KNOWLEDGE AND REPRESENTATION pdf

1: Materiality, Knowledge and Disclosure in Representations and Warranties - SGR

Knowledge representation and reasoning (KR, KRÂ², KR&R) is the field of artificial intelligence (AI) dedicated to representing information about the world in a form that a computer system can utilize to solve complex tasks such as diagnosing a medical condition or having a dialog in a natural language.

Knowledge independent elements of the optimization of self-instruction Before I move toward representation of knowledge, I would only shortly like to list some other principles of effective learning that are representation independent: In the process of learning, the coherent graph of semantic links between elements of the learned material is formed. This graph can be constructed by means of item-by-item build up; however, a standard textbook, hypertext document or interactive tutor will always have an efficacy advantage because of no implied granularity constraint. In other words, one should always learn first, and only then work with a repetition spacing system in order to retain the newly formed memory engrams over a longer period of time. An important note is due here: After all, it does not prevent an associative nature of knowledge as stored in memory. It affects only the way stimuli are presented at repetitions in order to maximize the memory effect. Active recall is more demanding from the memory consolidation standpoint, and it is closer to matching the real life situations with respect to what sort of synaptic connections should be strengthened. The art of mnemonics is as old as the art of learning, and professional mnemonists, with their trained memory capabilities, can truly leave an average mortal speechless in face of top mnemonic feats. Indeed, the mnemonic techniques are very easy to apply, and most of capable students, more or less consciously, use them in their daily routine. However, by conscious understanding of the rules and principles, even the best students can gain a great deal. The basic principle of mnemonics, from the neurobiological point of view, is to build memory images from as many previously stored engrams as it is only possible. As visual processing in the human brain seems to involve much more sophisticated circuitry than, for example, verbal processing, extensive use of visual imagery is a key to success. Instead of memorizing a nonsensic telephone number, the student can memorize a collection of visual scenes unequivocally mapped to numbers, and generate a unique, easily memorizable sequence of graphic events that can serve as an effective way of representing the number. Recalling the phone number can then be equivalent to invoking the stored visual event and translating it to the sequence of digits, or more often, to a sequence of two-figure numbers. As I will try to argue in later paragraphs, minimizing the number of synaptic connections involved in storing memories is the key to maximizing retention over a longer period. Representing new memories as easily recoverable composites of old memories serves exactly that end. To simplify the discussion of knowledge representation issues with respect to the complexity of neural connections involved in storing particular engram, I will shortly introduce a concept of a synaptic pattern. As early as since the introduction of sensitive neural activity measurement techniques, it has been known that memories can be associated with spatiotemporal patterns of synaptic activity, or synaptic patterns in short. There exists a substantial terminology confusion as far as naming the concept of the synaptic pattern is concerned. Therefore, it is worth noticing than in relevant literature, the notion of synaptic patterns, often devoid of its temporal component, might be used more or less synonymously with terms such as: As I will try to show in the chapter devoted to biological aspects of memory, the complexity of synaptic patterns is likely to be in strict correlation with item difficulty e. It is therefore central to understanding the principles of effective representation of knowledge in self-instruction systems based on active recall and repetition spacing. Items that do not comply with minimum complexity of synaptic patterns will, in the course of repetition, gradually experience loss of its components. In other words, memory will accomplish natural selection of the core synaptic pattern with elimination of all additional connections that do not get uniformly stimulated at repetitions. In later sections, I will use the term pattern extraction to describe the phenomenon of selecting the core synaptic pattern in the course of repetitions Components of effective knowledge representation in active recall systems All the principles of effective knowledge representation discussed in the following section come from years of experience in developing knowledge systems for my own use, as well as from a continuing opportunity to study the relationship between item difficulty and knowledge representation in knowledge systems developed by users of SuperMemo, many of whom seek professional advice before embarking on major knowledge system development, or equally often are eager to share their own experience and problems encountered while applying repetition spacing in all imaginable learning domains. The following are the five thematic groups related to knowledge representation in self-instruction systems based on active recall: In the five following sections I will address all the above thematic groups individually using examples from the aforementioned knowledge system on microeconomics Sequencing items in the stepwise process of acquiring associative knowledge The most important rule in sequencing items is inherently related to the learning process in general. The progression must go from basic concepts through foundations to more intricate and detailed issues. In all forms of learning, this principle derives from the need for comprehension, which obviously is greatly reduced when the student is thrown at a deep end at once. In application of repetition spacing with low forgetting index , this approach has also another important aspect. As in this form of learning, forgetting plays an insignificant role, the student is likely to experience the phenomenon of new pieces of knowledge nicely slotting in in the already established structure. This might affect the sequencing algorithm by moving it from truly basics-to-details approach to first-comprehended-first-memorized approach which dispenses with the need for full comprehension at the first contact with the newly learned material. This way the student will less likely experience the feeling of getting stuck because of his or her inability to comprehend some concept and, at the same time, reluctance to proceed any further out of fear of snowballing incomprehension. Thus, in the first pass of the material, the student will memorize only those components that have been comprehended, and hope that the slotting-in phenomenon will eliminate of comprehension gaps in the second pass. The basics-to-details approach may be combined successfully with the most-applicable approach, in which the student delves more rapidly into details of those parts of the material which are most frequently referenced in other parts. This enhances the slotting-in phenomenon, which is one of the strongest motivational factors in learning, providing the student with the sense of accomplishment. There is no cut-and-dried algorithm for optimally sequencing items then, however, in slightly more formal terms, the optimum sequencing might be defined as bicriterial optimization in which the following two factors are considered: Production is commonly assumed to be synonymous with manufacturing. However, a more accurate and useful definition of production, from the economic analysis standpoint, is as any activity that creates value. What is the name of an activity that creates value in economics? As the above definition provided an intuitive understanding and severed the link between production and manufacturing, the definition presented above, not