Selected AI categories and terms

AI techniques

AI techniques are different core algorithmic approaches used to implement AI functions. These techniques are defined below.

Bio-inspired approaches: a family of AI approaches inspired by biological systems, rather than a precise technique. These include genetic algorithms, which mimic genetic evolution mechanisms to better adapt decisions to new problems and new data, and swarm intelligence, where simple rules implemented by individual agents can lead to sophisticated and robust behavior via interaction at group-level.

Classification and regression trees: predictive models to support decision-making that use tree-like representations of facts and their possible consequences, sometimes referred to as decision trees. The outcome of a classification tree is a discrete value, such as the class to which data belongs. The outcome of a regression tree takes continuous values, such as the price of a house.

Deep learning: a machine learning approach that tries to understand the world in terms of a hierarchy of concepts. Most deep learning models are implemented by increasing the number of layers in a neural network.

Expert system: a computer system that solves complex problems within a specialized domain, usually requiring a high level of human intelligence and expertise. This expertise is expressed manually by human experts in the form of a set of rules which are simple logical tests.

Fuzzy logic: a decision-making approach which is not based on the usual "true or false" assessment, but rather on "degrees of truth" (where the "true" value ranges between completely true and completely false). Fuzzy logic relies on the principle that people make decisions based on imprecise and non-numerical information.

Instance-based learning: a family of machine learning algorithms that compare new problem with cases seen in training and can adapt the model to previously unseen data. It is called "instance-based" learning because it constructs hypotheses directly from the training instances themselves.

Latent representation: the mathematical representation of variables that are inferred rather than directly observed. Latent representation is applied in natural language processing, for example, where it is usually inferred from the statistical distribution of words, and in deep learning, where it is often used for performing transfer learning, i.e. knowledge gained while solving one problem and applying it to a different but related problem.

Logic programming: uses facts and rules to make decisions, without specifying additional intermediary steps, in order to achieve a particular goal.

Machine learning: an AI process that uses algorithms and statistical models to allow computers to make decisions without having to explicitly program it to perform the task. Machine learning algorithms build a model on sample data used as training data in order to identify and extract patterns from data, and therefore acquire their own knowledge. A typical example is a program that identifies and filters spam email.

Multi-task learning: a machine learning approach where a single model is used to solve multiple learning tasks at the same time, exploiting commonalities and differences between the various tasks.

Neural network: a learning process inspired by the neural structures of the brain. The network is a connected framework of many functions (neurons) working together to process multiple data inputs. The network is generally organized in successive layers of functions, each layer using the output of the previous one as an input.

Ontology engineering: a set of tasks related to the methodologies for building ontologies, namely the way concepts and their relationship in a particular domain are formally represented.

Probabilistic graphical models: a framework for representing complex domains using distribution of probabilities, where the models use a graph-based representation for defining the statistical dependence or independence relationships between data.

Probabilistic reasoning: an AI approach which combines deductive logic and probability theory to model logical relations under uncertainty in data.

Reinforcement learning: an area of machine learning that uses a system of reward and punishment for learning how to attain a complex objective. This approach seeks to incentivize software agents to learn correct decisions by trial and error and to pursue a long-term reward.

Rule learning: machine learning methods which identify and generalize automatically a set of rules to be used for prediction or classification of new unseen data. These rules are usually simple conditional tests.

Supervised learning: the most widely adopted form of machine learning. In supervised learning the expected grouping of the information in certain categories (output) is provided to the computer through examples of data (input) which have been manually categorized correctly and form the training dataset. Based on these examples of input-output, the AI system can categorize new, unseen data into the predefined categories.

Support vector machine: a supervised learning algorithm that analyzes labeled/grouped data, identifies the data points that are most challenging to group and, based on that, identifies how to separate the different groups and classify unseen data points. The name "support vector machine" comes from the boundary lines that separate the different groups of data.

Unsupervised learning: a type of machine learning algorithm that finds and analyzes hidden patterns or commonalities in data that has not been labeled or classified. Unlike supervised learning, the system has not been provided with a predefined set of classes, but rather identifies patterns and creates labels/groups in which it classifies the data.

AI functional applications

AI functional applications cover the functions performed by AI techniques, independent of the field of application. These functional applications are categorized as follows.

Augmented reality: this computer vision application provides an interactive experience of a real-world environment, where elements from the real-world are "augmented" by computer-generated sensory information and layered over with the natural environment.

Biometrics: deals with the recognition of people based on physiological characteristics, such as face, fingerprint, vascular pattern or iris, and behavioral traits, such as gait or speech. It combines computer vision with knowledge of human physiology and behavior.

Character recognition: the process of reading typed, handwritten or printed text and converting it into machine-encoded text. A subset of image recognition, it is also known as optical character recognition or reader (OCR).

Computer vision: an interdisciplinary field that deals with how computers see and understand digital images and videos. Computer vision spans all tasks performed by biological vision systems, including "seeing" or sensing a visual stimulus, understanding what is being seen, and extracting complex information into a form that can be used in other processes.

Distributed AI: systems consisting of distributed, multiple autonomous learning agents which process independently data and provide partial solutions which are then integrated, through communication nodes connecting the individual agents. Distributed AI systems can by design aim at solving complex learning and decision-making tasks, involving large data sets and requiring high computational power.

