



Towards a Reconciliation Between Reasoning and Learning - A Position Paper

Didier Dubois and Henri Prade^(✉)

IRIT - CNRS, 118 route de Narbonne, 31062 Toulouse Cedex 09, France
{dubois,prade}@irit.fr

Abstract. The paper first examines the contours of artificial intelligence (AI) at its beginnings, more than sixty years ago, and points out the important place that machine learning already had at that time. The ambition of AI of making machines capable of performing any information processing task that the human mind can do, means that AI should cover the two modes of human thinking: the instinctive (reactive) one and the deliberative one. This also corresponds to the difference between mastering a skill without being able to articulate it and holding some pieces of knowledge that one can use to explain and teach. In case a function-based representation applies to a considered AI problem, the respective merits of learning a universal approximation of the function vs. a rule-based representation are discussed, with a view to better draw the contours of AI. Moreover, the paper reviews the relative positions of knowledge and data in reasoning and learning, and advocates the need for bridging the two tasks. The paper is also a plea for a unified view of the various facets of AI as a science.

1 Introduction

What is artificial intelligence (AI) about? What are the research topics that belong to AI? What are the topics that stand outside? In other words, what are the contours of AI? Answers to these questions may have evolved with time, as did the issue of the proper way (if any) of doing AI. Indeed over time, AI has been successively dominated by logical approaches (until the mid 1990's) giving birth to the so-called “symbolic AI”, then by (Bayesian) probabilistic approaches, and since recently by another type of numerical approach, artificial neural networks. This state of facts has contributed to developing antagonistic feelings between different schools of thought, including claims of supremacy of some methods over others, rather than fostering attempts to understand the

A preliminary version of this paper was presented at the 2018 IJCAI-ECAI workshop “Learning and Reasoning: Principles & Applications to Everyday Spatial and Temporal Knowledge”, Stockholm, July 13–14.

potential complementarity of approaches. Moreover, when some breakthrough takes place in some sector of AI such as expert systems in the 1980's, or fuzzy logic in the 1990's (outside mainstream AI), or yet deep learning [51] nowadays, it is presented through its technological achievements rather than its actual scientific results. So we may even - provocatively - wonder: Is AI a science, or just a bunch of engineering tools? In fact, AI has developed over more than sixty years in several directions, and many different tools have been proposed for a variety of purposes. This increasing diversity, rather than being a valuable asset, may be harmful for an understanding of AI as a whole, all the more so as most AI researchers are highly specialized in some area and are largely ignoring the rest of the field.

Besides, beyond the phantasms and fears teased by the phrase '*artificial intelligence*', the meaning of words such as 'intelligence', 'learning', or 'reasoning' has a large spectrum and may refer to quite different facets of human mind activities, which contributes to blur the meaning of what we claim when we are using the acronym AI. Starting with 'intelligence', it is useful to remember the dichotomy popularized in [44] between two modes of thinking: "System 1" which is fast, instinctive and emotional, while "System 2" is slower, more deliberative, and more logical. See [76] for an illustration of similar ideas in the area of radiological diagnosis, where "super-experts" provide correct diagnosis, even on difficult cases, without any deliberation, while "ordinary experts" may hesitate, deliberate on the difficult cases and finally make a wrong diagnosis. Yet, a "super-expert" is able to explain what went wrong to an "ordinary expert" and what important features should have been noticed in the difficult cases.

Darwiche [21] has recently pointed out that what is achieved by deep learning corresponds to tasks that do not require much deliberation, at least for a top expert, and is far from covering all that may be expected from AI. In other words, the system is mastering skills rather than being also able to elaborate knowledge for thinking and communicating about its skills. This is the difference between an excellent driver (without teaching capability) and a driving instructor.

The intended purpose of this paper is to advocate in favor of a unified view of AI both in terms of problems and in terms of methods. The paper is organized as follows. First, in Sect. 2 a reminder on the history of the early years of AI emphasizes the idea that the diversity of AI has been there from its inception. Then Sect. 3 first discusses relations between a function-based view and a rule-based view of problems, in relation with "modeling versus explaining" concerns. The main paradigms of AI are then restated and the need for a variety of approaches ranging from logic to probability and beyond is highlighted. Section 4 reviews the roles of knowledge and data both in reasoning and in machine learning. Then, Sect. 5 points out problems where bridging reasoning and learning might be fruitful. Section 6 calls for a unified view of AI, a necessary condition for letting it become a mature science.

2 A Short Reminder of the Beginnings of AI

To have a better understanding of AI, it may be useful to have a historical view of the emergence of the main ideas underling it [53, 54, 64]. We only focus here on its beginnings. Still it is worth mentioning that exactly three hundreds years before the expression ‘artificial intelligence’ was coined, the English philosopher Thomas Hobbes of Malmesbury (1588–1679) described human thinking as a symbolic manipulation of terms similar to mathematical calculation [39]. Indeed, he wrote “*Per Ratiocinationem autem intelligo computationem.*” (or in English one year later “*By ratiocination I mean computation.*”) The text continues with “*Now to compute, is either to collect the sum of many things that are added together, or to know what remains when one thing is taken out of another. Ratiocination, therefore, is the same with addition and subtraction.*” One page after one reads: “*We must not therefore think that computation, that is, ratiocination, has place only in numbers, as if man were distinguished from other living creatures (which is said to have been the opinion of Pythagoras) by nothing but the faculty of numbering; for magnitude, body, motion, time, degrees of quality, action, conception, proportion, speech and names (in which all the kinds of philosophy consist) are capable of addition and subtraction.*” Such a description appears retrospectively quite consonant with what AI programs are trying to do!