quite consistent with the previous approach may have greater applicability in cases where the production process becomes the focus of more detailed analysis, esp. Because of the fundamental nature of the concept, the very exact imprint of the above definition may be considered a valuable asset in building more advanced facets of knowledge of economics. The difficulty with basic concepts is that they are so basic that they cannot be asked for in yet more basic terms. Questions like "what is production? Simplifying the answer to "series of activities" and following it with a number of items that define the "activities" is also inadmissible because of a number of adequate substitutes for "series of activities" like, for example, "any activity that creates value", and many more. Here, a very valuable, and often underappreciated tool of Cloze deletion comes handy. Consider the following wording of items: Production is a series of Production is a series of activities by which resource inputs are transformed through a recipe and technological process into outputs of The experience shows that despite seemingly high degree of dissociation between the particular components of the learned concept, the presented items appear to produce a solid imprint in the students memory that not only provides a firm support for stable comprehension, but also makes it possible to effortlessly recite the entire definition of production. E-factors for such constructed items typically fall into the range from 2. Despite a larger number of items, this is a sure guarantee to produce less workload than in the case of cramming the entire definition in the answer with E-factor most likely to drop below 1. For contrast, let as consider the following item that appeared to show a high degree of intractability because of the lack of respect for basics-to-details approach: What is discount rate? This made the definition of discount rate leave little semantic connotation forcing the student to assume syntactic approach to memorization which is nothing else than mindless cramming with very poor retention prospects. Techniques for minimizing the complexity of synaptic patterns as a key to keeping E-factors high The most important principle of effective knowledge representation in systems based on active recall and repetition spacing is minimization of the complexity of synaptic patterns involved in storing memory engrams Wozniak, This principle translates to keeping the content of question-answer items simple, specific, graphic, consistent, comprehensible and univocal. The main purpose of such an approach is to make sure that the spatiotemporal pattern of firing during the learning task is the same in each successive repetition. In other words, there should be minimum change to the synaptic pattern in the course of repetition as a result of pattern extraction. The entire concept of optimum repetition spacing is based on dealing with uniform pieces of information whose memory engrams are uniform and stable, and consequently can be treated as atomic entities. If the neuronal firing was to change its course over a number of repetition, a subset of synapses in the relevant synaptic pattern would not receive sufficient enhancement resulting in partial loss of the learned information. Using examples from the microeconomics database, I will show all the distinguishable facets of minimization of synaptic patterns involved in representing individual pieces of information stored in the database. Comprehension It has been strongly pointed in the preceding section that the basics-to-details approach is, among other things, supposed to ensure the maximum level of comprehension. Here I will only note that comprehension is indeed related to the minimization of the complexity of synaptic patters that is the subject of this chapter. Nonsensical phrases or concepts involve a much greater number of neurons in the process of learning. Low-level electrical measurements showed that the neural activity is higher in cases of memorizing nonsensical words as opposed to natural words. Similarly, PET scans show that the brain activity of people exhibiting high IQ is much lower during performing a learning tasks that it is the case for low IQ students. This was explained by psychologists as being related to the fact that well-established synaptic patters representing various professions are usually not matched by similar patterns that could be used easily used to represent names, esp. Minimum information principle Minimum information principle is the most obvious consequence of the approach based on minimum complexity of synaptic patterns. In order to keep the memory image of items simple, items should be simple themselves. Let us consider the shortcomings of competitive markets such as unequal distribution of income, imposition of production costs on the public, development of socially undesirable products, product proliferation, etc. The minimum information principles says that the question "What are the shortcomings of competitive markets? In such situations, solution comes from narrowing the focus of the question; the approach which often requires additional terminology and knowledge structuring, and is generally more demanding for the database developer. A typical question with a narrower focus might sound as follows: What problem with distribution of income appears on a competitive market? What is an example of imposing production costs on people who do not consume in competitive markets? The experience shows that a number of items that would glue the above granules in a coherent entirety is necessary. This can conveniently be accomplished by means of Cloze deletions discussed later in the chapter. The main shortcomings of competitive markets are: After all, as it will be shown later, enumerations are one of the trickiest obstacles to overcome in complying with the minimum information principle. The presented Cloze deletion serves as: A yet more complex knowledge structure appears in the analysis of tax revenues in an attempt to plot the Laffer curve for European countries in the years Upon the analysis, on the two ends of the spectrum, notable examples of two countries are worth considering: Naturally, a single item cramming all the above facts has little chance of passing the minimum information criterion. What was the average tax rate in Spain ? What was the change in the tax revenue in Spain ? Naturally, the understanding does not need examples. The theoretical implications of marginal tax revenue might be considered as a sufficient element of understanding; however, the usefulness of facts illustrating the theory has long been appreciated in education; I will therefore present for consideration an exemplary set of items acting as an associative glue for the discussed tax revenue case: In the years , the average tax rate and the tax revenue in Spain and Sweden were as follows: Items formulated in the above way appeared to produce very coherent memory engrams that showed above average retention rate despite the inherent intractability of numeric responses as in the first of the two presented examples. Narrowing by example Narrowing by example is a very efficient way of making the question-related stimulus more specific and thus more successful in imprinting durable memories. The concept of the price ceiling may be enhanced if a narrowing example of goods that might be subject to government imposed price ceiling is provided. Moreover, the example makes the definition of price ceiling more specific; hence the increased likelihood of diminished complexity of the synaptic pattern, and minimization of pattern extraction. What is the name of the price specified by the government above which goods cannot be sold e. Similarly, the illustration of competitive nature of pork versus beef helps narrowing by example in the definition of horizontal markets: What is the name of markets for products that can act as substitutes e. The next section will discuss the various aspects of information redundancy in knowledge representation. In that context, the technique of extending by example will be presented.

