Prototyping* which adds 5 new aspects (bubbles) of reference which, as a whole, go under the acronym *S.T.A.I.R.*:

- **Security**, protect information and ensure high levels of privacy



- design intelligent systems to anonymize data and maintain data integrity
- Protect your system from external attacks
- Conduct periodic security and privacy reviews
- **Transparency**, clearly supporting organizational decisions
 - Share key elements
 - Leverage comprehensible models and have intelligible explanations of the model's behavior
 - Train people on how to interpret suggestions
- **Accountability**, awareness of decision-making responsibility
 - Making it clear who is responsible
 - Make sure people are properly trained to use the results correctly
 - Keeping people at the heart of decision-making
- **Inclusiveness**, enhancement of different points of view
 - Human Experience, Automated Analytics
 - Attract a diverse pool of talent
 - research and use best practices, analytical techniques, and best tools
- **Reliability**, operating reliably, safely and consistently in all conditions
 - Review of support systems for anticipated and non-anticipated circumstances
 - Provide detailed explanations of how the system works
 - Appropriately report performance issues



Figure 6 - AgileAI Fast Prototyping



The second specialized practice is *Make-Measure-Learn*, which focuses on the effective implementation of AI solutions for business agility 6 operational steps:

1. *Minimum Viable Model (MVM)*, artificial intelligence is used to identify an organizational model (or practices) that can offer a convincing solution to the problem.
2. *Minimum Viable Product (MVP)* is the first implementation of MVM to assess its possible impacts. The implementation should be done in a "protected" context but representative of the specific organizational structure.
3. *Implementation*, which goes from the experimentation phase to a first concrete integration with existing processes. The goal is to identify the main critical issues of the model and take action in solving them.
4. *Data Flow*: the implementation data must arrive in the A.I. system suitably prepared, thanks to *data cleaning* and *dataset creation actions*.
5. *Production*, the new model is fully integrated into decision-making processes and moves on to its refinement, with particular attention to the impact on people.
6. *Refinement*, the model is continuously monitored, continuously updating it in relation to continuous changes to preserve and improve its effectiveness and efficiency.

These steps allow you to implement AgileAI, going through the different phases of the AgileAI Implementation toolkit, immediately providing concrete elements to support the most strategic and risky decision-making aspects, thus accelerating the organization's agility path.



4. Conclusions

As we hope it has emerged from reading the paper, the goal of its content is to develop a perspective on how artificial intelligence can impact organizations, defining a high-level approach for its implementation in its context.

If you are interested in the in-depth study and concrete application, you can contact us at the e-mail address info@agileconstellation.info or through our social channels that you can find on the official website.





AgileAI is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

The authors cannot be held responsible in any way for the use of what is reported in this document, in the annexed ones and in the related digital channels. The user has full responsibility for his own work and frees him from any type of direct and indirect task.

Based on available work on info@agileconstellation.info

Permissions in addition to the purposes of this license may be available from AgileConstellation.info