In the late 1940’s with the advent of cybernetics [96], the introduction of artificial neural networks [56]¹, the principle of synaptic plasticity [37] and the concept of computing machines [91] lead to the idea of thinking machines with learning capabilities. In 1950, the idea of machine intelligence appeared in a famous paper by Turing [92], while Shannon [89] was investigating the possibility of a program playing chess, and the young Zadeh [97] was already suggesting multiple-valued logic as a tool for the conception of thinking machines.

As it is well-known, the official birthday act of AI corresponds to a research program whose application for getting a financial support, was written in the summer of 1955, and entitled “A proposal for the Dartmouth summer research project on artificial intelligence” (thus putting the name of the new field in the title!); it was signed by the two fathers of AI, John McCarthy (1927–2011), and Marvin Minsky (1927–2016), and their two mentors Nathaniel Rochester (1919–2001) (who designed the IBM 701 computer and was also interested in neural network computational machines), and Claude Shannon (1916–2001) [55] (in 1950 he was already the founder of digital circuit design theory based on Boolean logic, the founder of information theory, but also the designer of an electromechanical mouse (Theseus) capable of searching through the corridors of a maze until reaching a target and of acquiring and using knowledge from past experience). Then a series of meetings was organized at Dartmouth College (Hanover, New Hampshire, USA) during the summer of 1956. At that time, McCarthy was already interested in symbolic logic representations, while Minsky

¹ One would notice the word ‘logical’ in the title of this pioneering paper.

had already built a neural network learning machine (he was also a friend of Rosenblatt [79] the inventor of perceptrons).

The interests of the six other participants can be roughly divided into reasoning and learning concerns, they were on the one hand Simon (1916–2001), Newell (1927–1992) [63] (together authors with John Clifford Shaw (1922–1991) of a program *The Logic Theorist* able to prove theorems in mathematical logic), and More [60] (a logician interested in natural deduction at that time), and on the other hand Samuel (1901–1990) [81] (author of programs for checkers, and later chess games), Selfridge (1926–2008) [84] (one of the fathers of pattern recognition), and Solomonoff (1926–2009) [90] (already author of a theory of probabilistic induction).

Interestingly enough, as it can be seen, these ten participants, with different backgrounds ranging from psychology to electrical engineering, physics and mathematics, were already the carriers of a large variety of research directions that are still present in modern AI, from machine learning to knowledge representation and reasoning.

3 Representing Functions and Beyond

There are two modes of representation of knowledge, that can be called respectively functional and logical. The first mode consists in building a large, often numerical, function that produces a result when triggered by some input. The second mode consists of separate, possibly related, chunks of explicit knowledge, expressed in some language. The current dominant machine learning paradigm (up to noticeable exceptions) has adopted the functional approach², which ensures impressive successes in tasks requiring reactivity, at the cost of losing explanatory power. Indeed, we can argue that what is learnt is know-how or skills, rather than knowledge. The other, logical, mode of representation, is much more adapted to the encoding of articulated knowledge, reasoning from it, and to the production of explanations via deliberation, but its connection to learning from data is for the most part still in infancy.

A simple starting point for discussing relationships between learning and reasoning is to compare the machineries of a classifier and a rule-based expert system, for diagnosis for instance. In both cases, a function-based view may apply. On the one hand, from a set of examples (of inputs and outputs of the function, such as pairs (symptoms, disease)) one can easily predict the disease corresponding to a new case via its input symptoms, after learning some function (e.g., using neural nets). On the other hand, one may have a set of expert rules stating that if the values of the inputs are such and such, the global evaluation should be in some subset. Such rules are mimicking the function. If collected from an expert, rules may turn out to be much less successful than the function learned from data. Clearly, the first view may provide better approximations and does not require the elicitation of expert rules, which is costly. However, the explanatory power will be poor in any case, because it will not be possible

² Still this function-based approach is often cast in a probabilistic modeling paradigm.

to answer “why not” questions and to articulate explanations based on causal relations. On the contrary, if causal knowledge is explicitly represented in the knowledge base, it has at least the merit of offering a basis for explanations (in a way that should be cognitively more appropriate for the end-user). It is moreover well-known that causal information cannot easily be extracted from data: only correlations can be laid bare if no extra information is added [66].

The fuzzy set literature offers early examples of the replacement of an automatic control law by a set of rules. Indeed Zadeh [98] proposed to use fuzzy expert rules for controlling complex non linear dynamic systems that might be difficult to model using a classical automatic control approach, while skilled humans can do the job. This was rapidly shown to be successful [52]. The fact of using fuzzy rules, rather than standard Boolean if-then rules, had the advantage of providing a basis for an interpolation mechanism, when an input was firing several rules to some degree. Although the approach was numerical and quite far from the symbolic logic-based AI mainstream trend in those times, it was perceived as an AI-inspired approach, since it was relying on the representation of expert know-how by chunks of knowledge, rather than on the derivation of a control law from the modeling of the physical system to be controlled (i.e., the classical control engineering paradigm). After some time, it was soon recognized that fuzzy rules could be learnt rather than obtained from experts, while keeping excellent results thanks to the property of universal approximation possessed by sets of fuzzy rules. Mathematical models of such fuzzy rules are in fact closely related to neural network radial basis functions. But, fuzzy rules thus obtained by learning may become hardly intelligible. This research trend, known under the names of ‘soft computing’ or ‘computational intelligence’, thus often drifted away from an important AI concern, the explainability power; see [27] for a discussion.

The long term ambition of AI is to make machines capable of performing any information processing task the human mind can perform. This certainly includes recognition, identification, decision and diagnosis tasks (including sophisticated ones). They are “System 1” tasks (using Kahneman terminology) as long as we do not need to explain and reason about obtained results. But there are other problems that are not fully of this kind, even if machine learning may also play a role in their solving. Consider for instance the solving of quadratic equations. Even if we could predict, in a bounded domain, by machine learning techniques, whether an equation has zero, one or two solutions and what are their values (with a good approximation) from a large amount of examples, the solving of such equations by discovering their analytical solution(s), via factorization through symbolic calculations, seems to be a more powerful way of handling of the problem (the machine could then teach students).