KNOWLEDGE AND REPRESENTATION pdf

2: Knowledge Representation and Semantics | AMIA

There is a familiar pattern in knowledge representation research in which the description of a new knowledge representation technology is followed by claims that the new ideas are in fact formally equivalent to an existing technology.

In any intelligent system, representing the knowledge is supposed to be an important technique to encode the knowledge. Knowledge is an useful term to judge the understanding of an individual on a given subject. In intelligent systems, domain is the main focused subject area. So, the system specifically focuses on acquiring the domain knowledge. Types of knowledge in AI Depending on the type of functionality, the knowledge in AI is categorized as: Procedural knowledge Procedural knowledge derives the information on the basis of rules, strategies, agendas and procedure. It describes how a problem can be solved. Procedural knowledge directs the steps on how to perform something. Heuristic knowledge Heuristic knowledge is based on thumb rule. It provides the information based on a thumb rule, which is useful in guiding the reasoning process. In this type, the knowledge representation is based on the strategies to solve the problems through the experience of past problems, compiled by an expert. Hence, it is also known as Shallow knowledge. Meta-knowledge This type gives an idea about the other types of knowledge that are suitable for solving problem. Meta-knowledge is helpful in enhancing the efficiency of problem solving through proper reasoning process. Structural knowledge Structural knowledge is associated with the information based on rules, sets, concepts and relationships. It provides the information necessary for developing the knowledge structures and overall mental model of the problem. Issues in knowledge representation The main objective of knowledge representation is to draw the conclusions from the knowledge, but there are many issues associated with the use of knowledge representation techniques. Some of them are listed below: Refer to the above diagram to refer to the following issues. Important attributes There are two attributes shown in the diagram, instance and isa. Since these attributes support property of inheritance, they are of prime importance. Relationships among attributes Basically, the attributes used to describe objects are nothing but the entities. However, the attributes of an object do not depend on the encoded specific knowledge. Choosing the granularity of representation While deciding the granularity of representation, it is necessary to know the following: What are the primitives and at what level should the knowledge be represented? What should be the number small or large of low-level primitives or high-level facts? High-level facts may be insufficient to draw the conclusion while Low-level primitives may require a lot of storage. Suppose that we are interested in following facts: Now, this could be represented as "Spotted agent John , object Alex " Such a representation can make it easy to answer questions such as: Suppose we want to know: Representing sets of objects. There are some properties of objects which satisfy the condition of a set together but not as individual; Example: Consider the assertion made in the sentences: Finding the right structure as needed To describe a particular situation, it is always important to find the access of right structure. This can be done by selecting an initial structure and then revising the choice. While selecting and reversing the right structure, it is necessary to solve following problem statements. They include the process on how to: Select an initial appropriate structure. Fill the necessary details from the current situations. Determine a better structure if the initially selected structure is not appropriate to fulfill other conditions. Find the solution if none of the available structures is appropriate. Create and remember a new structure for the given condition. Logic Representation Facts are the general statements that may be either True or False. Thus, logic can be used to represent such simple facts. To build a Logic-based representation: User has to define a set of primitive symbols along with the required semantics. The symbols are assigned together to define legal sentences in the language for representing TRUE facts. New logical statements are formed from the existing ones. The statements which can be either TRUE or false but not both , are called propositions. A declarative sentence expresses a statement with a proposition as content; Example: The declarative "Cotton is white" expresses that Cotton is white. So, the sentence "Cotton is white" is a true statement. What is Propositional Logic PL? Propositional logic is a study of propositions. Each proposition has either a true or a false value but not both at a time. Propositions is represented by variables.

There are two types of propositions: It does not contain any other preposition. Rocky is a dog. It contains more than one prepositions. Surendra is a boy and he likes chocolate. Connectives and the truth tables of compound prepositions are given below: Truth table for negation:

KNOWLEDGE AND REPRESENTATION pdf

3: Knowledge representation and reasoning - Wikipedia

Some Long-Term Problems that need Knowledge Representation â€¢ Read a chapter in a textbook and answer questions at the end of the chapter â€¢ Einstein in a box: The.

Expert systems[edit] One of the first examples of an expert system was MYCIN , an application to perform medical diagnosis. In the MYCIN example, the domain experts were medical doctors and the knowledge represented was their expertise in diagnosis. Expert systems were first developed in artificial intelligence laboratories as an attempt to understand complex human decision making. Based on positive results from these initial prototypes, the technology was adopted by the US business community and later worldwide in the s. The Stanford heuristic programming projects led by Edward Feigenbaum was one of the leaders in defining and developing the first expert systems. History[edit] In the earliest days of expert systems there was little or no formal process for the creation of the software. Researchers just sat down with domain experts and started programming, often developing the required tools e. As expert systems moved from academic prototypes to deployed business systems it was realized that a methodology was required to bring predictability and control to the process of building the software. There were essentially two approaches that were attempted: Use conventional software development methodologies Develop special methodologies tuned to the requirements of building expert systems Many of the early expert systems were developed by large consulting and system integration firms such as Andersen Consulting. These firms already had well tested conventional waterfall methodologies e. One trend in early expert systems development was to simply apply these waterfall methods to expert systems development. Another issue with using conventional methods to develop expert systems was that due to the unprecedented nature of expert systems they were one of the first applications to adopt rapid application development methods that feature iteration and prototyping as well as or instead of detailed analysis and design. In the s few conventional software methods supported this type of approach. The final issue with using conventional methods to develop expert systems was the need for knowledge acquisition. Knowledge acquisition refers to the process of gathering expert knowledge and capturing it in the form of rules and ontologies. Knowledge acquisition has special requirements beyond the conventional specification process used to capture most business requirements. These issues led to the second approach to knowledge engineering: Organic Knowledge OK system can be abstractly seen as advisory board of very different even adversary educators with mechanisms of cooperation that work together to suggest to the student the next game after constantly observing him. In addition to using all the principles of SE, the organic SE adds another special layer of tools. It simulates and incorporates both the knowledge of the software engineer and the knowledge of the domain expert. In a nutshell the organic approach is treating the problem and the solution process as a living organism or ecology of organisms if more appropriate. The organic solution is non-algorithmic and evolving using feedback and data-to-knowledge mechanisms. It is as if the solution is like a child â€" in the beginning having no knowledge except some basic mechanisms needed for evolution , and by process of feedback and Darwinian natural selection the solution gradually evolves into better and better reactions using the growing body of knowledge. If at all, they are an ecological system of many different and sometimes contradictory experts, called organs. But in reality the OK systems are so much more in almost every aspect that it is more correct to say they are the fulfillment of the ES vision.

