ARTIFICIAL INTELLIGENCE

Artificial Intelligence systems function as automated and semi-automated decision-making algorithms. The different types of AI (generative, symbolic, analytical) share ontologies to facilitate their integration and use at the enterprise level.

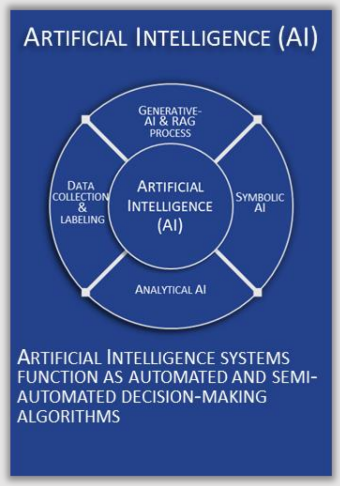

1. CONDITIONS OF SUCCESS

The interest in AI depends on the use cases of each company. Nevertheless, with broad application possibilities, significant gains are to be sought in all organizations. Indeed, AI covers a wide range of functionalities, such as:

- Creativity in communication and marketing, teaching, coaching, translation, text synthesis, report creation, financial optimization, customer tracking, trend calculations, pattern and video recognition, sound production, etc.

Beyond the specific case of a company, TRAIDA identifies two universal contributions of AI that do not depend on use cases. They form a strategic foundation so that stakeholders share certain fundamental objectives for the use of AI. Without this foundation, integrating AI into the organization encounters two riks:

- In the event of failure to implement AI in certain use cases, stakeholders may become demotivated. To counter this risk, it is important to have a framework that recalls the fundamental and shared objectives throughout the company.
- Poor implementation of AI leads to a misalignment with the company's fundamental objectives. Gains are then partially recognized by stakeholders. This context disrupts the organization and opens the debate towards questioning the profitability of AI. The strategic framework is necessary to counter this risk. It ensures that the contribution of AI for each use case aligns with the major objectives that bring stakeholders together.

To build this strategic foundation, the two universal contributions are as follows:

1. **Automate tasks**; that is, decision-making and the resulting actions.
2. **Accumulate and exploit knowledge**; in order to better control the organization.

These two contributions are identified by the majority of AI experts, but their formulation in the specific context of each company remains to be done. Indeed, automation is intimidating and requires an explanation to situate it within a framework of overall activity improvement. Similarly, knowledge management has been a recurring theme for decades, without much motivation. However, with AI, it becomes strategic and profitable.

By formalizing the two universal contributions of AI in terms that suit your company, you build your strategic AI framework. This is a document of a few pages, a sort of charter on the fundamental objectives of the company with AI.

To guide you in drafting this framework, the two universal contributions are detailed in the following section.

Automating tasks

Since its inception, computing has had the fundamental objective of automating tasks; AI is part of this trend. According to consulting firm McKinsey (2024), 70% of tasks could see a 50% productivity gain thanks to AI. Thus, it is not about 100% automation but about entrusting certain steps in work processes to AI.

This collaboration between humans and AI has consequences on the behavior of work processes. Four levels of analysis and execution must be distinguished: the design of the new process, the rules for triggering AI, its execution, and finally, its control:

1. **Design of the new process**: This involves determining the steps to entrust to AI in work processes. For each of these steps, a knowledge base is gathered to train the AI. Its execution must be consistent with the objectives of the step. This knowledge base consists of labeled or unlabeled data sets and documents already available in the company or newly created.
2. **Rules for triggering AI**: Before the AI-executed step is performed, the system checks that the conditions for its triggering are met. There is thus a context in which the AI must act; if this is not verified, the AI could be triggered incorrectly. In conventional programming logic, this preliminary check forms preconditions. They can be developed traditionally (algorithms) or rely on AI specialized in context management (symbolic AI, analytical AI, context learning, etc.).
3. **Execution**: Once the preconditions are verified, the AI executes the process step, which consists of two parts:

   a. Decision-making to build an action plan;

   b. Execution of this plan.

