Chapter 1

Introduction

The problem of semantic representation of natural language information is the central topic of this work. This task is important for the following scientific disciplines:

- the theoretical foundations of artificial intelligence (AI), concerning the knowledge representation problem itself;
- linguistics, in connection with the formal description of the semantics of natural language expressions and for the formalization of lexical knowledge;
- cognitive psychology, to model conceptual structures and the processes of reasoning;
- the development of natural language processing systems, e.g. question-answering systems or machine translation systems (especially for the creation of knowledge bases and large computational lexica).

The present work describes a comprehensive repertoire of representational means, allowing for an adequate description of the semantics of natural language expressions, be it "on paper" or in a computer. At the same time, the use of these representational means for the investigation of fundamental problems of natural language semantics will be demonstrated.

In dealing with the semantics of natural languages in general, one should be aware that the term **language** (Ge: "Sprache") has a twofold meaning. On the one hand, we have to investigate language as a system, e.g. the German or English languages, with its regularities being independent of actual speech acts (this system aspect has been called "la langue" by the Swiss linguist de Saussure). In this regard, the meaning of natural language expressions can be described independently of a specific context of utterance, and, therefore, we speak of the **primary meaning** or **core meaning**. The investigation of meaning in this sense is the topic of a special branch of linguistics and computational linguistics, known as the "Semantics", which is also the main topic of this book.

On the other hand, language expressions can be investigated with regard to their use in concrete utterances (this aspect is called "la parole" by de Saus-