KNOWLEDGE AND REPRESENTATION pdf

4: Knowledge engineering - Wikipedia

Knowledge representation and reasoning are the parts of AI that are concerned with how an agent uses what it knows in deciding what to do. It is the study of thinking as a computational process. The book introduces the symbolic structures invented for representing knowledge and the computational processes devised for reasoning with those.

Full text of the second edition of Artificial Intelligence: The role of representations in solving problems The general framework for solving problems by computer is given in Figure 1. To solve a problem, the designer of a system must flesh out the task and determine what constitutes a solution; represent the problem in a language with which a computer can reason; use the computer to compute an output, which is an answer presented to a user or a sequence of actions to be carried out in the environment; and interpret the output as a solution to the problem. Knowledge is the information about a domain that can be used to solve problems in that domain. To solve many problems requires much knowledge, and this knowledge must be represented in the computer. As part of designing a program to solve problems, we must define how the knowledge will be represented. A representation scheme is the form of the knowledge that is used in an agent. A representation of some piece of knowledge is the internal representation of the knowledge. A representation scheme specifies the form of the knowledge. A knowledge base is the representation of all of the knowledge that is stored by an agent. A good representation scheme is a compromise among many competing objectives. A representation should be rich enough to express the knowledge needed to solve the problem. It should be easy to see the relationship between the representation and the domain being represented, so that it is easy to determine whether the knowledge represented is correct. A small change in the problem should result in a small change in the representation of the problem. Many different representation schemes have been designed. Many of these start with some of these objectives and are then expanded to include the other objectives. For example, some are designed for learning and then expanded to allow richer problem solving and inference abilities. Some representation schemes are designed with expressiveness in mind, and then inference and learning are added on. Some schemes start from tractable inference and then are made more natural, and more able to be acquired. Some of the questions that must be considered when given a problem or a task are the following: What is a solution to the problem? How good must a solution be? How can the problem be represented? What distinctions in the world are needed to solve the problem? What specific knowledge about the world is required? How can an agent acquire the knowledge from experts or from experience? How can the knowledge be debugged, maintained, and improved? How can the agent compute an output that can be interpreted as a solution to the problem? Is worst-case performance or average-case performance the critical time to minimize? Is it important for a human to understand how the answer was derived? These issues are discussed in the next sections and arise in many of the representation schemes presented later in the book.

5: What is a Knowledge Representation?

propositional account of the knowledge that the overall process exhibits, and b) independent of such external semantic attribution, play a formal but causal and essential role in engendering the behaviour that manifests that knowledge.".