AI problems cannot always be viewed in terms of the function-based view mentioned above. There are cases where we do not have a function, only a one-to-many mapping, e.g., when finding all the solutions (if any) of a set of constraints. Apart from solving combinatorial problems, tasks such as reasoning about static or dynamical situations, or building action plans, or explaining results, commu-

nicating explanations pertaining to machine decisions in a meaningful way to an end-user, or analyzing arguments and determining their possible weakness, or understanding what is going on in a text, a dialog in natural language, in an image, a video, or finding relevant information and summarizing it are examples that may require capabilities beyond pure machine learning. This is why AI, over the years, has developed general representation settings and methods capable of handling large classes of situations, while mastering computation complexity. Thus, at least five general paradigms have emerged in AI:

- **Knowledge representation** with symbolic or numerical structured settings for representing knowledge or preferences, such as logical languages, graphical representations like Bayesian networks, or domain ontologies describing taxonomy of concepts. Dedicated settings have been also developed for the representation of temporal or spatial information, of uncertain pieces of information, or of independence relations.
- **Reasoning and decision** Different types of reasoning tasks, beyond classical deduction, have been formalized such as: non monotonic reasoning for dealing with exception-tolerant rules in the presence of incomplete information, or reasoning from inconsistent information, or belief revision, belief updating, information fusion in the presence of conflicts, or formal argumentation handling pros and cons, or yet reasoning directly from data (case-based reasoning, analogical reasoning, interpolation, extrapolation). Models for qualitative (or quantitative) decision from compact representations have been proposed for decision under uncertainty, multiple criteria, or group decisions.
- **General algorithms for problem solving** This covers a panoply of generic tools ranging from heuristic ordered search methods, general problem solver techniques, methods for handling constraints satisfaction problems, to efficient algorithms for classical logic inference (e.g., SAT methods), or for deduction in modal and other non-classical logics.
- **Learning** The word ‘learning’ also covers different problems, from the classification of new items based on a set of examples (and counter-examples), the induction of general laws describing concepts, the synthesis of a function by regression, the clustering of similar data (separating dissimilar data into different clusters) and the labelling of clusters, to reinforcement learning and to the discovery of regularities in data bases and data mining. Moreover, each of these problems can often be solved by a variety of methods.
- **Multiple agent AI** Under this umbrella, there are quite different problems such as: the cooperation between human or artificial agents and the organization of tasks for achieving collective goals, the modeling of BDI agents (Belief, Desire, Intention), possibly in situations of dialogue (where, e.g., agents, which have different information items at their disposal, do not pursue the same goals, and try to guess the intentions of the other ones), or the study of the emergence of collective behaviors from the behaviors of elementary agents.

4 Reasoning with Knowledge or with Data

In the above research areas, knowledge and data are often handled separately. In fact, AI traditionally deals with knowledge rather than with data, with the important exception of machine learning, whose aim can sometimes be viewed as changing data into knowledge. Indeed, basic knowledge is obtained from data by induction, while prior background knowledge may help learning machineries. These remarks suggest that the joint handling of knowledge and data is a general issue, and that combining reasoning and learning methods should be seriously considered.

Rule-based systems, or ontologies expressed by means of description logics, or yet Bayesian networks, represent background knowledge that is useful to make prediction from facts and data. In these reasoning tasks, knowledge as well as data is often pervaded with uncertainty. This has been extensively investigated.

Data, provided that they are reliable, are positive in nature since their existence manifests the *actual* possibility of what is observed or reported. This contrasts with knowledge that delimit the extent of what is *potentially* possible by specifying what is impossible (which has thus a negative flavor). This is why reasoning from both knowledge and data goes much beyond the application of generic knowledge to factual data as in expert systems, and even the separate treatment of knowledge and data in description logics via ‘TBox’ and ‘ABox’ [4]. It is a complex issue, which has received little attention until now [93].

As pointed out in [71], reasoning directly with data has been much less studied. The idea of similarity naturally applies to data and gives birth to specific forms of reasoning such as case-based reasoning [45], case-based decision [35], or even case-based argumentation. “Betweenness” and similarity are at the basis of interpolation mechanisms, while analogical reasoning, which may be both a matter of similarity and dissimilarity, provides a mechanism for extrapolation. A well-known way of handling similarity and interpolation is to use fuzzy rules (where fuzzy set membership degrees capture the idea of similarity w.r.t. the core value(s) of the fuzzy set) [67]. Besides, analogical reasoning, based on analogical proportions (i.e., statements of the form “ a is to b as c is to d ”, where items a, b, c, d are represented in terms of Boolean, nominal or numerical variables), which can be logically represented [28, 58, 72], provides an extrapolation mechanism that from three items a, b, c described by complete vectors, amounts to inferring the missing value(s) in incomplete vector d , providing that a, b, c, d makes an analogical proportion component-wise on the known part of d ; this was successfully applied to classification [14, 18, 57], and more recently to preference learning [13, 32].

Lastly, the ideas of interpolation and extrapolation closely related to analogical proportion-based inference seem to be of crucial importance in many numerical domains. They can be applied to symbolic settings in the case of propositional categorization rules, using relations of betweenness and parallelism respectively, under a conceptual spaces semantics [83]; see [82] for an illustration.

5 Issues in Learning: Incomplete Data and Representation Formats

The need for reasoning from incomplete, uncertain, vague, or inconsistent information, has led to the development of new approaches beyond logic and probability. Incompleteness is a well-known phenomenon in classical logic. However, many reasoning problems exceed the capabilities of classical logic (initially developed in relation with the foundations of mathematics where statements are true or false, and there is no uncertainty in principle). As for probability theory, single probability distributions, often modeled by Bayesian networks are not fully appropriate for handling incomplete information nor epistemic uncertainty. There are different, but related, frameworks for modeling ill-known probabilities that were developed in the last 50 years by the Artificial Intelligence community at large [95]: belief functions and evidence theory (which may be viewed as a randomization of the set-based approach to incomplete information), imprecise probability theory [3,94] (which uses convex families of probability functions) and quantitative possibility theory (which is the simplest model since one of the lower and the upper probability bounds is trivial).

