

# Artificial Intelligence

Artificial intelligence (AI) is problem solving realized by means of artifacts. It usually involves applications such as:

- advanced (web) search engines and recommendation systems
- understanding human speech and generative or creative tools
- automated decision-making and strategic gaming

AI arguably covers four general categories: representation, reasoning, learning and search; language understanding being a special case of representation, planning a special case of reasoning and table lookup a special case of search. Machine learning methods are useful on large problems, which is becoming increasingly important as applications such as language understanding are moving into real-world situations outside the lab<sup>1</sup>.

AI methods can also be used to solve fierce combinatorial problems<sup>2</sup>. These are problems that can be represented as a search tree using nodes representing entities. Current (non AI) methods work well on small problems but often fail when applied to larger real-world problems because there are too many options in the search trees that must be explored<sup>3</sup>. Where the number of potential options cannot be kept manageable, solving problems amounts to finding different solution strategies that work. This requires methods for searching through very large trees, possibly involving ideas from machine learning such as neural networks or evolutionary algorithms. More generally, machine learning methods can be applied to virtually any problem where the solution space is finite (finding a path through a graph).

More abstract AI systems evolve towards cognitive systems<sup>4</sup> that are able to independently develop strategies for solving problems. To this end, systems are equipped with capabilities for context comprehension, interaction, adaptation, adjunction and learning. Cognitive systems realize areas or spaces with interconnected items of knowledge and representations of problem solving processes<sup>5</sup>. Cognitive systems are complex systems providing computational representations of processes to augment computational capacities<sup>6</sup>.

---

1 Calling for a better understanding of why certain learning algorithms work well on some types of problems, but not others. Furthermore, there is a need for a better grasp on issues such as learning from large databases, learning in multiple domains and learning task-specific knowledge.

2 Combinatorics involves enumeration (counting) of specified structures (arrangements or configurations), associated with finite systems (or infinitely countable with discrete setting). It is about structures that satisfy certain given criteria and construction of these structures, in order to find "best" solutions among several possibilities.

3 Searching for solutions to a problem among all possible alternatives is an important capability but one that is increasingly difficult due to its complexity. Brute-force searches require enumerating all alternatives, which often proves problematic if not impossible even on simple problems, whereas other approaches are so specialized that they have little value outside their specific domain (and sometimes not even there). Brute-force approaches have been successfully applied to optimization problems where only a limited variety of desirable solutions are available, but there are many applications that require solving very large problems with potentially countless solutions. In the end, any feasible algorithm will require shortcuts that often involve approximations or heuristics.

4 The term is used to indicate systems mimicking functions associated with intelligence, such as learning and problem solving. Non-trivial cognitive systems are capable of learning and adapting to changing environment. High-level cognitive systems may show various degrees of intelligence.

5 Such as the Semantic Web.

6 Such as Knowledge Graphs, Ontologies and Blackboards.

## Semantic Web

The goal of the Semantic Web is to make Internet data machine-readable (*FAIR*). Advances in ICT and AI (knowledge graphs, property graphs, deep learning, automated knowledge base construction, language models, computer vision and multi modality) have resulted in:

- machine driven / automated *FAIR* of semantically rich datasets
- knowledge modeling<sup>7</sup> and reasoning<sup>8</sup> efforts crystallizing into ontologies<sup>9</sup>
- controlled vocabularies and conceptual models such as CIDOC-CRM, the Europeana Data Model, and FRBRoo, with numerous 5-star Linked Data datasets populating the Linked Open Data cloud

Semantic web technologies offer promising ways to deal with information that is subjective, vague, fragmentary, uncertain and contradictory indeed.

Currently, there are many interlinked sources such as DBpedia<sup>10</sup>, GeoNames<sup>11</sup>, Wikidata<sup>12</sup>, GRID<sup>13</sup>, KnowWhereGraph<sup>14</sup> and Open Icecat<sup>15</sup>, and their number continues to grow every day.

