Chapter 4
Knowledge Representation

The explanation of the term *knowledge* very often involves using the term *information*. Loosely stated, knowledge is information in support of or in conflict with a certain hypothesis, or it serves to resolve a problem or answer a specific question. Specific knowledge that results from information processing may be either expected or it may be new and surprising. The initially gathered information is often fragmented and unstructured, and in that form it is not suitable for further exchange and processing across different systems. Moreover, one does not usually have an a priori understanding of what the atoms of knowledge are, how they are connected, and how one can retrieve or deduce new knowledge from them. In order to answer some of these important questions, the next section begins by examining different definitions of knowledge, followed by a discussion of knowledge organization, and concluding with practical applications of knowledge representations.

4.1 Definition of Knowledge

Knowledge and *concept* are among the most abstract terms in human vocabulary. Similarly to the term *concept*, all the characteristics of knowledge cannot be captured within a single definition. The ancient philosopher Plato described knowledge as "*justified true belief*". According to Plato, a person knows a proposition to be true if (and only if) he or she believes in the truth of the proposition and at the same time has justification for doing so. In the following centuries, many definitions and theories of knowledge were proposed; however, not a single one has been widely agreed upon.

Similarly to the definitions of concepts, attempts to define knowledge were also made in other disciplines besides philosophy. These definitions are often tailored to meet the specificities of the field, which also involves the use of specific terminology. For example, from the viewpoint of cognitive psychology, knowledge is considered a "*cognitive subject matter content*" (Merrill 2000). When it comes to computer science, a very pragmatic position towards explaining knowledge is

G. Jakus et al., *Concepts, Ontologies, and Knowledge Representation*,
SpringerBriefs in Computer Science, DOI: 10.1007/978-1-4614-7822-5_4,

taken, as it is often considered that knowledge must have an applicable, functional or even predictive value. Such position is reflected also in the following definitions of knowledge from the field of computer science:

- "*Whatever can be ascribed to an agent, such that its behavior can be computed according to the principle of rationality*" (Newell 1982);
- "*Whole body of data and information that people bring to bear to practical use in action, in order to carry out tasks and create new information*" (Schreiber et al. 2000);
- "*Conceptual models of systems and principles*" that explain "*functioning, causes and effects, form, features and may have a predictive nature*" (Halladay and Milligan 2004).

The definition of knowledge is partly demanding because it depends on the context. This can be illustrated, for example, by taking the definition of (Schreiber et al. 2000) which defines knowledge in terms of data and information. It is, however, difficult to make a clear distinction between data ("*the uninterpreted signals that reach our senses*"), information ("*data equipped with meaning*") and knowledge itself, because such distinction depends on the respective circumstances. In certain situations, one person's knowledge can, namely, represent completely meaningless data for someone else. For example, all knowledge a skilled chess player has on opening strategies makes very little sense to someone who does not know how the individual chess pieces move.

If it is so difficult to define knowledge, why should one take the trouble of defining it? Even without an explicit definition, knowledge can be recognized from observing the activity of entities (human or software agents) that are capable of particular actions, and the effects of such activity (Newell 1982; Guarino and Giaretta 1995; Schreiber et al. 2000). For example, we are able to differentiate between knowledgeable and ignorant people by simply observing their actions to achieve a pursued goal. When, for example, witnessing a person waving around with a computer mouse, one can conclude that he or she does not have the knowledge on how to use this device to interact with the computer.

Despite the diversity of views regarding the nature of knowledge, there seems to be a broader agreement on distinguishing different types of knowledge. The matter was first discussed by philosophers Bertrand Russell distinguishing between *declarative knowledge* and *knowledge by acquaintance* (Russell 1912), and Gilbert Ryle distinguishing between *declarative* and *procedural knowledge* (Ryle 1949).

Declarative knowledge describes facts or the understanding that something is true. Such knowledge can be expressed verbally, for example using declarative sentences, such as "*London hosted the* 2012 *Summer Olympic Games*". Declarative knowledge is also referred to as descriptive or propositional knowledge or "knowledge that".

Procedural knowledge, on the other hand, refers to the ability or possession of a skill to perform a task in an efficient way. It may not always be possible to verbalize procedural knowledge, as it can sometimes only be recognized "in action" or by its effect. The difference between declarative and procedural knowledge is, for example, reflected in the ability to swim. You may know how to move arms and legs, but until you actually apply this declarative knowledge in (deep enough) water, you cannot actually consider yourself a swimmer. Procedural knowledge is also referred to as imperative knowledge or "know-how".

Knowledge by acquaintance refers to familiarity with someone or something gained through experience. For example, in order *to know* human feelings such as love or fear, one must experience them first. The meaning of the verb "*to know*" used in the previous sentence refers to knowledge by acquaintance. Knowledge by acquaintance is also referred to as personal knowledge or "knowing of".