During process execution, points of collaboration with a human operator are possible, for example, in the form of a communication space in a workflow messenger. Execution under AI's responsibility can also act in the physical world with robot commands. The action plan constructed by AI then orchestrates its own executions, as well as those of human operators and robots.

4. **Control**: A fourth step ensures that the actions taken by AI meet the company's expectations. Two complementary intervention levels are provided:

   a. A verification that the AI's response conforms to an expected context at the process step's exit (guardrail). In conventional programming, these are postconditions. Just like preconditions, they can be developed in classical language or rely on AI. Barriers against hallucinations can also be provided by leveraging LLM's RAG (Retrieval Augmented Generation) technologies.

   b. Overall supervision of process behavior using AI dedicated to observing the information system. This AI is trained with the process specifications and all regulatory documentation. All process executions feed a knowledge base that serves as a source of information for analyzing this general AI (also called Trusted-AI, Second Brain, Nerve Center, etc.).

Accumulating knowledge

Companies manage an increasing amount of information. According to the observations of many experts(*), it is likely that only 30% to 40% of knowledge is recorded in databases. The rest, i.e., 60% to 70%, forms tacit or informal knowledge that is exclusively present in the minds of human actors.

(*) https://link.springer.com/article/10.1007/s12144-023-04994-3

Knowledge management is not a new topic for companies. It is an essential objective of computing. However, to go beyond databases and multimedia storage spaces, it is necessary to focus on transforming tacit knowledge into digital contents. Unfortunately, projects that encourage actors to document their practices in writing encounter this question:

- "*Why devote efforts to formalizing documentary repositories if they cannot be automatically exploited to improve processes?*".

In other words, once the know-how and practices are written down, it remains indispensable to employ consultants to analyze them in order to draw conclusions. This work is long, costly, often approximate, and not very compatible with the regular and necessary updates of the repositories.

Ultimately, many companies abandon overly ambitious projects for formalizing tacit knowledge. However, with AI, the potential for this formalization changes for two reasons:

1. The training of AIs is all the better when there is a formalization of tacit knowledge. It allows them to better understand the organization. In other words, AI systems that operate on 30% to 40% of the knowledge available in databases are much less efficient than those that also absorb tacit knowledge. The potential for process automation increases exponentially with each additional formalization of tacit knowledge.
2. AI enables the automatic exploitation of knowledge formalized in writing. It thus substitutes for consultants to accelerate the transition from writing to action. The barrier of the lack of profitability of documentation initiatives is then lifted.

In this context, actors must improve the quality of their writing to better formalize their work practices. For its part, the organization must explain the objectives and career plans with AI to maintain maximum trust in knowledge sharing.

The strategic foundation for profitability

The two universal contributions we have just described are fundamental for increasing the profitability of AI, particularly for its use at the enterprise level. Indeed, beyond a quick gain from carefully selected initial use cases, deploying AI across all processes of the organization naturally encounters difficulties and obstacles. In other words, you need to be precise about your fundamental objectives with AI:

- **Design work processes by entrusting certain steps to AI**: Implement a method for designing processes with preconditions, postconditions, training AI for action plan calculation, and then executing these plans collaboratively with humans and robots, and using AI for supervising the behavior of the information system.
- **Encourage teams to formalize their know-how and practices in writing**: Implement a method to improve writing quality and the automatic exploitation of this new knowledge by AI.

In addition to these two fundamental objectives, it should be noted that AI systems need high-quality data to function properly. TRAIDA is based on an architecture centered around the semantic platform with the MDM, ODS, and EKG repositories. They form the third fundamental objective in the technical field (see the respective TRAIDA cards).

In the continuation of this card's description, we detail the categories of AI to provide a better understanding of the solutions offered by the market. Although this general technical culture is useful, the management of the AI approach relies mainly on the strategic framework we have just described. Its implementation does not depend on technical mastery. To successfully scale AI in your company, you need to engage your stakeholders on this framework to build clear and widely shared commitment.

2. IMPORTANCE OF THIS CARD FOR YOUR TRANSFORMATIVE AI

The use of AI at the enterprise level involves several types of technology. The boundaries between these types evolve with the state of the art. Given the strength of innovations, they could merge into a universal user solution. This is referred to as Artificial General Intelligence (AGI), which is predicted to arrive within a decade.