The traditional approach for going from data to knowledge is to resort to statistical inferential methods. However, these methods used to assume data that are precise and in sufficient quantity. The recent concern with big data seems to even strengthen the relevance of probability theory and statistics. However there are a number of circumstances where data is missing or is of poor quality, especially if one tries to collect information for building machines or algorithms supposed to face very complex or unexpected situations (e.g., autonomous vehicles in crowded areas). The concern of Artificial Intelligence for reasoning about partial knowledge has led to a questioning of traditional statistical methods when data is of poor quality [19,38,42,43].

Besides, the fact that we may have to work with incomplete relational data and that knowledge may also be uncertain has motivated the development of a new probabilistic programming language first called “Probabilistic Similarity Logic”, and then “Probabilistic Soft Logic” (PSL, for short) where each ground atom in a rule has a truth value in $[0, 1]$. It uses the Łukasiewicz t-norm and co-t-norm to handle the fuzzy logical connectives [5,33,34]. We are close to representation concerns of fuzzy answer set programs [61]. Besides, there is a need for combining symbolic reasoning with the subsymbolic vector representation of neural networks in order to use gradient descent for training the neural network to infer facts from an incomplete knowledge base, using similarity between vectors [16,17,78].

Machine learning may find some advantages to use advanced representation formats as target languages, such as weighted logics [26] (Markov logic, probabilistic logic programs, multi-valued logics, possibilistic logic, etc.). For instance, qualitative possibility theory extends classical logic by attaching lower bounds of necessity degrees and captures nonmonotonic reasoning, while generalized possibilistic logic [30] is more powerful and can capture answer-set programming, or reason about the ignorance of an agent. Can such kinds of qualitative uncertainty

modeling, or yet fuzzy or uncertain description logics, uncertainty representation formalisms, weighted logics, be used more extensively in machine learning? Various answers and proposals can be found in [48–50, 86, 88]. This also raises the question of extending version space learning [59] to such new representation schemes [41, 73, 75].

If-then rules, in classical logic formats, are a popular representation format in relational learning [80]. Association rules have logical and statistical bases; they are rules with exceptions completed by confidence and support degrees [1, 36]. But, other types of rules may be of interest. Mining genuine default rules that obey Kraus, Lehmann and Magidor postulates [47] for nonmonotonic reasoning relies on the discovery of big-stepped probabilities [8] in a database [9]. Multiple threshold rules, i.e., rules describing how a global evaluation depends on multiple criteria evaluations on linearly ordered scales, such as, e.g., selection rules of the form “if $x_1 \geq a_1$ and \dots and $x_n \geq a_n$ then $y \geq b$ ” play a central role in ordinal classification [46] and can be represented by Sugeno integrals or their extensions [15, 74]. Gradual rules, i.e., statements of the form “the more x is A , the more y is B ”, where A , and B are fuzzy sets, are another representation format of interest [65, 87]. Other types of fuzzy rules may provide a rule-based interpretation [20] for neural nets, which may be also related to non-monotonic inference [7, 22]. All these examples indicate the variety of rules that makes sense and be considered both in reasoning and in learning.

Another trend of research has been also motivated by the extraction of symbolic knowledge from neural networks [22] under the form of nonmonotonic rules. The goal of a neuro-symbolic integration has been pursued with the proposal of a connectionist modal logic, where extended modal logic programs are translated into neural network ensembles, thus providing a neural net view of, e.g., the muddy children problem [24]. Following a similar line of thought, the same authors translate a logic program encoding an argumentation network, which is then turned into a neural network for arguments [23]. A more recent series of works [25, 85, 86] propose another form of integration between logic and neural nets using a so-called “Real Logic”, implemented in deep Tensor Neural Networks, for integrating deductive reasoning and machine learning. The semantics of the logical constants is in terms of vectors of real numbers, and first order logic formulas have degrees of truth in $[0, 1]$ handled with Łukasiewicz multiple-valued logic connectives. Somewhat related is a work on ontology reasoning [40] where the goal is to generate a neural network with binary outputs that, given a database storing tuples of the form (subject, predicate, object), is able, for any input literal, to decide the entailment problem for a logic program describing the ontology. Others look for an exact representation of a binarized neural network as a Boolean formula [62].

The use of degrees of truth multiple-valued logic raises the question of the exact meaning of these degrees. In relation with this kind of work, some have advocated a non-probabilistic view of uncertainty [11], but strangely enough without any reference to the other uncertainty representation frameworks! Maybe more promising is the line of research initiated a long time ago by

Pinkas [68, 69] where the idea of penalty logic (related to belief functions [31]) has been developed in relation with neural networks, where penalty weights reflect priorities attached to logical constraints to be satisfied by a neural network [70]. Penalty logics and Markov logic [77] are also closely related to possibilistic logic [30].

Another intriguing question would be to explore possible relations between spikes neurons [12], which are physiologically more plausible than classical artificial neural networks, and fire when conjunctions of thresholds are reached, with Sugeno integrals (then viewed as a System 1-like black box) and their logical counterparts [29] (corresponding to a System 2-like representation).

6 Conclusion

Knowledge representation and reasoning on the one hand, and machine learning on the other hand, have been developed largely as independent research trends in artificial intelligence in the last three decades. Yet, reasoning and learning are two basic capabilities of the human mind that do interact. Similarly the two corresponding AI research areas may benefit from mutual exchanges. Current learning methods derive know-how from data in the form of complex functions involving many tuning parameters, but they should also aim at producing articulated knowledge, so that repositories, storing interpretable chunks of information, could be fed from data. More precisely, a number of logical-like formalisms, whose explanatory capabilities could be exploited, have been developed in the last 30 years (non-monotonic logics, modal logics, logic programming, probabilistic and possibilistic logics, many-valued logics, etc.) that could be used as target languages for learning techniques, without restricting to first-order logic, nor to Bayes nets.