However, the real added value (utility) of knowledge graphs emerges when data in local domains (representing proprietary knowledge) are “adapted” and then “connected” to the common domain (the open global knowledge; Linked Data<sup>16</sup>).

Another important adding value feature of graph databases is inferencing capability, where new knowledge is being realized from facts. Facts, when properly embedded, make search results more ‘applicable’, opening new avenues for ‘actionable’ (useful) insights. Moreover text mining techniques extracting new facts from free-flowing texts may embed data directly in graph databases. Conversely, modern text analysis technology makes considerable use of large knowledge graphs providing background knowledge, human-like concept and entity awareness and more accurate interpretation of texts.

---

7 Knowledge representations aim at representing information about the world in forms computational systems use to solving complex tasks or engaging in natural language dialog.

8 Automated reasoning systems (inference engines, theorem provers, classifiers) involve problem solving, representation scheme formalisms (semantic nets, systems architecture, frames, rules and ontologies) and logic, to automate various kinds of reasoning, such as application of rules or relations of sets and subsets.

9 Ontologies realize representations, formal naming and definition of categories, properties and relations between concepts, data and entities, substantiating one, many or all domains of discourse. More simply, ontologies are a way of showing properties of subject areas and how these are related, by defining sets of concepts and categories that represent the subjects.

10 The DBpedia Association has built up an extensive network of international members comprising representatives from industry and academia. All DBpedia members support and encourage DBpedia and actively use the Open Knowledge Graph to enrich their datasets.

11 The GeoNames geographical database covers all countries and contains over eleven million place names that are available for download free of charge.

12 Wikidata is a free, collaborative, multilingual, secondary database, collecting structured data to provide support for Wikipedia, Wikimedia Commons, the other wikis of the Wikimedia movement and to anyone in the world.

13 Global Research Identifier Database is an international database of institutions engaged in academic research.

14 KnowWhereGraph is an integrated 12 billion triples strong knowledge graph on 30 data layers at the intersection between humans and their environment using Semantic Web and Linked Data technologies.

15 Open Icecat is a multilingual open catalog containing product data sheets, related digital assets and usage statistics.

16 Linked Data is structured data supporting semantic queries. It builds upon standard Web technologies. The idea behind Linked Data is for the Internet to become a global database.

## ICT Technologies

**Knowledge graphs** represent networks of ‘real-world’ objects and illustrate the relationships between them. Data are usually stored in graph databases and visualized as graph structures, prompting the term knowledge “graph”. Knowledge graphs are made up of three main components: nodes, edges and labels. Any term can be a node. Edges define relationships between nodes.

In knowledge representation and reasoning, knowledge graphs are knowledge bases that use graph-structured data models or topologies to embed data. Knowledge graphs are used to store interlinked descriptions of entities – objects, events, situations or abstract concepts – while also encoding the semantics (usage) underlying the terminology used. Knowledge graphs are associated with linked open data projects in the frame of the Semantic Web.

**Ontologies** represent the backbone of the formal semantics of a knowledge graph. They can be seen as data schemata of graphs. More precisely, knowledge graphs may include ontologies allowing agents (both analogue and digital) to understand and reason about their contents.

Knowledge graphs formally represent semantics by describing entities and their relationships using ontologies as a schema layer. In so doing, knowledge graphs allow logical inferencing to retrieve implicit knowledge rather than merely permitting queries requesting explicit knowledge.

The **blackboard** architectural model consists of a common knowledge base, the blackboard, that is systematically updated by a diverse group of specialist knowledge sources, starting with a problem specification and ending with a solution. Knowledge sources update the blackboard with (partial) solutions when internal constraints match certain blackboard states. In this way, the specialist knowledge sources are said to work together to solve a problem.

The blackboard model is specifically designed as a way to handle complex, ill-defined problems, where the solution is the sum of its parts, needs not necessarily be fully or even sufficiently correct and is not per se time critical<sup>17</sup>.

---

<sup>17</sup> More is better, mistakes are not expensive and there is plenty of time.