The best option to understand what Knowledge Representation is simply to mention what it is intended for. Its mission is to make knowledge as explicit as possible. This is necessary because knowledge is stored in implicit form, i. To facilitate knowledge sharing it is necessary to make it explicit. Tacit Knowledge Tacit knowledge is what an agent obtains when it observes its environment and makes internal representations of what it perceives. Here "agent" stands for an entity capable of election. Agent choices are built from its internal representation, its model of the world. The model captures what there is and how it works, thus allowing the agent to predict what would happen if it does something or not, a complete view on that from a Systems Theory perspective [Klir92] is shown in Figure. Knowledge viewed from Systems Theory perspective In other words, tacit knowledge allows an agent to choose the best options that, hopefully, will help it achieve its goals. These goals are unimportant from a generic point of view. They might range from survival to booking a ticket, passing through getting a favourable transoceanic export rate, for instance. For social agents, tacit knowledge is also stored distributed in common habits established in a community [Polanyi97 , Tsoukas96]. The same principles apply, although from the perspective of the whole community as an agent. It can be also considered tacit because it is not explicitly represented in the community. It is distributed while agent act collectively, for example by imitation. This process is also known as socialisation [Nonaka95], a complete view of the tacit-explicit knowledge cycle is shown in Figure. Human natural languages are an example of tacit shared knowledge. Although a part of natural languages can be formalised, humans acquire natural language abilities mainly by imitation. The Knowledge Spiral Knowledge Sharing Knowledge is exchanged between social agents because this way each agent gets access to more than the knowledge it has been able to build up. Obviously, each agent, and the community as a whole, is then more prepared to make the correct choices. Agents have access to more than individual experiences and even unprecedented situations can be resolved satisfactorily. However, tacit social knowledge is exchanged inefficiently. The exchange mechanisms, e. This reduces knowledge propagation in space and time. To overcome these limitations, some agents have developed ways to make knowledge explicit and encode it in more perdurable form. Human languages are an example of this. They have written forms as perdurable encoding, always with some kind of physical support. Moreover, technology advances have also allowed perdurable encoding of languages oral form. Other kinds of agents have developed external and perdurable knowledge formats. For instance, cells DNA can be considered an encoding of how to reproduce a cell thus allowing its perpetuation. Generally, all these knowledge transmission mechanisms are studied by Semiotics. They are based on signs, their basic components. More details about Semiotics are shown in Figure and the Semiotics section. Knowledge exchange mechanisms based on signs and thus studied by Semiotics As has been said, explicit knowledge can overcome space-time limitations of tacit knowledge exchange. Perdurable encoding and transmission mechanisms allow that it can be acquired a long time after its encoding and far away its origination point. However, many of these representations carry interpretation ambiguities. This is because they are not wholly formalised. They are so expressive that some exchanged knowledge can be acquired at the destination leading to a different piece of knowledge. However, this cannot be considered a bad property. Ambiguity provides easy adaptation of the representation mechanisms to new situations. For instance, metaphor produces a new interpretation of a previous representation inside a particular context. Ambiguity also allows exploration of new possibilities because knowledge is not confined in a restrictive immutable form. A good example of this advantage is DNA. Information The previous risk of misunderstanding during knowledge exchange is the reason why this kind of encoded "knowledge" should not be considered knowledge. It is more appropriate to say that it is information. The encoded knowledge can be completely lost if the receiver agent cannot understand it. For instance, if two agents exchange a written message but the second one does not understand the used language, nothing of the originally codified knowledge can be retrieved. Therefore, to be completely strict, there does not exist more knowledge than tacit knowledge. Information is the small part of it an agent is able to articulate [Stenmark02]. When an agent receives some information, it uses its tacit knowledge to interpret it and, possibly, this may lead to a change in the tacit knowledge it possesses. However, this categorisation is in practice relaxed. There are different kinds of information and normally when the exchanged information is rich enough it is considered knowledge. Rich information has embedded enough contextual information to facilitate its full interpretation. Moreover, some encoding restrictions must be imposed in order to guarantee, to some extent, a final interpretation near to the original encoded knowledge. In the opposite side of information rich enough to be considered knowledge, there is data. It is de-contextualised information, i. Knowledge Formalisation As mentioned before, despite ambiguity advantages, sometimes it is necessary to exchange knowledge as reliably as possible. This has been a clear requirement in human societies for a long time. Indeed, Socrates can be considered a starting point in this formal knowledge exchange research, but roots could be extended even before. From these remote times, humans have developed many representation formalisms. All them define their own set of shared constraints that must be incorporated as tacit knowledge in knowledge emitters and receiver. Once a formalism has been incorporated in the tacit knowledge of a community, this community can share information in a so direct and rich way that it can be considered knowledge exchange. These formalisms can be very simple, for instance defining a set of reserved natural language words with an agreed community meaning. Then, community agents can share knowledge interchanging messages that use these agreed words. This is an example of a purely textual formalism, but there are also graphical ones. They are called diagrammatic formalisms and they are quite simple and easier to interpret, for instance Conceptual Maps [Novak84]. However, the more powerful formalisms use techniques that are more sophisticated. They are mainly based on mathematics, philosophy and cognitive science. These disciplines provide basic ideas of how we perceive and model the world. Thus, they set a base that we naturally share, although not in an obvious way. Mathematics provides a compact set of principles widely shared among human society. This shared common base allows the construction of very powerful expressions. These expressions have clear meaning for those that incorporate the used part of the shared mathematical base into their tacit knowledge. Meanwhile, philosophy studies the nature of knowledge, how we create and manage it. Some techniques have been developed that capture a part of our brain operation. Most of them use mathematical tools to some extent. For instance, logic and ontology are two building blocks of Knowledge Representation. On the other hand, there are also attempts to explain mathematics from a philosophical point of view [Lofting]. Despite all the possibilities of advanced representation formalism, it is important to remark that tacit knowledge is richer than any description of it. Knowledge Representations As has been shown along the previous sections, the final objective of knowledge representations is to make knowledge explicit. Knowledge can be shared less ambiguously in its explicit form and this became especially important when machines started to be applied to facilitate knowledge management. Nowadays, Knowledge Representation is a multidisciplinary field that applies theories and techniques from: Therefore, Knowledge Representation can be defined as the application of logic and ontology to the task of constructing computable models of some domain [Sowa00]. Logic and Ontology provide the formalisation mechanisms required to make expressive models easily sharable and computer aware. Finally, thanks to computational resources, great quantities of knowledge expressed this way can be automated. Thus, the full potential of knowledge accumulations can be exploited. However, computers play only the role of powerful processors of more or less rich information sources. The final interpretation of the results is carried out by the agents that motivate this processing, in this case human users of the knowledge management systems. At this point, it is important to remark that the possibilities of the application of actual Knowledge Representation techniques are enormous. Knowledge is always more than the sum of its parts and Knowledge Representation provides the tools needed to manage accumulations of knowledge and the World Wide Web is becoming the biggest accumulation of knowledge ever faced by humanity. These possibilities will be more deeply explored in the next State of the Art sections, devoted to Web Technologies and the Semantic Web. Principles In addition to the previous definition, Knowledge Representation can be also described by the five fundamental roles that it plays in artificial intelligence; they are the Knowledge

KNOWLEDGE AND REPRESENTATION pdf

Representation principles [Davis93]: A knowledge representation is a surrogate: Symbols are surrogates for the external things. Symbols and links between them form a model of the external system that can be manipulated to simulate it or reason about it. A knowledge representation is a set of ontological commitments: Ontology is the study of existence. Thus, ontology determines the categories of things that exist or may exist in an application domain. Those categories set the ontological commitments of the application designer or knowledge engineer.

6: Knowledge structuring for learning

Knowledge representation is at the very core of a radical idea for understanding intelligence. Instead of trying to understand or build brains from the bottom up, its goal is to understand and build intelligent behavior from the top down, putting the focus on what an agent needs to know in order to behave intelligently, how this knowledge can be represented symbolically, and how automated.