Interfacing classifiers with human users may require some ability to provide high level explanations about recommendations or decisions that are understandable by an end-user. Reasoning methods should handle knowledge and information extracted from data. The joint use of (supervised or unsupervised) machine learning techniques and of inference machineries raises new issues. There is a number of other points, worth mentioning, which have not be addressed in the above discussions:

- *Teachability* A related issue is more generally how to move from machine learning models to knowledge communicated to humans, about the way the machine proceeds when solving problems.
- *Using prior knowledge* Another issue is a more systematic exploitation of symbolic background knowledge in machine learning devices. Can prior causal knowledge help exploiting data and getting rid of spurious correlations? Can an argumentation-based view of learning be developed?
- *Representation learning* Data representation impacts the performance of machine learning algorithms [10]. In that respect, what may be, for instance, the role of vector space embeddings, or conceptual spaces?

- *Unification of learning paradigms* Would it be possible to bridge learning paradigms from transduction to inductive logic programming? Even including formal concept analysis, or rough set theory?

This paper has especially advocated the interest of a cooperation between two basic areas of AI: knowledge representation and reasoning on the one hand and machine learning on the other hand, reflecting the natural cooperation between two modes, respectively reactive and deliberative, of human intelligence. It is also a plea for maintaining a unified view of AI, all facets of which have been present from the very beginning, as recalled in Sect. 2 of this paper. It is time that AI comes of age as a genuine science, which means ending unproductive rivalries between different approaches, and fostering a better shared understanding of the basics of AI through open-minded studies bridging sub-areas in a constructive way. In the same spirit, a plea for a unified view of computer science can be found in [6]. Mixing, bridging, hybridizing advanced ideas in knowledge representation, reasoning, and machine learning or data mining should renew basic research in AI and contribute in the long term to a more unified view of AI methodology. The interested reader may follow the work in progress of the group “Amel” [2] aiming at a better mutual understanding of research trends in knowledge representation, reasoning and machine learning, and how they could cooperate.

Acknowledgements. The authors thank Emiliano Lorini, Dominique Longin, Gilles Richard, Steven Schockaert, Mathieu Serrurier for useful exchanges on some of the issues surveyed in this paper. This work was partially supported by ANR-11-LABX-0040-CIMI (Centre International de Mathématiques et d’Informatique) within the program ANR-11-IDEX-0002-02, project ISIPA.

References

1. Agrawal, R., Imielinski, T., Swami, A.N.: Mining association rules between sets of items in large databases. In: Buneman, P., Jajodia, S. (eds.) Proceedings 1993 ACM SIGMOD International Conference on Management of Data, Washington, DC, 26–28 May 1993, pp. 207–216. ACM Press (1993)
2. Amel, K.R.: From shallow to deep interactions between knowledge representation, reasoning and machine learning. In: BenAmor, N., Theobald, M. (eds.) Proceedings 13th International Conference Scala Uncertainty Mgmt (SUM 2019), Compiègne, LNCS, 16–18 December 2019. Springer, Heidelberg (2019)
3. Augustin, T., Coolen, F.P.A., De Cooman, G., Troffaes, M.C.M.: Introduction to Imprecise Probabilities. Wiley, Hoboken (2014)
4. Baader, F., Horrocks, I., Lutz, C., Sattler, U.: An Introduction to Description Logic. Cambridge University Press, Cambridge (2017)
5. Bach, S.H., Broecheler, M., Huang, B., Getoor, L.: Hinge-loss Markov random fields and probabilistic soft logic. *J. Mach. Learn. Res.* **18**, 109:1–109:67 (2017)
6. Bajcsy, R., Reynolds, C.W.: Computer science: the science of and about information and computation. *Commun. ACM* **45**(3), 94–98 (2002)
7. Balkenius, C., Gärdenfors, P.: Nonmonotonic inferences in neural networks. In: Proceedings 2nd International Conference on Principle of Knowledge Representation and Reasoning (KR 1991), Cambridge, MA, pp. 32–39 (1991)