The use of knowledge adds a great value to computer systems. However, this introduces the question of how to represent and organize knowledge in a uniform manner to make it suitable for use and sharing. In order to provide answers to this important issue, knowledge representation and organization issues are discussed next.

4.2 Representing and Organizing Knowledge

4.2.1 Knowledge Representation

The term "knowledge representation" refers to "*using formal symbols to represent a collection of propositions believed by some putative agent*" (Brachman and Levesque 2004). As such, knowledge representation acts as an internal representation of reality inside an intelligent agent. As every practical instance of knowledge representation contains only a limited number of propositions about the world, it can only approximate reality and, in addition, it inevitably gives more focus to some things and at the same time neglects others. By choosing a specific type of knowledge representation, the intelligent agent is, therefore, bound to use a specific set of terms which determine how and what to perceive of reality. These terms are also referred to as "ontological commitments" (Davis et al. 1993).

As already discussed in the chapter on ontologies, providing a committing set of terms used to represent reality is, in fact, the task of ontologies. For that reason, ontologies can be considered the heart of every knowledge representation. However, as an actual ontology cannot represent all the propositions about the world (as their number is quite possibly infinite), in most cases, the knowledge contained in an ontology is not sufficient for an intelligent agent to carry out its tasks. To be of practical use for knowledge representation, ontologies therefore must not only serve as a source of explicitly recorded knowledge, but must also provide the means to create new knowledge by manipulating the existing knowledge through a process referred to as *reasoning*.

To establish a basis for reasoning, ontologies must be "embedded" into a suitable framework which includes a formal system of logic and an efficient computational environment (Davis et al. 1993; Sowa 2010). Besides supplying the symbols and formal structure for representation, *logic* also provides the rules and operations that can be used on symbols to create new knowledge through the process of reasoning. *Computational environment* provides the means for computationally efficient reasoning and use of represented knowledge in practical applications.

4.2.1.1 Formalisms for Representing Knowledge

Every day, we exchange our knowledge with each other and with machines. We do not perform this activity by actually exchanging concrete objects with each other or with a machine, but by exchanging surrogate representations. An important role of knowledge representations is, therefore, also to provide a medium for human expression and communication (Davis et al. 1993). The most widespread medium of such type is *natural language*. Although natural languages are very expressive, they are, however, often ambiguous, inconsistent, underspecified and difficult to model, and are as such, therefore, not appropriate to be used for the representation of knowledge and its manipulation within computer systems. For this purpose, other more "computer friendly" formalisms are much more appropriate, for example semantic networks, frames, description logics, conceptual graphs and fuzzy logic.

Semantic networks are graphical structures that are particularly suitable for representing static world knowledge. The nodes in semantic networks represent concepts or objects connected with binary semantic relations acting as graph edges.

Frames are knowledge representation structures influenced by the organization of human memory. A frame is a "*remembered framework*" suitable for representing "*stereotyped situations*" (like, for example, going to a birthday party) (Minsky 1975). A frame consists of the so-called *slots* which can be interpreted as some sort of properties or attributes that can be assigned values or references to other frames. The content of the slots representing facts is fixed, while the content of other slots, also referred to as "terminals", can be customized to meet the circumstances of the reality that is being represented.

Frames can be considered as an evolution of semantic networks. As slots can contain references to other frames, the resulting associations among the frames form a network of nodes (frames) and relations (linking slots). Nevertheless, frames build on classic semantic networks by introducing default slot values and procedures that enable the dynamic assignment of slot values under particular conditions and circumstances. A frame representing knowledge on one of the authors of this survey is shown in Fig. 4.1.

Description logics are a family of knowledge representation languages that can be used to *describe* an application domain through formal, *logic*-based semantics (Baader et al. 2008). By providing the latter, description logics build upon semantic networks and frames. Description logics have good computational properties, which in turn enables efficient reasoning.

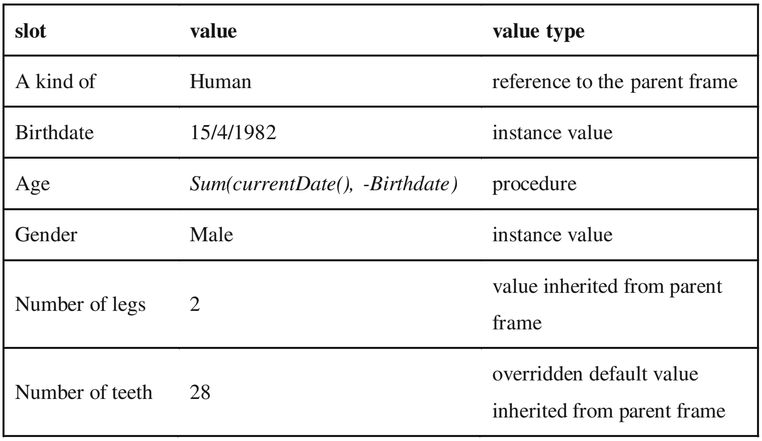

slot	value	value type
A kind of	Human	reference to the parent frame
Birthdate	15/4/1982	instance value
Age	*Sum(currentDate(), -Birthdate)*	procedure
Gender	Male	instance value
Number of legs	2	value inherited from parent frame
Number of teeth	28	overridden default value inherited from parent frame