A KR is a Surrogate Any intelligent entity that wishes to reason about its world encounters an important, inescapable fact: A program or person engaged in planning the assembly of a bicycle, for instance, may have to reason about entities like wheels, chains, sprockets, handle bars, etc. This unavoidable dichotomy is a fundamental rationale and role for a representation: Operations on and with representations substitute for operations on the real thing, i. In this view reasoning itself is in part a surrogate for action in the world, when we can not or do not yet want to take that action. The first question about any surrogate is its intended identity: There must be some form of correspondence specified between the surrogate and its intended referent in the world; the correspondence is the semantics for the representation. The second question is fidelity: What attributes of the original does it capture and make explicit, and which does it omit? Perfect fidelity is in general impossible, both in practice and in principle. It is impossible in principle because any thing other than the thing itself is necessarily different from the thing itself in location if nothing else. Put the other way around, the only completely accurate representation of an object is the object itself. All other representations are inaccurate; they inevitably contain simplifying assumptions and possibly artifacts. Two minor elaborations extend this view of representations as surrogates. First, it appears to serve equally well for intangible objects as it does for tangible objects like gear wheels: Representations function as surrogates for abstract notions like actions, processes, beliefs, causality, categories, etc. Second, formal objects can of course exist inside the machine with perfect fidelity: Mathematical entities, for example, can be captured exactly, precisely because they are formal objects. Since almost any reasoning task will encounter the need to deal with natural objects i. Imperfect surrogates mean incorrect inferences are inevitable Two important consequences follow from the inevitability of imperfect surrogates. One consequence is that in describing the natural world, we must inevitably lie, by omission at least. At a minimum we must omit some of the effectively limitless complexity of the natural world; our descriptions may in addition introduce artifacts not present in the world. The second and more important consequence is that all sufficiently broad-based reasoning about the natural world must eventually reach conclusions that are incorrect, independent of the reasoning process used and independent of the representation employed. Sound reasoning cannot save us: If the world model is somehow wrong and it must be some conclusions will be incorrect, no matter how carefully drawn. A better representation cannot save us: The significance of the error may of course vary; indeed much of the art of selecting a good representation is in finding one that minimizes or perhaps even eliminates error for the specific task at hand. But the unavoidable imperfection of surrogates means we can supply at least one guarantee for any entity reasoning in any fashion about the natural world: Drawing only sound inferences thus does not free reasoning from error; it can only ensure that inference is not the source of that error. Given that broad based reasoning will inevitably be wrong, the step from sound inference to other models of inference is thus not a move from total accuracy to error, but is instead a question of balancing off the possibility of one more source of error against the gains e. We do not suggest that unsound reasoning ought to be embraced casually, but do claim that, given the inevitability of error even with sound reasoning, it makes sense to evaluate pragmatically the relative costs and benefits that come from using both sound and unsound reasoning methods. A KR is a Set of Ontological Commitments If, as we have argued, all representations are imperfect approximations to reality, each approximation attending to some things and ignoring others, then in selecting any representation we are in the very same act unavoidably making a set of decisions about how and what to see in the world. That is, selecting a representation means making a set of ontological commitments. It is unavoidably so because of the inevitable imperfections of representations. It is usefully so because judicious selection of commitments provides the opportunity to focus attention on aspects of the world we believe to be relevant. The focusing effect is an essential part of what a representation offers, because the complexity of the natural world is overwhelming. We and our reasoning machines need guidance in deciding what in the world to attend to and what to ignore. The glasses supplied by a representation can provide that guidance: In telling us what and how to see, they allow us to cope with what would otherwise be untenable complexity and detail. Hence the ontological commitment made by a representation can be one of the most important contributions it offers. There is a long history of work attempting to build good ontologies for a variety of task domains, including early work on an ontology for liquids [12] , the lumped element model widely used in representing electronic circuits e. Each of these offers a way to see some part of the world. The lumped element model, for instance, suggests that we think of circuits in terms of components with connections between them, with signals flowing instantaneously along the connections. This is a useful view, but not the only possible one. A different ontology arises if we need to attend to the electrodynamics in the device: Ontologies can of course be written down in a wide variety of languages and notations e. Simply put, the important part is notions like connections and components, not whether we choose to write them as predicates or LISP constructs. The commitment we make by selecting one or another ontology can produce a sharply different view of the task at hand. Consider the difference that arises in selecting the lumped element view of a circuit rather than the electrodynamic view of the same device. As a second example, medical diagnosis viewed in terms of rules e. Commitment begins with the earliest choices The INTERNIST example also demonstrates that there is significant and unavoidable ontological commitment even at the level of the familiar representation technologies. Logic, rules, frames, etc. Logic, for instance, involves a fairly minimal commitment to viewing the world in terms of individual entities and relations between them. Rule-based systems view the world in terms of attribute-object-value triples and the rules of plausible inference that connect them, while frames have us thinking in terms of prototypical objects. Each of these thus supplies its own view of what is important to attend to, and each suggests, conversely, that anything not easily seen in those terms may be ignored. This is of course not guaranteed to be correct, since anything ignored may later prove to be relevant. But the task is hopeless in principle--every representation ignores something about the world--hence the best we can do is start with a good guess. The existing representation technologies supply one set of guesses about what to attend to and what to ignore. Selecting any of them thus involves a degree of ontological commitment: The commitments accumulate in layers The ontologic commitment of a representation thus begins at the level of the representation technologies and accumulates from there. Additional layers of commitment are made as we put the technology to work. At the most fundamental level, the decision to view diagnosis in terms of frames suggests thinking in terms of prototypes, defaults, and a taxonomic hierarchy. But prototypes of what, and how shall the taxonomy be organized? An early description of the system [21] shows how these questions were answered in the task at hand, supplying the second layer of commitment: The prototypes are thus intended to capture prototypical diseases e. This is a sensible and intuitive set of choices but clearly not the only way to apply frames to the task; hence it is another layer of ontological commitment. At the third and in this case final layer, this set of choices is instantiated: Ontologic questions that arise even at this level can be quite fundamental. Consider for example determining which of the following are to be considered diseases i. The ontologic commitment here is sufficiently obvious and sufficiently important that it is often a subject of debate in the field itself, quite independent of building automated reasoners. Similar sorts of decisions have to be made with all the representation technologies, because each of them supplies only a first order guess about how to see the world: Similarly logic tells us to view the world in terms of individuals and relations, but does not specify which individuals and relations to use. Commitment to a particular view of the world thus starts with the choice of a representation technology, and accumulates as subsequent choices are made about how to see the world in those terms. A KR is not a data structure Note that at each layer, even the first e. Part of what makes a language representational is that it carries meaning [13] , i. That correspondence in turn carries with it constraint. A semantic net, for example, is a representation, while a graph is a data structure. They are different kinds of entities, even though one is invariably used to implement the other, precisely because the net has should have a semantics. That semantics will be manifest in part because it constrains the network topology: While every representation must be implemented in the machine by some data structure, the representational property is in the correspondence to something in the world and in the constraint that correspondence imposes. This role comes about because the initial conception of a representation is typically motivated by some insight indicating how people reason intelligently, or by some belief about what it means to reason intelligently at all. The theory is fragmentary in two distinct senses: Where the sanctioned inferences indicate what can be inferred at all, the recommended inferences are concerned with what should be inferred. Guidance is needed because the set of sanctioned inferences is typically far too large to be used indiscriminantly. Where the ontology we examined earlier tells us how to see, the recommended inferences suggest how to reason. We begin with the first of these components, examining two of several fundamentally different conceptions of intelligent reasoning that have been explored in AI. Those conceptions and their underlying assumptions demonstrate the broad range of views on the question and set important context for the remaining components. What is intelligent reasoning? What are the essential, defining properties of intelligent reasoning? As a consequence of the relative youth of AI as a discipline, insights about the nature of intelligent reasoning have often come from work in other fields. Five fields--mathematical logic, psychology, biology, statistics, and economics--have provided the inspiration for five distinguishable notions of what constitutes intelligent reasoning Table I. One view, historically derived from mathematical logic, makes the assumption that intelligent reasoning is some variety of formal calculation, typically deduction; the modern exemplars of this view in AI are the logicists. A second view, rooted in work in psychology, sees reasoning as a characteristic human behavior and has given rise to both the extensive work on human problem solving and the large collection of knowledge-based systems. Researchers working on several varieties of connectionism are the current descendants of this line of work. A fourth approach, derived from probability theory, adds to logic the notion of uncertainty, yielding a view in which reasoning intelligently means obeying the axioms of probability theory. A fifth view, from economics, adds the further ingredient of values and preferences, leading to a view of intelligent reasoning defined by adherence to the tenets of utility theory. Views of intelligent reasoning and their intellectual origins Briefly exploring the historical development of the first two of these views will illustrate the different conceptions they have of the fundamental nature of intelligent reasoning and will demonstrate the deep-seated differences in mindset that arise as a consequence. By the time of Leibnitz the agenda is quite specific and telling: In the 19th century Boole provided the basis for propositional calculus in his "Laws of Thought;" later work by Frege and Peano provided additional foundation for the modern form of predicate calculus.