8. Benferhat, S., Dubois, D., Prade, H.: Possibilistic and standard probabilistic semantics of conditional knowledge bases. *J. Log. Comput.* **9**(6), 873–895 (1999)
9. Benferhat, S., Dubois, D., Lagrue, S., Prade, H.: A big-stepped probability approach for discovering default rules. *Int. J. Uncert. Fuzz. Knowl.-based Syst.* **11**(Suppl.-1), 1–14 (2003)
10. Bengio, Y., Courville, A., Vincent, P.: Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(8), 1798–1828 (2013)
11. Besold, T.R., Garcez, A.D.A., Stenning, K., van der Torre, L., van Lambalgen, M.: Reasoning in non-probabilistic uncertainty: logic programming and neural-symbolic computing as examples. *Minds Mach.* **27**(1), 37–77 (2017)
12. Bichler, O., Querlioz, D., Thorpe, S.J., Bourgoin, J.-P., Gamrat, C.: Extraction of temporally correlated features from dynamic vision sensors with spike-timing-dependent plasticity. *Neural Netw.* **32**, 339–348 (2012)
13. Bounhas, M., Pirlot, M., Prade, H., Sobrie, O.: Comparison of analogy-based methods for predicting preferences. In: BenAmor, N., Theobald, M. (eds.) *Proceedings 13th International Conference on Scala Uncertainty Mgmt (SUM 2019)*, Compiègne, LNCS, 16–18 December. Springer, Heidelberg (2019)
14. Bounhas, M., Prade, H., Richard, G.: Analogy-based classifiers for nominal or numerical data. *Int. J. Approx. Reasoning* **91**, 36–55 (2017)
15. Brabant, Q., Couceiro, M., Dubois, D., Prade, H., Rico, A.: Extracting decision rules from qualitative data via sugeno utility functionals. In: Medina, J., Ojeda-Aciego, M., Verdegay, J.L., Pelta, D.A., Cabrera, I.P., Bouchon-Meunier, B., Yager, R.R. (eds.) *IPMU 2018. CCIS*, vol. 853, pp. 253–265. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-91473-2_22
16. Cohen, W.W.: TensorLog: a differentiable deductive database. *CoRR*, abs/1605.06523 (2016)
17. Cohen, W.W., Yang, F., Mazaitis, K.: TensorLog: deep learning meets probabilistic DBs. *CoRR*, abs/1707.05390 (2017)
18. Couceiro, M., Hug, N., Prade, H., Richard, G.: Analogy-preserving functions: a way to extend Boolean samples. In: *Proceedings 26th International Joint Conference on Artificial Intelligence (IJCAI 2017)*, Melbourne, 19–25 August, pp. 1575–1581 (2017)
19. Couso, I., Dubois, D.: A general framework for maximizing likelihood under incomplete data. *Int. J. Approx. Reasoning* **93**, 238–260 (2018)
20. d’Alché-Buc, F., Andrés, V., Nadal, J.-P.: Rule extraction with fuzzy neural network. *Int. J. Neural Syst.* **5**(1), 1–11 (1994)
21. Darwiche, A.: Human-level intelligence or animal-like abilities?. *CoRR*, abs/1707.04327 (2017)
22. d’Avila Garcez, A.S., Broda, K., Gabbay, D.M.: Symbolic knowledge extraction from trained neural networks: a sound approach. *Artif. Intell.* **125**(1–2), 155–207 (2001)
23. d’Avila Garcez, A.S., Gabbay, D.M., Lamb, L.C.: Value-based argumentation frameworks as neural-symbolic learning systems. *J. Logic Comput.* **15**(6), 1041–1058 (2005)
24. d’Avila Garcez, A.S., Lamb, L.C., Gabbay, D.M.: Connectionist modal logic: representing modalities in neural networks. *Theor. Comput. Sci.* **371**(1–2), 34–53 (2007)
25. Donadello, I., Serafini, L., Garcez, A.D.A.: Logic tensor networks for semantic image interpretation. In: Sierra, C. (ed) *Proceedings 26th International Joint Conference on Artificial Intelligence (IJCAI 2017)*, Melbourne, 19–25 August 2017, pp. 1596–1602 (2017)

26. Dubois, D., Godo, L., Prade, H.: Weighted logics for artificial intelligence - an introductory discussion. *Int. J. Approx. Reasoning* **55**(9), 1819–1829 (2014)
27. Dubois, D., Prade, H.: Soft computing, fuzzy logic, and artificial intelligence. *Soft Comput.* **2**(1), 7–11 (1998)
28. Dubois, D., Prade, H., Richard, G.: Multiple-valued extensions of analogical proportions. *Fuzzy Sets Syst.* **292**, 193–202 (2016)
29. Dubois, D., Prade, H., Rico, A.: The logical encoding of Sugeno integrals. *Fuzzy Sets Syst.* **241**, 61–75 (2014)
30. Dubois, D., Prade, H., Schockaert, S.: Generalized possibilistic logic: foundations and applications to qualitative reasoning about uncertainty. *Artif. Intell.* **252**, 139–174 (2017)
31. Dupin de Saint-Cyr, F., Lang, J., Schiex, T.: Penalty logic and its link with Dempster-Shafer theory. In: de Mántaras, R.L., Poole, D. (eds.) *Proceedings 10th Annual Conference on Uncertainty in Artificial Intelligence (UAI 1994)*, Seattle, 29–31 July, pp. 204–211 (1994)
32. Fahandar, M.A., Hüllermeier, E.: Learning to rank based on analogical reasoning. In: *Proceedings 32th National Conference on Artificial Intelligence (AAAI 2018)*, New Orleans, 2–7 February 2018 (2018)
33. Fakhraei, S., Raschid, L., Getoor, L.: Drug-target interaction prediction for drug repurposing with probabilistic similarity logic. In: *SIGKDD 12th International Workshop on Data Mining in Bioinformatics (BIOKDD)*. ACM (2013)
34. Farnadi, G., Bach, S.H., Moens, M.F., Getoor, L., De Cock, M.: Extending PSL with fuzzy quantifiers. In: *Papers from the 2014 AAAI Workshop Statistical Relational Artificial Intelligence*, Québec City, 27 July, pp. WS-14-13, 35–37 (2014)
35. Gilboa, I., Schmeidler, D.: Case-based decision theory. *Q. J. Econ.* **110**, 605–639 (1995)
36. Hájek, P., Havránek, T.: *Mechanising Hypothesis Formation - Mathematical Foundations for a General Theory*. Springer, Heidelberg (1978). <https://doi.org/10.1007/978-3-642-66943-9>
37. Hebb, D.O.: *The Organization of Behaviour*. Wiley, Hoboken (1949)
38. Heitjan, D., Rubin, D.: Ignorability and coarse ckata. *Ann. Statist.* **19**, 2244–2253 (1991)
39. Hobbes, T.: *Elements of philosophy, the first section, concerning body*. In: Molesworth, W. (ed.) *The English works of Thomas Hobbes of Malmesbury*, vol. 1. John Bohn, London, 1839. English translation of "Elementa Philosophiae I. De Corpore" (1655)
40. Hohenecker, P., Lukasiewicz, T.: *Ontology reasoning with deep neural networks*. CoRR, abs/1808.07980 (2018)
41. Hüllermeier, E.: Inducing fuzzy concepts through extended version space learning. In: Bilgiç, T., De Baets, B., Kaynak, O. (eds.) *IFSA 2003. LNCS*, vol. 2715, pp. 677–684. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-44967-1_81
42. Hüllermeier, E.: Learning from imprecise and fuzzy observations: data disambiguation through generalized loss minimization. *Int. J. Approx. Reasoning* **55**(7), 1519–1534 (2014)
43. Jaeger, M.: Ignorability in statistical and probabilistic inference. *JAIR* **24**, 889–917 (2005)
44. Kahneman, D.: *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York (2011)
45. Kolodner, J.L.: *Case-Based Reasoning*. Morgan Kaufmann, Burlington (1993)
46. Kotlowski, W., Slowinski, R.: On nonparametric ordinal classification with monotonicity constraints. *IEEE Trans. Knowl. Data Eng.* **25**(11), 2576–2589 (2013)

47. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. *Artif. Intell.* **44**, 167–207 (1990)
48. Kuzelka, O., Davis, J., Schockaert, S.: Encoding Markov logic networks in possibilistic logic. In: Meila, M., Heskes, T. (eds.) *Proceedings 31st Conference on Uncertainty in Artificial Intelligence (UAI 2015)*, Amsterdam, 12–16 July 2015, pp. 454–463. AUAI Press (2015)
49. Kuzelka, O., Davis, J., Schockaert, S.: Learning possibilistic logic theories from default rules. In: Kambhampati, S. (ed.) *Proceedings 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*, New York, 9–15 July 2016, pp. 1167–1173 (2016)
50. Kuzelka, O., Davis, J., Schockaert, S.: Induction of interpretable possibilistic logic theories from relational data. In: Sierra, C. (ed.) *Proceedings 26th International Joint Conference on Artificial Intelligence, IJCAI 2017*, Melbourne, 19–25 August 2017, pp. 1153–1159 (2017)
51. LeCun, Y., Bengio, Y., Hinton, G.E.: Deep learning. *Nature* **521**(7553), 436–444 (2015)
52. Mamdani, E.H., Assilian, S.: An experiment in linguistic synthesis with a fuzzy logic controller. *Int. J. Man-Mach. Stu.* **7**, 1–13 (1975)
53. Marquis, P., Papini, O., Prade, H.: *Éléments pour une histoire de l'intelligence artificielle*. In: *Panorama de l'Intelligence Artificielle. Ses Bases Méthodologiques, ses Développements*, vol. I, pp. 1–39. Cepaduès (2014)
54. Marquis, P., Papini, O., Prade, H.: Some elements for a prehistory of Artificial Intelligence in the last four centuries. In: *Proceedings 21st European Conference on Artificial Intelligence (ECAI 2014)*, Prague, pp. 609–614. IOS Press (2014)
55. McCarthy, J., Minsky, M., Roch-ester, N., Shannon, C.E.: A proposal for the Dartmouth summer research project on artificial intelligence, august 31, 1955. *AI Mag.* **27**(4), 12–14 (2006)
56. McCulloch, W.S., Pitts, W.: A logical calculus of ideas immanent in nervous activity. *Bull. Math. Biophys.* **5**, 115–133 (1943)
57. Miclet, L., Bayouh, S., Delhay, A.: Analogical dissimilarity: definition, algorithms and two experiments in machine learning. *JAIR* **32**, 793–824 (2008)
58. Miclet, L., Prade, H.: Handling analogical proportions in classical logic and fuzzy logics settings. In: Sossai, C., Chemello, G. (eds.) *ECSQARU 2009. LNCS (LNAI)*, vol. 5590, pp. 638–650. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02906-6_55
59. Mitchell, T.: *Version spaces: an approach to concept learning*. Ph.D. thesis, Stanford (1979)
60. More, T.: On the construction of Venn diagrams. *J. Symb. Logic* **24**(4), 303–304 (1959)
61. Mushthofa, M., Schockaert, S., De Cock, M.: Solving disjunctive fuzzy answer set programs. In: Calimeri, F., Ianni, G., Truszczyński, M. (eds.) *LPNMR 2015. LNCS (LNAI)*, vol. 9345, pp. 453–466. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23264-5_38
62. Narodytska, N.: Formal analysis of deep binarized neural networks. In: Lang, J. (ed.) *Proceedings 27th International Joint Conference Artificial Intelligence (IJCAI 2018)*, Stockholm, 13–19 July 2018, pp. 5692–5696 (2018)
63. Newell, A., Simon, H.A.: The logic theory machine. a complex information processing system. In: *Proceedings IRE Transactions on Information Theory (IT-2)*, The Rand Corporation, Santa Monica, Ca, 1956. Report P-868, 15 June 1956, pp. 61–79, September 1956