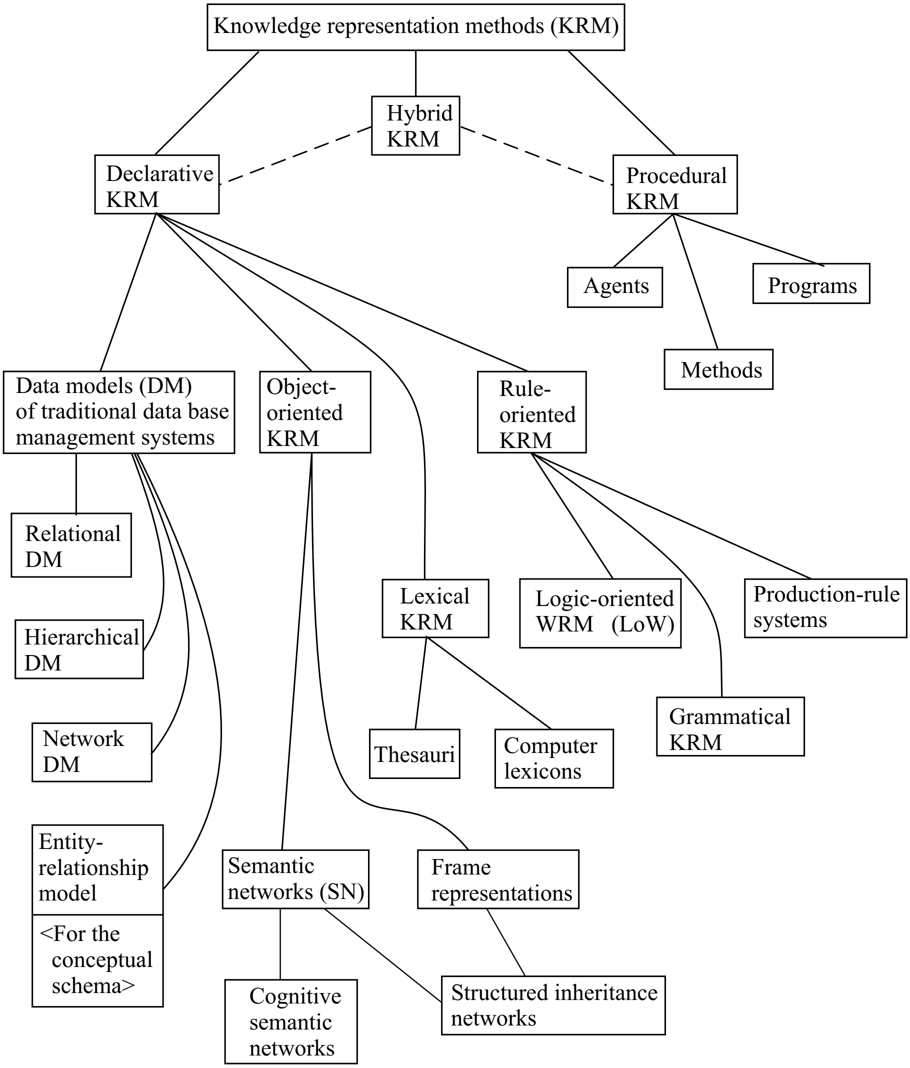

Figure 1.1. Overview of knowledge representation models

sure). It is connected with a specific context of utterance or a specific dialogue situation. Researching this aspect of meaning, one has to take into consideration the intentions of the speaker/writer (what the speaker/writer really means) as well as the effects on the hearer/reader (what is achieved with the hearer/reader). In general, such an utterance has various **side meanings** or **secondary meanings** apart from its primary meaning (its propositional content).

The investigation of these aspects is the subject of **pragmatics** or **speech act theories** (see [229]), which will be only touched upon in this work.

It is important to note that the representational means described in this book are language-neutral and thus provide a kind of semantic **interlingua**. The general paradigm of knowledge representation they are embedded in is the **semantic network** paradigm (see Chap. 2). Its position in the world of knowledge representation methods is shown in Fig. 1.1. **Multi**layered Extended Semantic **Net**works (acronym: **MultiNet**s) are based on the following main components (a detailed description is given in Part II; for a short description of MultiNet see [113]):

- Representatives of concepts (the nodes of the network);
- Functions and relations (providing the arcs between the nodes);
- Sorts and features representing semantic classes ("normal" labels of nodes);
- Multidimensionally organized attributes of nodes (labels of nodes which are the basis for the discrimination of different layers in the network);
- Methods of encapsulation (used for the partitioning of the network);
- Axiomatic rules (used for the inferential connection of nodes and for the formal definition of relations and functions).

Unfortunately, there is no generally accepted definition to determine the **adequacy** of a system of representational means. From a theoretical point of view, it would be very desirable to have a finite and manageable number of basic conditions or criteria from which we could "automatically" derive an appropriate set of semantically primitive representational means (e.g. the sorts, relations, functions, and semantic features proposed in this book). It would be very helpful if these criteria allowed us to decide whether certain representational elements are admissible or not, whether they are necessary or not, etc. In reality, such a complete set of criteria does not exist (at least not at this moment). In addition, no system of semantically primitive representational means can fully cope with the richness of nuances and the diversity of natural language, because it necessarily results from classification, generalization, and therefore also from coarsening. But, just as no linguistic theory can do without a classifying and coarsening concept formation, no natural language processing system that has to be realized in practice and that has to use a large knowledge base can do without a classifying and systematizing repertoire of representational means. The epistemically and cognitively fundamental relationships mirrored in the set of representational means are important for another reason as well. They are carriers of the most important inferential mechanisms connected with conceptual reasoning. To renounce them would make it necessary to interconnect all natural language constructs standing in logical relation to each other, leading to an unmanageable combinatorial explosion.

It has to be emphasized that the term **semantically primitive representational means** does not imply that every concept can be decomposed and reduced to the meaning of semantically primitive elements (as attempted in [276]). Rather, these representational means are used as irreducible concepts on a metalevel to classify the concepts on the semantic level and to describe their fundamental interrelationships and the inferential connections between them. Since every natural language is both **language** and **metalanguage** at the same time, all formally defined concepts can be described in natural language again, something that produces a complicated hierarchy of language layers. To avoid an infinite iteration, this hierarchy is closed by formal constructs, which are described by logical methods.

In the following list, we propose a set of criteria which should be fulfilled by every system of representational means in order to provide the basic elements of a formal description of natural language semantics:

a) Principal (global) requirements

- **Universality** – The representational means must be defined independently of a specific natural language or the application domain, and should not be “tailored” ad hoc to a special field of discourse.
- **Cognitive adequacy** – The representational means must allow for an adequate modeling of human conceptual structures (as far as they are known) and of their manifestation in the semantics of the natural language expressions describing them. These models must be **concept-centered**, i.e. every concept should have a unique representative through which all information belonging to it is accessible.[1]
- **Interoperability** – The representational means must be applicable to theoretical investigations of semantically oriented disciplines or computational linguistics, as well as to the specification of formal interfaces to the components of applied AI systems. They should be usable for the construction of a computational lexicon as well as for expressing the results of the syntactic-semantic analysis, as building blocks for the formal language used in the inference machine, and as a basis for text generation. Only in this way can the results of the above-mentioned disciplines or system components build on each other and the necessary integration be achieved.
- **Homogeneity** – It must be possible to describe the meanings of words or the meanings of sentences or texts (dialogues) with the same means used for the description of logical rules governing the formal processes of reasoning.