This potential for convergence reinforces the idea that companies should be interested in all types of AI, not just those highlighted by current technological trends. The effect of compounded interests should be sought by combining AIs together, as well as use cases that rely on different AIs.

To take this path toward unification, it is essential to rely on semantic data management that promotes the sharing of ontologies at the enterprise level; that is, by all use cases that rely on AI. It should be noted that an ontology is first a catalog of business concepts that the company uses to execute its activity. Then, these business concepts are organized in a hierarchy to describe their classifications and specializations in the sense of the object-oriented approach. Finally, the ontology itself emerges with the development of relationships between business concepts. This semantic modeling is the foundation for providing information to AIs and accumulating knowledge.

We now describe the different types of AI according to the current state of the art.

GENERATIVE AI & RAG PROCESS

This AI has become popular with the general public through ChatGPT by OpenAI. Since then, several similar solutions have emerged, such as Meta's LLAMA or Anthropic's Claude.

These are LLM (Large Language Model) AIs that generate text from prompts and knowledge bases. Contrary to the common belief that this AI merely predicts words to complete texts, deeper practice suggests that its functioning is more subtle. Indeed, beyond text translation and synthesis, it is possible to engage in a structured and intelligent conversation with the AI to receive help with goals that have nothing to do with text completion.

For teams within the company, it is important to experience this level of intelligence, but to achieve this, it is necessary to formulate prompts and enhance the AI's intelligence with your own knowledge bases. The AI then becomes a sort of piano, the use of which varies greatly depending on the musician's mastery. LLM is different from the application that use it as the core. For LLM it is just to do text completion - autoregressive text prediction. This kind of language model is enough to solve any natural language processing problem. The application like ChatGPT of course will do more things than just the language model (text completion), for example RLHF (reinforcement learning from human feedback), various NLP tasks, prompt analysis and guardrail.

At this point, the question of intelligence must be addressed. A lack of preparation to answer it leaves the field open to criticism from users resistant to change. Let's take the example of this negative assertion:

- "*No AI can replace human intelligence because it only recycles existing information of varying quality*."

Within the TRAIDA framework, we respond as follows:

- Intelligence consists of recycling existing knowledge to assemble and modify it so that new information emerges. The creation of knowledge ex nihilo is no longer merely intelligence but genius. At this stage, such a level of contribution is not expected within the company.
- The more powerful the learning, the more intelligence is unleashed. This power relies on two pillars for the company: the formalization of tacit knowledge and the high quality of data repositories. In other words, the perceived intelligence level of available AIs within the company reflects its own intelligence level in formalizing tacit knowledge and mastering data (ontologies). For example, using ChatGPT in its raw version does not know your company's data, so its intelligence level is close to zero for helping you in your management processes. You will need to provide it with knowledge about your activity to obtain relevant results.

This remark also applies to use cases that seem elementary, such as using ChatGPT for text translation. Indeed, two approaches are possible:

- Basic: a simple prompt with a copy-paste of the text to be translated.
- Intelligent: creation of an AI assistant to train it on a translation style, with a glossary of important terms, a directive requiring the AI not to add, change, or remove ideas, and finally, a request to add an explanatory note for the most delicate points of the translation. Moreover, this AI assistant would

be positioned as an expert in translation in the target language and an expert in the subject area of the texts to be translated (e.g., an expert in Greco-Roman history).

The translation quality then differs significantly. The first approach is even risky because the AI could engage in unnecessary hallucinations. The second approach guides the AI much better and improves the quality of the final result.

Thus, when choosing an LLM solution, it is important to study the creation of AI assistants. It is not enough to let your users create prompts; you need to organize a catalog of AI assistants so that each of them can enhance their skills in their respective fields.