7: Knowledge Representation

A Representation is an assertion of fact true on the date that the party makes the Representation. A Warranty is a promise of indemnity if the Representation is inaccurate. Buyer: The Buyer wants comprehensive representations and warranties that are not qualified by knowledge or materiality.

These systems featured data structures for planning and decomposition. The system would begin with a goal. It would then decompose that goal into sub-goals and then set out to construct strategies that could accomplish each subgoal. However, the amorphous problem definitions for systems such as GPS meant that they worked only for very constrained toy domains e. In order to tackle non-toy problems, AI researchers such as Ed Feigenbaum and Frederick Hayes-Roth realized that it was necessary to focus systems on more constrained problems. It was the failure of these efforts that led to the cognitive revolution in psychology and to the phase of AI focused on knowledge representation that resulted in expert systems in the s and 80s, production systems , frame languages , etc. Rather than general problem solvers, AI changed its focus to expert systems that could match human competence on a specific task, such as medical diagnosis. Expert systems gave us the terminology still in use today where AI systems are divided into a Knowledge Base with facts about the world and rules and an inference engine that applies the rules to the knowledge base in order to answer questions and solve problems. In these early systems the knowledge base tended to be a fairly flat structure, essentially assertions about the values of variables used by the rules. A frame is similar to an object class: It is an abstract description of a category describing things in the world, problems, and potential solutions. Frames were originally used on systems geared toward human interaction, e. Frames were good for representing the real world, described as classes, subclasses, slots data values with various constraints on possible values. Rules were good for representing and utilizing complex logic such as the process to make a medical diagnosis. Integrated systems were developed that combined Frames and Rules. KEE had a complete rule engine with forward and backward chaining. It also had a complete frame based knowledge base with triggers, slots data values , inheritance, and message passing. Although message passing originated in the object-oriented community rather than AI it was quickly embraced by AI researchers as well in environments such as KEE and in the operating systems for Lisp machines from Symbolics , Xerox , and Texas Instruments. At the same time as this was occurring, there was another strain of research which was less commercially focused and was driven by mathematical logic and automated theorem proving. This reasoner is called the classifier. In this way the classifier can function as an inference engine, deducing new facts from an existing knowledge base. The classifier can also provide consistency checking on a knowledge base which in the case of KL-ONE languages is also referred to as an Ontology. One of the first realizations learned from trying to make software that can function with human natural language was that humans regularly draw on an extensive foundation of knowledge about the real world that we simply take for granted but that is not at all obvious to an artificial agent. Basic principles of common sense physics, causality, intentions, etc. An example is the frame problem , that in an event driven logic there need to be axioms that state things maintain position from one moment to the next unless they are moved by some external force. In order to make a true artificial intelligence agent that can converse with humans using natural language and can process basic statements and questions about the world, it is essential to represent this kind of knowledge. Cyc established its own Frame language and had large numbers of analysts document various areas of common sense reasoning in that language. The knowledge recorded in Cyc included common sense models of time, causality, physics, intentions, and many others. Currently one of the most active areas of knowledge representation research are projects associated with the semantic web. The semantic web seeks to add a layer of semantics meaning on top of the current Internet. Rather than indexing web sites and pages via keywords, the semantic web creates large ontologies of concepts. Searching for a concept will be more effective than traditional text only searches. Frame languages and automatic classification play a big part in the vision for the future semantic web. The automatic classification gives developers technology to provide order on a constantly evolving network of knowledge. Defining ontologies that are static and incapable of evolving on the fly would be very limiting for