Fig. 4.1 Representing knowledge using a frame

Conceptual graphs (Sowa 1984) are a logical formalism based on existential graphs (Peirce 1909) and semantic networks. Conceptual graphs are used to represent knowledge in a "*logically precise, humanly readable, and computationally tractable*" form that can, in addition, be represented in different notations including human language.

Fuzzy logic is an extension of the traditional logical systems offering the framework for knowledge representation in environments characterized by uncertainty and imprecision. In contrast to other typical approaches to knowledge representation which only enable exact knowledge representation and drawing conclusions based on the latter, fuzzy logic functions in a way that is much more similar to that of the human mind.

4.2.2 Knowledge Organization

When representing knowledge, two basic levels of representation can be recognized corresponding to the content and the structure (or the organization) of knowledge (van Geenen 2004; Stillings et al. 1995). The content of knowledge refers to using formal symbols to represent the concepts of reality. Such content is organized by means of *knowledge structures*. The two levels of knowledge representation can also be looked upon as a distinction between the semantics and the syntax of knowledge representations. While syntax deals with the way of arranging symbols used to represent knowledge (i.e. knowledge structures), semantics addresses the meaning of these symbols and their arrangements (i.e. knowledge content).

The difficulty connected with defining knowledge is to a large extent related to the content of knowledge, as this is domain- and context-dependant. On the other hand, knowledge structures can, in general, be applied to a wide range of domains (Schreiber et al. 2000). Some examples of such knowledge structures are presented in the continuation of this survey.

4.2.2.1 Imitation of Human Memory

Perhaps the most natural approach to the organization of knowledge is to organize it similarly to the way in which it is organized in human memory. As proposed in Merrill (2000), knowledge can be represented by using the so-called *knowledge objects* which consist of *knowledge components*. From this perspective, four types of knowledge objects are essential for organizing knowledge:

- *entities*—representing things or objects;
- *actions*—representing procedures that can be performed on, to or with entities or their parts;
- *processes*—representing the events that occur as a result of an action, and
- *properties*—representing descriptors for entities, actions, or processes.

Knowledge components are, in general, used to name, describe, or illustrate the parent knowledge object or its parts. In addition, knowledge components also define the relations of a knowledge object to the components of other objects. For example, the components of a knowledge object representing a process (or knowledge about how something works) describe:

- the conditions for executing the process (e.g. matching particular property values),
- the outcomes of process execution (e.g. property values that are changed as a result of the process) and
- other processes that are potentially triggered by the respective process.

A generic knowledge structure that can be used to represent a process is illustrated in Fig. 4.2.

The application of the presented model can be explained with (the activity of) illuminating a room. The activity starts by acting on the light switch which serves as a controller of the entity "the light". The entity contains the property "state" with two possible assigned values, i.e. "on" or "off", which can be portrayed, respectively, with a bulb with or without emerging light rays. The value of the property "state" is the condition for the execution of the process of "turning on the light". If the value of this property at the moment of triggering the process on the controller is "on", then the process will not be executed. In the case when the value of the property "state" is "off", the process would alter it to "on". In addition, the process of turning on the light could also trigger another process, for example, turning on another light in the room.

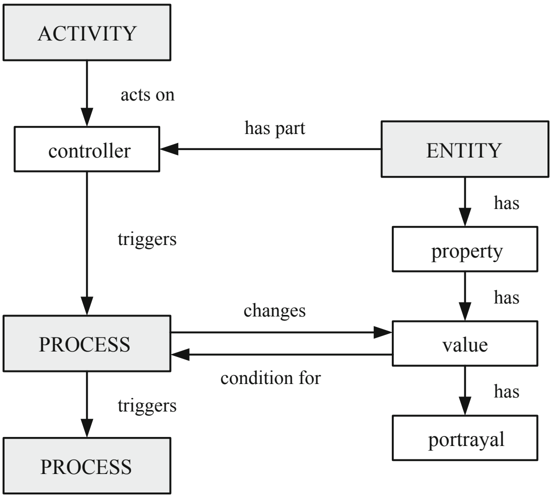

Fig. 4.2 Process knowledge structure (Merrill 2000)

4.2.2.2 WordNet

WordNet was already briefly presented in the chapter on concepts. There we mentioned that this lexical database consists of sets of synonym words, referred to as synsets, which express concepts lexically. Since individual words can have more than one meaning (phenomena referred to as *polysemy*), they can appear in more than one synset. Polysemic relations between the individual senses of a word have an important impact on the structure of WordNet, as they impose the organization of synsets in the form of a semantic network.