Image and video segmentation: the process of breaking down a digital image into multiple segments or analyzing the images constituting a video, assigning a label to every pixel in an image, in order to simplify or change the representation of an image into something that is more meaningful and easier to analyze. This process is typically used to locate objects and boundaries (lines, curves, etc.) in images.

Information extraction: the task of extracting structured information from unstructured or semi-structured textual sources.

Knowledge representation and reasoning: the field dedicated to representing information about the world usable by a computer to solve complex tasks. These representations are usually based on the way humans represent knowledge, reason (for instance through rules and building relations of sets and sub-sets) and solve problems.

Natural language processing: use of algorithms to analyze human (natural) language data so that computers can understand what humans have written or said and further interact with them.

Object tracking: the process of locating one or more moving objects over time in a video.

Planning/scheduling: the realization of strategies or action sequences for execution by intelligent agents, such as autonomous robots and unmanned vehicles.

Predictive analytics: the process of making predictions about future or otherwise unknown events using a variety of statistical techniques to analyze current and historical facts.

Robotics: the design, construction and operation of machines able to follow step-by-step instructions or perform complex actions automatically and with a certain level of autonomy. Robotics combines hardware with the implementation of AI techniques to perform these tasks.

Scene understanding: the process, often in real-time, of perceiving, analyzing and elaborating an interpretation of a scene and objects in context with respect to the 3D structure of the scene, its layout, and the spatial, functional, and semantic relationships between objects.

Semantics: the automatic recognition and disambiguation of topics and concepts in raw text, image or video, and the application of reasoning for further identifying new associations and facts.

Sentiment analysis: the identification, extraction, analysis and categorization of affective state or opinion from text, social media activity, audio, video or biometric sensors information.

Speech processing: systems involving analysis of speech signals, including speech recognition, natural language processing and speech synthesis.

Speech recognition: the process of identifying words in spoken language and of translating them into text.

Speech synthesis: the artificial production of human speech.

Speaker recognition: the identification of a person from the characteristics of their voice.

Speech-to-speech application: an end-to-end systems where the input and output are a raw audio voice signal, which can be different (another voice or another language) or enhanced (de-noised).

AI application fields

AI technologies can be applied to multiple fields, as summarized below.

Banking and finance: Machine learning is already deeply integrated into many aspects of financial systems, from the approval of loans, to the management of assets and the assessment of risks. Automated trading systems involve the use of complex AI algorithms to make extremely fast trading decisions. Modern fraud detection systems actively learn new potential security threats. AI is predicted to have an impact on financial customer services in the future, with specialized chatbots and voice assistants, recommendation systems for financial products and for improving safety by exploiting advances in biometric systems.

Business: AI techniques are already commonly used for improving marketing and advertising, personalization and product recommendations. Many companies rely on AI algorithms to identify trends and insights into customer data and to make faster decisions with the objective of following their impact on the market in real-time.

Document management and publishing: Over the past two decades, AI has been continuously improving automatic data extraction, structuring and conversion of documents (including automatic translation). Improved document clustering and advanced data analytics are expected to better exploit the huge volume of documents that exist. AI-powered document management systems could also enhance security and protect customer data.

Industry and manufacturing: AI is likely to have major impact on industry and manufacturing. Predictive maintenance is expected to limit costs related to unplanned downtime and malfunction. AI algorithms should also help companies to cope with the increasing complexity of products, engineering processes and quality regulations. Improved robots are expected to handle more cognitive tasks and make autonomous decisions. Generative design systems are able to quickly generate, explore and optimize design alternatives from a set of high-level design goals. Continuous monitoring of the market by AI tools could help proactively to optimize staffing, inventory, energy consumption and the supply of raw materials.

Life and medical sciences: Automatic diagnostic systems are a very promising application of new machine learning techniques. Recent results have shown that it is possible to surpass human expert accuracy for several narrow tasks, such as detection of melanoma or risks of atherosclerosis in arteries. Drug personalization is also frequently cited as a key marker of progress driven by AI. The availability of large amounts of clinical data mean AI is predicted to improve drug discovery and reduce development costs by helping select the most promising hypotheses and focus on more targeted research.

Security: Cyber-security (spam filtering, intrusion-detection) has benefited from machine learning since the 1990s. Automated surveillance is developing quickly, sometimes in conjunction with smart city technologies. AI techniques such as face detection, behavior and crowd analysis are mature enough to make surveillance cameras more "active" without the need for human supervision. Predictive policing technology has started to be used in several U.S. states and the U.K. and AI techniques are also integrated in mass surveillance programs. AI is also considered as a new enabler for a vast range of military requirements, including intelligence, surveillance, reconnaissance, logistics, battlefield planning, weapons systems and defense/offense decisions.

Telecommunications: AI is expected to drive new opportunities in telecoms by helping to improve network performance, thanks to anomaly detection and prediction of service degradations, and also by optimizing customer services.

Transportation: Fuzzy logic and other AI approaches have been used in transportation since the 1980s. It is widely predicted that autonomous vehicles will save costs, lower emissions and enhance road safety, and that AI will improve traffic management by reducing traffic jams and make possible crewless cargo ships and fully automated package delivery.