---

[1] In computer science, this characterization would be called **object-orientedness**.

- **Communicability** – No single person is able to construct a large knowledge base or a complex applied AI system like a question-answering system or a theoretical edifice that covers all the semantic phenomena of a natural language. To accomplish this, whole teams are required whose members cooperate effectively. This necessitates, among other things, that they have a common understanding of the representational means of a knowledge representation system to be used by all of them. Hence the requirement that the definitions also be intuitively intelligible.
- **Practicability** – Every knowledge representation system designed for a real-world application has to fulfill certain pragmatic requirements, i.e. it must be technically tractable and effectively implementable. Of what use, for instance, is the most fine-grained semantic representation if no one is able to provide the corresponding background knowledge necessary for the syntactic-semantic analysis to disambiguate the theoretically possible variants of meaning, or if the representation and processing of knowledge can not be effectively implemented or dealt with in such a highly differentiated system?
  This requirement has also a quantitative aspect. The usefulness of a KRS should be proved by applying its expressional means to the description of thousands, or tens of thousands, of concepts. It is of little use to demonstrate the functioning of a KRS with a few examples if the representational principles proposed cannot be practically maintained during the treatment of a large stock of knowledge.
- **Automatability** – The predefined repertoire of expressional means should permit automatic processing of knowledge, and especially automatic knowledge acquisition from natural language sources.

b) Internal, structural requirements

- **Completeness** – There should be no meaning which cannot be represented with the representational means. It must be emphasized that this requirement does not concern completeness in the logical sense, i.e. that every true expression which can be formulated in the representational language should also be derivable.
- **Optimal granularity** – On the one hand, different meanings must be mapped into different structures; on the other hand, to keep a system manageable, not every fine semantic nuance can be mirrored in a KRS.
- **Consistency** – Pieces of information logically contradicting each other must not be derivable from one another. For equivalent meanings, however, it is precisely this mutual derivability that must be warranted. It follows that inference rules and the definitions of representational means (carried in Multi-

Net essentially by the R-Axioms) must be adapted to each other in such a way that the kernel of a knowledge base must be **globally consistent**. When knowledge about concrete concepts or concrete facts is added, the knowledge base must only be **locally consistent** (see Sect. 13.1). This means that knowledge pieces contradicting themselves in one part of the knowledge base must not affect other parts which are not semantically connected with it.

- **Multidimensionality** – The qualitative distinction of different aspects of knowledge (immanent vs. situational knowledge, intensional vs. extensional aspect, quality vs. quantity, etc.) must be mirrored in the assignment of concepts to different layers of representation.
- **Local interpretability** – The basic constructs should be logically interpretable by themselves, and independent of their embedding in the context of the knowledge base as a whole.

One question often arises in connection with knowledge representation: Does a **canonical meaning representation** exist, i.e. are we able to define a general function which maps semantically equivalent NL-expressions into identical meaning representations? As already stated by Woods [281], there are theoretical reasons why such a canonical representation does not exist at all, since such representations do not exist even for formal languages essentially weaker than natural languages (cf. the undecidability of the word problem or of the problem of simplification for symbolic mathematical expressions [213]). What can be achieved, however, is a certain **normalization** of the meaning structures of natural language expressions. Thus the great variety of semantic structures can be reduced by identifying the representations of semantically (nearly) equivalent sentences (e.g. active vs. passive voice), or by ignoring the differences in the topic-focus structure of sentences (see [231], [90]).