However, polysemy is not the only relation that affects the topology of WordNet. By its nature, polysemy links lexical units representing concepts rather than linking concepts directly. The relations among concepts in WordNet are established primarily by arranging the concepts into a hierarchy in which a concept that is subordinate to another concept is considered more specific. As such, it inherits general knowledge from its superordinate concept, so that it only needs to enclose the respective specific knowledge. Such organization of knowledge is influenced by the organization of human semantic memory (Fellbaum 2010), which is in turn similar to the approach presented in the previous section.

In the case of nouns and verbs, the described hierarchy is based on the hypernym relation, where each subordinate concept is “a kind of” its superordinate concept. Figure 4.3 shows the hypernymic hierarchy for the concept *homo*. The synonyms forming a synset are presented in strong text, while their definitions (descriptions) are written in parentheses.

With some proposed ways of characterizing knowledge organization already presented, the following section focuses on various possibilities of knowledge use in intelligent computer systems.

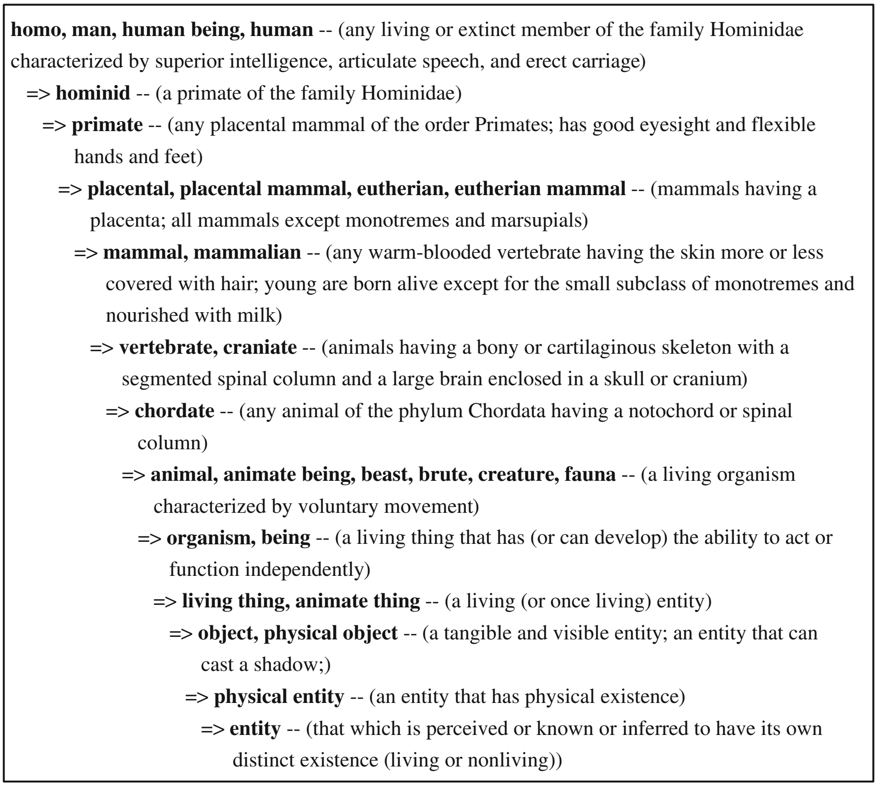

```
homo, man, human being, human -- (any living or extinct member of the family Hominidae
characterized by superior intelligence, articulate speech, and erect carriage)
 => hominid -- (a primate of the family Hominidae)
   => primate -- (any placental mammal of the order Primates; has good eyesight and flexible
       hands and feet)
     => placental, placental mammal, eutherian, eutherian mammal -- (mammals having a
         placenta; all mammals except monotremes and marsupials)
       => mammal, mammalian -- (any warm-blooded vertebrate having the skin more or less
           covered with hair; young are born alive except for the small subclass of monotremes and
           nourished with milk)
         => vertebrate, craniate -- (animals having a bony or cartilaginous skeleton with a
             segmented spinal column and a large brain enclosed in a skull or cranium)
           => chordate -- (any animal of the phylum Chordata having a notochord or spinal
               column)
             => animal, animate being, beast, brute, creature, fauna -- (a living organism
                 characterized by voluntary movement)
               => organism, being -- (a living thing that has (or can develop) the ability to act or
                   function independently)
                 => living thing, animate thing -- (a living (or once living) entity)
                   => object, physical object -- (a tangible and visible entity; an entity that can
                       cast a shadow;)
                     => physical entity -- (an entity that has physical existence)
                       => entity -- (that which is perceived or known or inferred to have its own
                           distinct existence (living or nonliving))
```

Fig. 4.3 Example of a hierarchy of synsets from WordNet (Princeton 2010)

4.3 Knowledge Use

The main use of knowledge by humans (Merrill 2000) as well as by software agents (Newell 1982) is, clearly, to support intelligent behavior essential for solving problems. The process of integrating knowledge into computer systems that are designed to imitate problem solving that normally requires human experts is referred to as *knowledge engineering* (Feigenbaum and McCorduck 1983), while the respective systems are known as *knowledge-based systems* (KBS) or *expert systems*.