Retrieval Augmented Generation (RAG)

AI Training is Conducted at Two Levels:

1. **General**: In batch mode, on large quantities of information through the loading of datasets and documents. The freshness level of the information used must be consistent with the batch mode. For example, if training is to be conducted every month, information with a lifespan of just one day should not be considered. Given the operational and financial cost of training AIs, batch mode is indispensable.
2. **Augmented**: In real-time mode, by retrieving the most up-to-date data possible from databases. This is where RAG (Retrieval-Augmented Generation) mode comes into play. It involves a technical decoupling to enrich the information flow sent to the AI (via the prompt or API activation) and control the information flow returned by the AI, for instance, to identify hallucinations. The operational and financial cost is indexed to the number of requests made.

Examples of solutions: Haystack (deepset), Replika.

Multimodal generation

Generative AI also works for formats other than text, such as images, photos, and more recently videos, as well as sounds and music. The application domains of multimodal AI are very broad, ranging from marketing content production to the creation of cinematic works, comics, music, etc.

The training process for these AIs is less within the control of companies compared to text generation. The large volumes of photos, videos, and sounds used for this training are managed by solution providers.

Nevertheless, a company with a significant repository of multimedia data can also train its own assistants to achieve AI personalization in its context. For example, a repository of users conversations (audio files) from a call center in a very specific language can be used to train a speech recognition AI.

Knowledge governance

The data used to train AIs must be archived for two reasons:

1. To keep a record of the information used in case of an audit.
2. To be able to duplicate this archive and update certain data in order to reset the AI assistant.

This second point is essential for unlearning outdated or incorrectly loaded data in the AI.

SYMBOLIC AI

Symbolic AI relies on formal rules for decision-making calculations. Unlike generative AI, the learning process is reduced to creating these rules, then organizing them into bundles, lineages, and inference logics. The rules receive input data (facts) and execute within the framework of a context of data shared across multiple executions.

Technical solutions for symbolic AI are known as expert systems or rule engines. They are useful when the knowledge domain can be formally described by rules. They are also used to increase the level of abstraction of certain programming logics, for example, for implementing preconditions and postconditions that govern the execution of services or steps in processes.

Unlike generative AI, the results of symbolic AI are auditable since it is sufficient to list the rules that were traversed to obtain an answer. Moreover, the deterministic functioning eliminates any risk of hallucination, though not of bugs if the rules are poorly organized.

The impact of this type of AI is more technical than operational. Indeed, they face a complex engineering barrier in rule modeling. Thus, it is a slower and more limited integration mode compared to what is possible with generative AI.

It is entirely possible to train a generative AI with rules described in natural language. However, with a large volume of rules (several hundred), maintenance difficulties are likely to arise with hallucination effects that are difficult to correct. When the number of rules is fewer than ten and the documentation format is clear and unambiguous, the LLM can act as a tactical rule manager.

Examples of solutions: Prolog, Drools.

ANALYTICAL AI

Analytical AI aims to predict outcomes from datasets. For example, in the medical field, a file that references diseases per patient along with associated symptoms is a dataset for training an analytical AI. Once training is successful, the symptoms of a new patient can be submitted to this AI to diagnose the potential disease. Similarly, in the healthcare domain, the principle of analyzing radiological or MRI images with AI is based on the same type of training, i.e., with a set of images that provide the disease result.

This type of AI uses machine learning technology, with numerous solutions whose performances vary depending on the targeted use cases. Their operation relies on statistical analysis, which remains more deterministic than that of LLMs. Thus, generative AI can also be used with the perspective of analytical AI but for limited datasets to avoid hallucinations. For example, based on a dozen criteria from a quote request, it is possible to determine the email response to send to prospects. It is sufficient to provide the generative AI with a few dozen examples of quote analysis along with the reference email used to achieve the desired automation. The complexity of cases here is much lower and less sensitive than for the medical domain. LLMs are fundamentally statistical in nature as well. The primary distinction lies in the sheer scale of LLMs compared to traditional statistical machine learning models.

Examples of solutions: Scikit-Learn, TensorFlow, RapidMiner.