64. Nilsson, N.J.: *The Quest for Artificial Intelligence : A History of Ideas and Achievements*. Cambridge University Press, Cambridge (2010)
65. Nin, J., Laurent, A., Poncet, P.: Speed up gradual rule mining from stream data: A B-tree and OWA-based approach. *J. Intell. Inf. Syst.* **35**(3), 447–463 (2010)
66. Pearl, J.: *Causality*, vol. 2000, 2nd edn. Cambridge University Press, Cambridge (2009)
67. Perfilieva, I., Dubois, D., Prade, H., Esteva, F., Godo, L., Hodáková, P.: Interpolation of fuzzy data: analytical approach and overview. *Fuzzy Sets Syst.* **192**, 134–158 (2012)
68. Pinkas, G.: Propositional non-monotonic reasoning and inconsistency in symmetric neural networks. In: Mylopoulos, J., Reiter, R. (eds.) *Proceedings 12th International Joint Conference on Artificial Intelligence*, Sydney, 24–30 August 1991, pp. 525–531. Morgan Kaufmann (1991)
69. Pinkas, G.: Reasoning, nonmonotonicity and learning in connectionist networks that capture propositional knowledge. *Artif. Intell.* **77**(2), 203–247 (1995)
70. Pinkas, G., Cohen, S.: High-order networks that learn to satisfy logic constraints. *FLAP J. Appl. Logics IfCoLoG J. Logics Appl.* **6**(4), 653–694 (2019)
71. Prade, H.: Reasoning with data - a new challenge for AI? In: Schockaert, S., Senellart, P. (eds.) *SUM 2016. LNCS (LNAI)*, vol. 9858, pp. 274–288. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45856-4_19
72. Prade, H., Richard, G.: From analogical proportion to logical proportions. *Logica Universalis* **7**(4), 441–505 (2013)
73. Prade, H., Rico, A., Serrurier, M.: Elicitation of sugeno integrals: a version space learning perspective. In: Rauch, J., Raś, Z.W., Berka, P., Elomaa, T. (eds.) *ISMIS 2009. LNCS (LNAI)*, vol. 5722, pp. 392–401. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-04125-9_42
74. Prade, H., Rico, A., Serrurier, M., Raufaste, E.: Eliciting sugeno integrals: methodology and a case study. In: Sossai, C., Chemello, G. (eds.) *ECSQARU 2009. LNCS (LNAI)*, vol. 5590, pp. 712–723. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02906-6_61
75. Prade, H., Serrurier, M.: Bipolar version space learning. *Int. J. Intell. Syst.* **23**, 1135–1152 (2008)
76. Raufaste, E.: *Les Mécanismes Cognitifs du Diagnostic Médical : Optimisation et Expertise*. Presses Universitaires de France (PUF), Paris (2001)
77. Richardson, M., Domingos, P.M.: Markov logic networks. *Mach. Learn.* **62**(1–2), 107–136 (2006)
78. Rocktäschel, T., Riedel, S.: End-to-end differentiable proving. In: Guyon, I., et al. (eds.) *Proceedings 31st Annual Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, 4–9 December 2017, pp. 3791–3803 (2017)
79. Rosenblatt, F.: The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* **65**(6), 386–408 (1958)
80. Rückert, U., De Raedt, L.: An experimental evaluation of simplicity in rule learning. *Artif. Intell.* **172**(1), 19–28 (2008)
81. Samuel, A.: Some studies in machine learning using the game of checkers. *IBM J.* **3**, 210–229 (1959)
82. Schockaert, S., Prade, H.: Interpolation and extrapolation in conceptual spaces: a case study in the music domain. In: Rudolph, S., Gutierrez, C. (eds.) *RR 2011. LNCS*, vol. 6902, pp. 217–231. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-23580-1_16

83. Schockaert, S., Prade, H.: Interpolative and extrapolative reasoning in propositional theories using qualitative knowledge about conceptual spaces. *Artif. Intell.* **202**, 86–131 (2013)
84. Selfridge, O.G.: Pandemonium: a paradigm for learning. In: Blake, D.V., Uttley, A.M. (ed) *Symposium on Mechanisation of Thought Processes*, London, 24–27 November 1959, vol. 1958, pp. 511–529 (1959)
85. Serafini, L., Garcez, A.S.A.: Logic tensor networks: deep learning and logical reasoning from data and knowledge. In: Besold, T.R., Lamb, L.C., Serafini, L., Tabor, W. (eds.) *Proceedings 11th International Workshop on Neural-Symbolic Learning and Reasoning (NeSy 2016)*, New York City, 16–17 July 2016, vol. 1768 of *CEUR Workshop Proceedings* (2016)
86. Serafini, L., Donadello, I., Garcez, A.S.A.: Learning and reasoning in logic tensor networks: theory and application to semantic image interpretation. In: Seffah, A., Penzenstadler, B., Alves, C., Peng, X. (eds.) *Proceedings Symposium on Applied Computing (SAC 2017)*, Marrakech, 3–7 April 2017, pp. 125–130. *ACM* (2017)
87. Serrurier, M., Dubois, D., Prade, H., Sudkamp, T.: Learning fuzzy rules with their implication operators. *Data Knowl. Eng.* **60**(1), 71–89 (2007)
88. Serrurier, M., Prade, H.: Introducing possibilistic logic in ILP for dealing with exceptions. *Artif. Intell.* **171**(16–17), 939–950 (2007)
89. Shannon, C.E.: Programming a computer for playing chess. *Philos. Mag.* (7th series) **XLI** (314), 256–275 (1950)
90. Solomonoff, R.J.: *An inductive inference machine*. Tech. Res. Group, New York City (1956)
91. Turing, A.M.: *Intelligent machinery*. Technical report, National Physical Laboratory, London, 1948. Also. In: *Machine Intelligence*, vol. 5, pp. 3–23. *Edinburgh University Press* (1969)
92. Turing, A.M.: Computing machinery and intelligence. *Mind* **59**, 433–460 (1950)
93. Ughetto, L., Dubois, D., Prade, H.: Implicative and conjunctive fuzzy rules - a tool for reasoning from knowledge and examples. In: Hendler, J., Subramanian, D. (eds.) *Proceedings 16th National Conference on Artificial Intelligence*, Orlando, 18–22 July 1999, pp. 214–219 (1999)
94. Walley, P.: *Statistical Reasoning with Imprecise Probabilities*. *Chapman and Hall*, London (1991)
95. Walley, P.: Measures of uncertainty in expert systems. *Artif. Intell.* **83**(1), 1–58 (1996)
96. Wiener, N.: *Cybernetics or Control and Communication in the Animal and the Machine*. *Wiley*, Hoboken (1949)
97. Zadeh, L.A.: Thinking machines - a new field in electrical engineering. *Columbia Eng. Q.* **3**, 12–13 (1950)
98. Zadeh, L.A.: Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Trans. Syst. Man Cybern.* **3**(1), 28–44 (1973)