A typical knowledge-based system consist at least of:

- a *knowledge base* containing organized expert knowledge represented by using one of knowledge representation formalisms; and
- a *reasoning engine* with mechanisms for automated reasoning. The task of the reasoning engine is to derive new conclusions from the knowledge in the knowledge base in order to imitate the problem solving process of a human expert.

4.3.1 Knowledge Acquisition

Successful problem solving requires a high level of expertise. The value of a knowledge-based system is, therefore, closely related to the quality and the extent of knowledge stored in its knowledge base. The process of extracting, structuring and organizing knowledge to be used in knowledge-based systems is referred to as *knowledge acquisition* (Waterman 1985). The aim of knowledge acquisition is to acquire and structure data required for creating a knowledge model intended to be used for automated problem-solving that is expected to provide similar results compared to those provided by domain experts.

It is precisely the latter who are the most valuable source of knowledge. Knowledge acquisition through direct interaction with human experts is referred to as *knowledge elicitation*. Techniques and methods for knowledge elicitation include interviewing, brainstorming, protocols, laddering, observations, sorting, and many others (Medesker et al. 1995; Schreiber et al. 2000). The purpose of these techniques is to acquire different types of knowledge an expert possesses, including the knowledge an expert is not consciously aware of and it is therefore difficult to reach and articulate (Studer et al. 1998). An example of such knowledge is procedural knowledge which is manifested in the respective expert's skills.

Besides obtaining knowledge from human experts, it can also be extracted from digital sources, such as, for example, electronic documents, databases and the internet, by using the techniques of *knowledge discovery*.

4.3.2 Knowledge Modeling

The data gathered in knowledge acquisition are structured and represented, for example, in the form of annotated documents or diagrams. The structured data are first validated and then used to build a knowledge model in a process referred to as *knowledge modeling*. The knowledge model is stored in a knowledge base by using one of the knowledge representation formalisms that enable the knowledge to be interpretable by a reasoning engine.

The main characteristics of knowledge modeling are the following (Studer et al. 1998):

- The process of knowledge modeling is never completed. As the resulting knowledge models only an approximation of the real world, there is always room for further improvement.
- Consequently, knowledge modeling is often carried out in a number of iterations. In each iteration, the current version of the model serves as a starting point for further refinements and modifications or even acquisition of new knowledge.
- Because knowledge modeling is subject to subjective and consequently potentially faulty interpretations of knowledge engineers, the evolving knowledge models should be revised and evaluated with respect to reality in each stage of the process.

Various methodologies for building knowledge models have been proposed. The majority of them focuses on building ontologies, which is why some of them were already listed in the section on ontologies. CommonKADS methodology (Schreiber et al. 2000), an example of a methodology exceeding the scope of ontologies, is presented in the continuation.

The knowledge model described in CommonKADS captures three categories of knowledge required to solve a particular problem: domain knowledge, inference knowledge and task knowledge. In the continuation of this survey, the three categories are described by using an example of a simple medical diagnosis application (Fig. 4.4).

Domain knowledge corresponds to an ontology with domain-specific terms which can be used as a static knowledge base reusable for solving diverse tasks within an application. In the medical diagnosis application, domain knowledge would, for example, include the definitions of symptoms, diseases, and tests to confirm the diseases, and in addition, the relations among the above-mentioned elements.

Inference knowledge describes the reasoning primitives, or *inferences*, that can be used to carry out the reasoning process applied on domain knowledge. The knowledge model of the medical application, presented in Fig. 4.4, includes two such inferences. The inference "hypothesize" relates the observed symptoms with a probable disease, while the inference "verify" determines the tests that are able to confirm whether the specified symptoms are, in fact, caused by the assumed disease.