DATA COLLECTION & LABELING

This final topic does not address a specific type of AI but rather the labeling of data necessary for learning processes. This labeling occurs at several levels:

- During the collection of tacit knowledge to attach it to ontologies. This information capture can be directly integrated into production applications. For example, an insurance operator analyzing a claim file has a button on the UI to enter a text explaining how they conduct their analysis. This knowledge is then linked to the type of claim being managed (ontology). A quality control process is initiated so that the AI manager for the claims department can assess the relevance of this new knowledge and accumulate it for future AI training enrichment.
- During the preparation of datasets, for example, to indicate the location of a problem on each photo of an electronic board with a manufacturing defect. A platform is needed to manage this enrichment process.
- During the evaluation of AI responses to improve future results. For example, the image generation solution MidJourney operates within Discord. It offers the user a series of four images to choose the best one. The user's response is given naturally in the discussion thread. By making this choice, the user helps improve the AI. This evaluation principle can be replicated in production applications. For instance, at the end of a process executed with AI, the system proposes a quiz to the user to

evaluate the positive and negative points of the work. This is similar to the performance evaluation of an employee but applied to AI within the execution of a process.

These data labeling mechanisms prepare the training and allow for the evaluation of results to continuously improve the AI (human in the loop). They should be integrated as closely as possible to applications and process execution.

Examples of solutions : Labelbox, Amazon SageMaker Ground Truth, Prodigy, SuperAnnotate.

3. BLUEPRINT

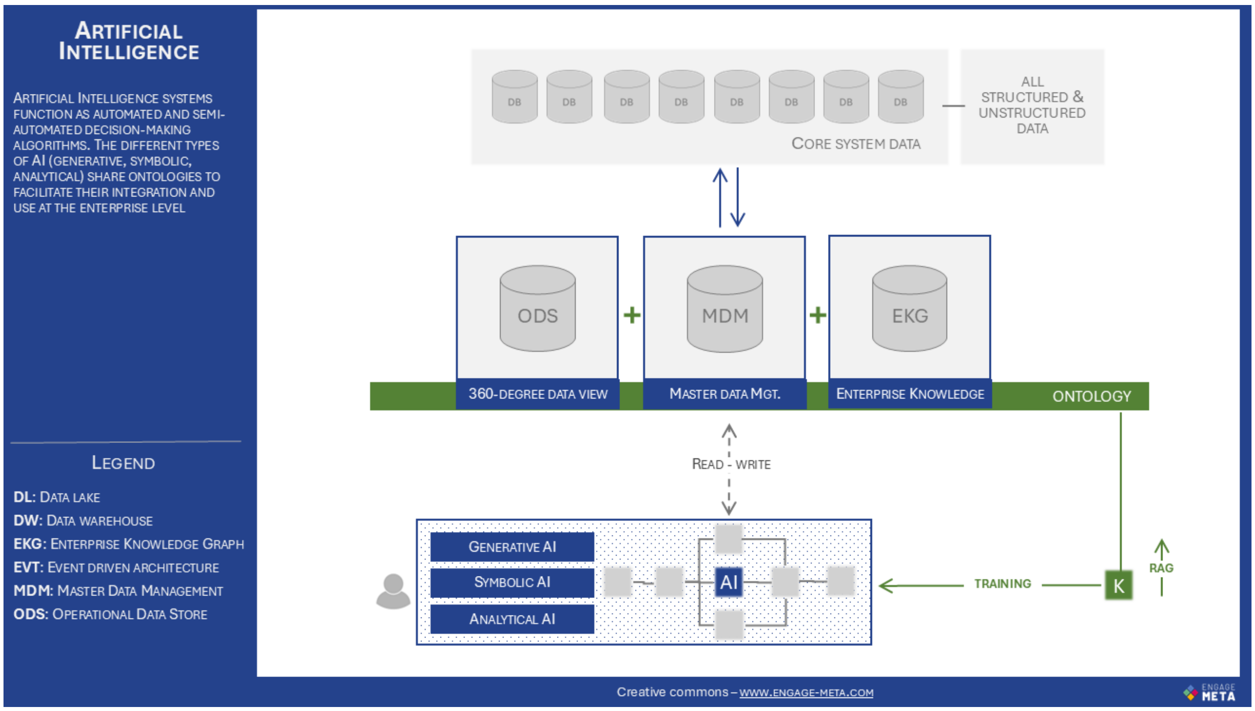

4. YOUR SITUATION & OBJECTIVES