Internet-based systems. The classifier technology provides the ability to deal with the dynamic environment of the Internet. The Resource Description Framework RDF provides the basic capability to define classes, subclasses, and properties of objects. The Web Ontology Language OWL provides additional levels of semantics and enables integration with classification engines. The justification for knowledge representation is that conventional procedural code is not the best formalism to use to solve complex problems. Knowledge representation makes complex software easier to define and maintain than procedural code and can be used in expert systems. For example, talking to experts in terms of business rules rather than code lessens the semantic gap between users and developers and makes development of complex systems more practical. Knowledge representation goes hand in hand with automated reasoning because one of the main purposes of explicitly representing knowledge is to be able to reason about that knowledge, to make inferences, assert new knowledge, etc. Virtually all knowledge representation languages have a reasoning or inference engine as part of the system. The ultimate knowledge representation formalism in terms of expressive power and compactness is First Order Logic FOL. There is no more powerful formalism than that used by mathematicians to define general propositions about the world. However, FOL has two drawbacks as a knowledge representation formalism: First order logic can be intimidating even for many software developers. Languages which do not have the complete formal power of FOL can still provide close to the same expressive power with a user interface that is more practical for the average developer to understand. The issue of practicality of implementation is that FOL in some ways is too expressive. With FOL it is possible to create statements e. Thus, a subset of FOL can be both easier to use and more practical to implement. This was a driving motivation behind rule-based expert systems. The history of most of the early AI knowledge representation formalisms; from databases to semantic nets to theorem provers and production systems can be viewed as various design decisions on whether to emphasize expressive power or computability and efficiency. It is a set of ontological commitments, i. In what terms should I think about the world? It is a fragmentary theory of intelligent reasoning, expressed in terms of three components: It is a medium for pragmatically efficient computation, i. One contribution to this pragmatic efficiency is supplied by the guidance a representation provides for organizing information so as to facilitate making the recommended inferences. It is a medium of human expression, i. Knowledge representation and reasoning are a key enabling technology for the Semantic web. Languages based on the Frame model with automatic classification provide a layer of semantics on top of the existing Internet. Rather than searching via text strings as is typical today, it will be possible to define logical queries and find pages that map to those queries. Classifiers focus on the subsumption relations in a knowledge base rather than rules. A classifier can infer new classes and dynamically change the ontology as new information becomes available. This capability is ideal for the ever-changing and evolving information space of the Internet. The Resource Description Framework RDF provides the basic capabilities to define knowledge-based objects on the Internet with basic features such as Is-A relations and object properties. What is the underlying framework used to represent knowledge? Semantic networks were one of the first knowledge representation primitives. Also, data structures and algorithms for general fast search. In this area, there is a strong overlap with research in data structures and algorithms in computer science. In early systems, the Lisp programming language, which was modeled after the lambda calculus , was often used as a form of functional knowledge representation. Frames and Rules were the next kind of primitive. Frame languages had various mechanisms for expressing and enforcing constraints on frame data. All data in frames are stored in slots. Slots are analogous to relations in entity-relation modeling and to object properties in object-oriented modeling. The most well known example is Prolog, but there are also many special purpose theorem proving environments. These environments can validate logical models and can deduce new theories from existing models. Essentially they automate the process a logician would go through in analyzing a model. Theorem proving technology had some specific practical applications in the areas of software engineering. For example, it is possible to prove that a software program rigidly adheres to a formal logical specification. This is also known as the issue of reflection in computer science. It refers to the capability of a formalism to have access to information about its own state. An example would be the meta-object protocol in Smalltalk and CLOS that gives developers run time access to the class objects and enables them to dynamically redefine the structure of the knowledge base even at run time. Meta-representation means the knowledge representation language is itself expressed in that language. For example, in most Frame based environments all frames would be instances of a frame class. That class object can be inspected at run time, so that the object can understand and even change its internal structure or the structure of other parts of the model. In rule-based environments, the rules were also usually instances of rule classes. Part of the meta protocol for rules were the meta rules that prioritized rule firing. Traditional logic requires additional axioms and constraints to deal with the real world as opposed to the world of mathematics. Also, it is often useful to associate degrees of confidence with a statement. This was one of the early innovations from expert systems research which migrated to some commercial tools, the ability to associate certainty factors with rules and conclusions. Later research in this area is known as fuzzy logic. Universals are general statements about the world such as "All humans are mortal". Facts are specific examples of universals such as "Socrates is a human and therefore mortal". In logical terms definitions and universals are about universal quantification while facts and defaults are about existential quantifications. All forms of knowledge representation must deal with this aspect and most do so with some variant of set theory, modeling universals as sets and subsets and definitions as elements in those sets. Non-monotonic reasoning allows various kinds of hypothetical reasoning. The system associates facts asserted with the rules and facts used to justify them and as those facts change updates the dependent knowledge as well. In rule based systems this capability is known as a truth maintenance system.

8: What is Knowledge Representation? - Definition from Techopedia

John F. Sowa, Knowledge Representation: Logical, Philosophical, and Computational Foundations, Brooks Cole Publishing Co., Pacific Grove, CA, Â©Actual.

9: Knowledge Representation Book

Knowledge Representation furthers research on formalisms and tools for semantically precise and computable knowledge representation in relation to biomedicine.

KNOWLEDGE AND REPRESENTATION pdf

Open secrets Stephen Yenser Physical chemistry multiple choice questions with answers Learn Spanish the Fast and Fun Way w/4 CDs Aerobella (Little Ones) Prejudice and the teaching of history, by C. L. Hannam. Memories of the Baron de Tott, on the Turks and the Tartars. The Video Avenger (Twist-A-Plot #7) Berlin since the walls end Dbq 19 causes of world war 1 A World of Art and Museums Monster Crappie (Instructional Fishing Series, Volume 5) Securities Exchange Act of 1934 Development bank of ethiopia lease financing The Character of kinship. The girl on the train kickass Appendix analogies and quantitative comparisons Suffering and the soul Set free by forgiveness 365 ways to survive as a parent Memorial culture : the material response to loss Sorcerers Moon (Boreal Moon Tale Series) Only Couples Need Apply (Your Truly) The suborder ensifera Rachmaninoff piano concerto no 2 in c minor The certified six sigma yellow belt handbook Point of power 31: I need not worry, God is in control V. 2. Morphology. The White Horse Battery Critical thinking tactics for nurses 365 Ways To Live To 100 How modern interrogators have adapted to Miranda Iata dangerous goods list How Old Brye came down from the mountain Writing ing and research Basket of Secrets (Heartsong Presents #620) Foreword : A brief history of ornamental horticulture State level public enterprises in Sikkim Manual of Strabismus Surgery The abscence of Mr. Glass G.K. Chesterton Lowndes and Rudolf