Task knowledge describes the goals of applying knowledge in the application as well as strategies to accomplish these goals. Such strategies are described with the aid of several levels of decompositions, through which complex tasks are broken-down into more basic tasks which are eventually associated to inferences. Beside the decomposition process, task knowledge also defines the way in which the tasks

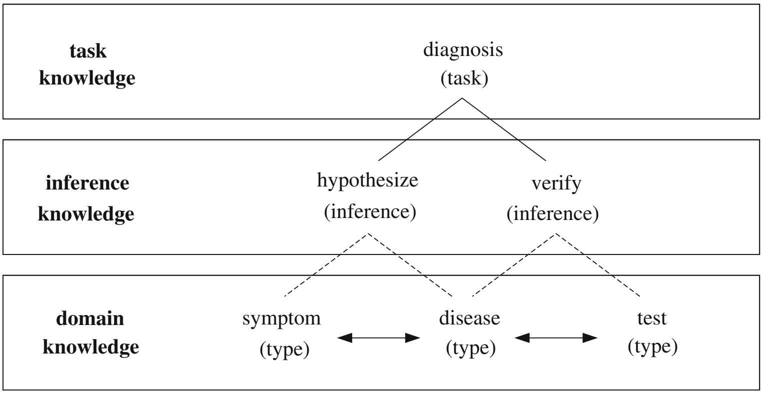

Fig. 4.4 The knowledge model in CommonKADS (Schreiber et al. 2000)

are carried out. The top-level “diagnosis” task of the medical application from Fig. 4.4 can, for example, be carried out by repeatedly invoking the sequence of both inferences from the inference layer.

4.3.3 Reasoning

Once knowledge is acquired, modeled and stored in a knowledge base, it is ready to be used for problem solving. The process of problem solving requires deriving conclusions reached by reasoning over explicitly represented knowledge. Conclusions can be reached by using different methods and strategies, and can thus be supported with different rationales on why a particular conclusion was selected over a wide range of others. In the continuation, the most common types of reasoning are presented, i.e. deduction, induction, abductive reasoning and reasoning by analogy.

Deduction is a type of reasoning which necessarily derives a conclusion from the given premises. If the premises are true, then the conclusions derived from the deduction process are also true. In general, deductive reasoning derives specific conclusions from more general evidence, for example:

Premise 1:	*All men are mortal*
Premise 2:	*Socrates is a man*
Conclusion:	*Socrates is mortal*

Induction can be in many ways considered an opposite of deduction. To start with, it is a type of reasoning which draws *general* conclusions based on the abstraction of observations of (many) individual *specific* instances, for example:

Premise 1:	*Socrates is mortal*
Premise 2:	*Socrates is a man*
Conclusion:	*All men are mortal*

As opposed to deductive reasoning, inductive reasoning does not guarantee the truth of the conclusions, even if all the premises are true. This can be manifested in the philosophical *problem of induction*, famously illustrated by the historic *black swan problem*:

Premise:	*All swans we have seen so far were white*
Conclusion:	*All swans are white.**

This conclusion is, of course, false, since black swans were discovered in the eighteenth century.

Abductive reasoning is a type of reasoning which does not draw certain conclusions but rather yields hypotheses or explanations of observation. As there can be an infinite number of explanations, abductive reasoning attempts to bring forward a single explanation by invalidating alternative explanations.

Abductive reasoning is, for example, very common in medicine. When diagnosing a patient, many possible diseases fit the displayed symptoms, but one of them is considered as more probable than others.

Analogical reasoning compares specific details of two concepts and concludes that if the examined concepts are alike in those details, then they can possibly be alike also in (some of) the others. Analogical reasoning can be considered as a form of inductive reasoning, as it does not assure the truth of the conclusions but rather extends, although perhaps inaccurately, our understanding on previously unknown concepts. This is illustrated in the following example, where (at least if referring to cartoon characters) the analogical reasoning yields a false conclusion.

Premise 1:	*Tom is a cat and Tom catches mice*
Premise 2:	*Garfield is a cat*
Conclusion:	*Garfield catches mice*.*

Analogical reasoning derives particular conclusions from particular premises and in this respect differs from the other three types of reasoning, where at least one of the propositions is general.

Solving complex problems often requires the application of a combination of different types of reasoning. Diagnosing a disease in the medical application from Fig. 4.4, for example, involves an instance of abductive reasoning ("hypothesize") and an instance of deductive reasoning ("verify"), which is illustrated in Fig. 4.5.

4.3.4 Applications of Knowledge-Based Systems

One of the common uses of knowledge-based systems is to provide non-professional users professional guidance when it is difficult to provide the actual support of an expert. The user interacts with a knowledge-based system (KBS) through a user interface which allows the user to issue queries to the KBS, answer additional questions asked by the KBS and receive solutions or advice (Fig. 4.6).

Knowledge-based systems are currently used to assist the users in many fields, for example in:

- strategy games. (A good example demonstrating the power of knowledge when used alongside the processing power of a machine is computer chess. The knowledge base of a chess game contains the strategies and moves that can be used to simulate the opposing "expert" player. Today, only few people can win against the computer in chess, and they can only succeed by taking advantage of the knowledge on how computer plans its next moves.);
- recommender systems, for example, recommending movies (Movielens 2012; (IMDb 2012), songs (Last.fm 2012; Pandora 2012) and shopping items (Amazon 2012) or suggesting social connections (Facebook 2012; LinkedIn 2012);

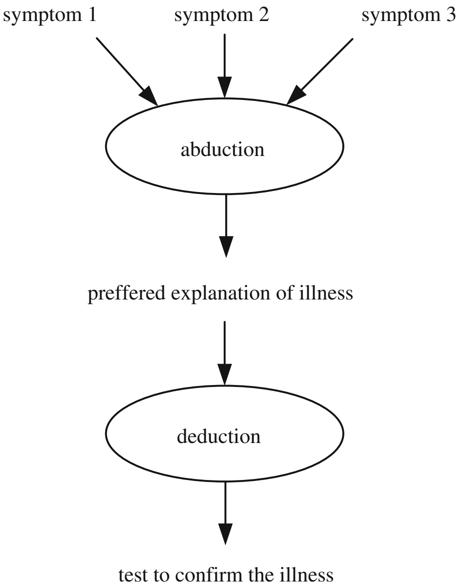

Fig. 4.5 Using a combination of different types of reasoning for problem solving

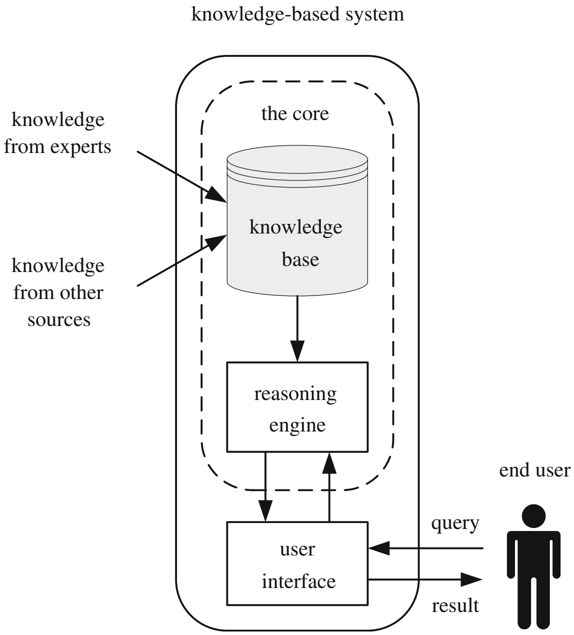

Fig. 4.6 The general structure of a knowledge-based system

- interactive applications (answering engines (Wolfram 2012), virtual interactive assistants (Siri 2012), tutorial applications (Hatzilygeroudis and Prentzas 2004), etc.).

The use of knowledge-based systems is not limited to the facilitation of tasks of non-professional users, but may also be used to support expert work. This is especially useful in the fields where vast knowledge is required to solve a problem and/or many combinations of input variables (observations) are common, which can lead to seemingly unpredictable conclusions. Such fields, for example, include:

- various kinds of diagnoses, for example machine fault diagnosis, e.g. (Jain et al. 2008), and medical diagnosis, e.g. (Miller et al. 1982);
- complex decision support systems, for example in industrial production, e.g. (Manohar et al. 1999), and agricultural production, e.g. (Cohen and Shoshany 2002); and
- finance, for example in financial analysis (Matsatsinis et al. 1997).

A more extensive review of knowledge-based systems methodologies and applications can be found in (Liao 2005).

The final goal of every application of knowledge, regardless of the nature of the agents using it (be it human or software), is solving some sort of a problem. This chapter presented the most important topics related to the use of knowledge by software agents with the goal of autonomous problem solving. The topics presented focused on how to achieve this goal through the definitions of knowledge, various representations and organizations of knowledge acquired from human experts and other resources, and modeling the acquired knowledge in order to be stored and reused for further manipulation. By using different types of reasoning, the manipulation of explicitly presented knowledge produces new knowledge which is required in order to answer a particular question or, in general, to solve a specific problem.

References

Amazon.com (2012) Amazon.com: online shopping for electronics, apparel, computers, books, DVDs & more.www.amazon.com/. Accessed 25 Sep 2012

Apple (2012) iPhone 4S—Ask Siri to help you get things done. http://www.apple.com/iphone/features/siri.html. Accessed 25 Sep 2012

Baader F, Horrocks I, Sattler U (2008) Description logics. In: van Harmelen F et al (eds) foundations of artificial intelligence. Handbook of knowledge representation, vol 3. Elsevier, Amsterdam, pp 135–179

Brachman RJ, Levesque HJ (2004) Knowledge representation and reasoning. Morgan Kaufmann, Los Altos

Cohen Y, Shoshany M (2002) A national knowledge-based crop recognition in Mediterranean environment. Int J Appl Earth Obs Geoinf 4(1):75–87

Davis R, Shrobe H, Szolovits P (1993) What is a knowledge representation? Artif Intell Mag 14(1):17–33

Facebook (2012) http://www.facebook.com/. Accessed 25 Sep 2012

Feigenbaum EA, McCorduck P (1983) The fifth generation: artificial intelligence and Japan's computer challenge to the world. Addison-Wesley Publishing, Reading

Fellbaum C (2010) WordNet. In: Poli R et al (eds) Theory and applications of ontology: computer applications, Springer Science+Business Media B.V., pp 231–243

Guarino N, Giaretta P (1995) Ontologies and knowledge bases, towards a terminological clarification. In: Mars NJI (ed) Towards very large knowledge bases. IOS Press, Amsterdam, pp 25–32

Halladay S, Milligan C (2004) The application of network science principles to knowledge simulation. In: Proceedings of the 37th annual Hawaii international conference on system sciences, Big Island, 5–8 Jan 2004

Hatzilygeroudis I, Prentzas J (2004) Using a hybrid rule-based approach in developing an intelligent tutoring system with knowledge acquisition and update capabilities. Expert Syst Appl 26(4):477–492

IMDb (2012) Movies, TV and celebrities. http://www.imdb.com/. Accessed 25 Sep 2012

Jain MB, Srinivas MB, Jain A (2008) A novel web based expert system architecture for on-line and off-line fault diagnosis and control (FDC) of transformers. In: Proceedings of the TENCON 2008—2008 IEEE region 10 conference, Hyderabad, 19–21 Nov 2008

Last.fm (2012) http://www.last.fm/. Accessed 25 Sep 2012

Liao S-H (2005) Expert system methodologies and applications—a decade review from 1995 to 2004. Expert Syst Appl 28(1):93–103

LinkedIn (2012) World's largest professional network | LinkedIn. http://www.linkedin.com/. Accessed 25 Sep 2012

Manohar PA, Shivathaya SS, Ferry M (1999) Design of an expert system for the optimization of steel compositions and process route. Expert Syst Appl 17(2):129–134

Matsatsinis NF, Doumpos M, Zopounidis C (1997) Knowledge acquisition and representation for expert systems in the field of financial analysis. Expert Syst Appl 12(2):247–262

Medesker L, Tan M, Turban E (1995) Knowledge acquisition from multiple experts: problems and issues. Expert Syst Appl 9(1):35–40

Merrill MD (2000) Knowledge objects and mental models. In: Proceedings of the international workshop on advanced learning technologies, Palmerston, 12 Apr–12 June, pp 244–246

Miller RA, Pople HE Jr, Myers JD (1982) Internist-1, an experimental computer-based diagnostic consultant for general internal medicine. N Engl J Med 307(8):468–476

Minsky M (1975) A Framework for representing knowledge. In: Winston PH (ed) The psychology of computer vision. McGraw-Hill, New York, pp 211–277

Movielens (2012) Movielens—movie recommendations. http://movielens.umn.edu/. Accessed 25 Sep 2012

Newell A (1982) The knowledge level. Artif Intell 18(1):87–127

Pandora Internet Radio (2012) http://www.pandora.com/. Accessed 25 Sep 2012

Peirce CS (1909) Existential graphs, manuscript 514. http://www.jfsowa.com/peirce/ms514.htm. Accessed 25 Sep 2012

Princeton University (2010) About WordNet. http://wordnet.princeton.edu. Accessed 25 Sep 2012

Russell B (1912) The problems of philosophy. Home University Library, London

Ryle G (1949) The concept of mind. Hutchinson Press, London

Schreiber G, Akkermans H, Anjewierden A, de Hoog R, Shadbolt N, Van de Velde W, Wielinga B (2000) Knowledge engineering and management: the CommonKADS methodology. MIT Press, Cambridge

Sowa JF (1984) Conceptual structures: information processing in mind and machine. Addison-Wesley Publishing, Reading

Sowa JF (2010) Ontology. http://www.jfsowa.com/ontology/index.htm. Last update 29 Nov 2010. Accessed 25 Sep 2012

Stillings NA, Weisler SW, Chase CH, Feinstein MH, Garfield JL, Rissland EL (1995) Cognitive science: an introduction, 2nd edn. MIT Press, Cambridge

Studer R, Benjamins VR, Fensel D (1998) Knowledge engineering: principles and methods. Data Knowl Eng 25(1–2):161–197

van Geenen E (2004) Knowledge structures and the usability of knowledge systems. Eburon Uitgeverij, Delft

Waterman DA (1985) A guide to expert systems. Addison-Wesley Publishing, Reading

Wolfram Research (2012) Wolfram|Alpha: computational knowledge engine. http://www.wolframalpha.com/. Accessed 25 Sep 2012