In the present work, we prefer a semiformal, content-oriented definition of the representational means. This is a necessary precondition for a completely formal treatment of their semantics; for, how could one define formally what is not completely understood conceptually? When describing the relations and functions of MultiNet, we use logical expressions which give a starting point for the inferential interlinking of meaning structures and thus provide a basis for the definition of an operational semantics (see Sect. 13.2). However, it has to be conceded that an entirely formal description of the semantics of the representational means has still to be worked out on the basis of the more content-oriented definitions in this book. Basically, three different methods can be taken into account for this purpose:

a) a **model-theoretic extensional method**, as used in logic and logic-oriented semantic theories. This approach is already problematic because many natural language concepts, and also the proposed expressional means of MultiNet, can be interpreted extensionally only with great difficulty, if at all. What are the extensions of "*religion*", "*illness*", "*abstract*", "*physical*", "*intension*", etc. or how do we treat modal restrictions of temporal relationships like "*possibly after the dinner*" extensionally?

b) a **procedural method**, as it is used in natural language interfaces to databases or in robotics, where meaning representations of natural language queries are mapped onto procedural expressions of the target system (e.g. onto retrieval procedures of a database management system or onto actions of a robot); this method also has only restricted applicability;

c) a **use-theoretic method**, where the meaning of concepts and semantic primitives is defined by their interrelation among themselves and by their proper use in the language game or in a question-answering game ("meaning as use").

We believe that for the foundation of meaning representation, as well as for theoretical investigations, the latter method, which dates back to Wittgenstein [279], is the most appropriate one. A purely procedural explanation of concept meanings is at best apt for restricted applications (e.g. for the above-mentioned natural language interfaces or for interpreting natural language commands to robots).[2] As a basic assumption discerning a) from b) and supporting c), we cite the following thesis:
"*Concepts essentially do not work as classificators during language understanding, thus discerning between "meant" and "non-meant" (this approach is typical of an extension-based model-theoretic semantics); they rather are connectives receiving their full potential in their mutual interconnections and enabling us to experience reality and to communicate our experiences to others.*" [242]

Furthermore, the truth or falsity of sentences does not play such a central role for understanding natural language as assigned to these categories in logic-oriented (extensional) theories of semantics. Human beings are often not able to decide on the truth or falsity of a proposition or on the applicability of a concept to a real object, even if the utterance in question has been understood (see the discussion in Sect. 15.3).

On the basis of this argument, the present work prefers method c), which, according to Wittgenstein, can be described as a question-answering game (or

---

[2] This method is actually used by our research group for realizing natural language access to the Internet and natural language interfaces to traditional databases (see [115]).

language game) governed by its own rules and manifested in the correct interplay between question and answer. This method is most clearly realized in the paradigm of a question-answering system of artificial intelligence (see Part II and [111]). For better understanding of our concern, this paradigm can be thought of as an integrated system into which the knowledge representation methods of MultiNet are embedded. The question-answering system does not have to be an implemented AI system; it can also be imagined as an abstract functional model into which the essential processes of language understanding are integrated. Because of this double interpretability, we deliberately use the same abbreviation **QAS** throughout the book for both terms, **question-answering game** and **question-answering system**.[3] Those familiar with the methods of artificial intelligence and automatic knowledge processing may associate a question-answering system with a QAS; those approaching the problem of meaning representation from linguistics or psychology may interpret the abbreviation QAS as “language game”.

It should be stated that human beings apparently have all three methods at their disposal to support the symbolic conceptual system that is closely connected with natural language. They are able to link words or concepts with objects of the world (analogously to model-theoretic/extensional semantics of formal theories where predicates are mapped into sets of individuals in an artificial “world”, i.e. into a universe of a predefined algebraic structure); human beings are also able to translate language expressions (e.g. the command “*Stand up!*”) into actions, i.e. into contractions of their muscles (“procedural” semantics); finally, they are able to interconnect concepts in a dialogue in a correct way without resorting to the first two methods (use-theoretic semantics).

Although the book mainly deals with knowledge representation and not with knowledge processing, it might be useful for the understanding of the whole system to explain the embedding of the knowledge base into a QAS. The functional diagram given in Fig. 1.2 can be seen as representative of both a technical question-answering system and a “natural” question-answering process. It comprises all the components characteristic of a question-answering game where the knowledge base (which can be built using the representational means of MultiNet) plays a central role.

Information formulated in natural language and given to a computer or a person must first be analyzed to determine its meaning, which has to be ex-

---

[3] In German we coined the ambiguous abbreviation **FAS** for **QAS** after the corresponding terms **Frage-Antwort-Spiel** (En: **question-answering game**) and **Frage-Antwort-System** (En: **question-answering system**) which have the same initial letters in their components and thus better mirror this double interpretability.

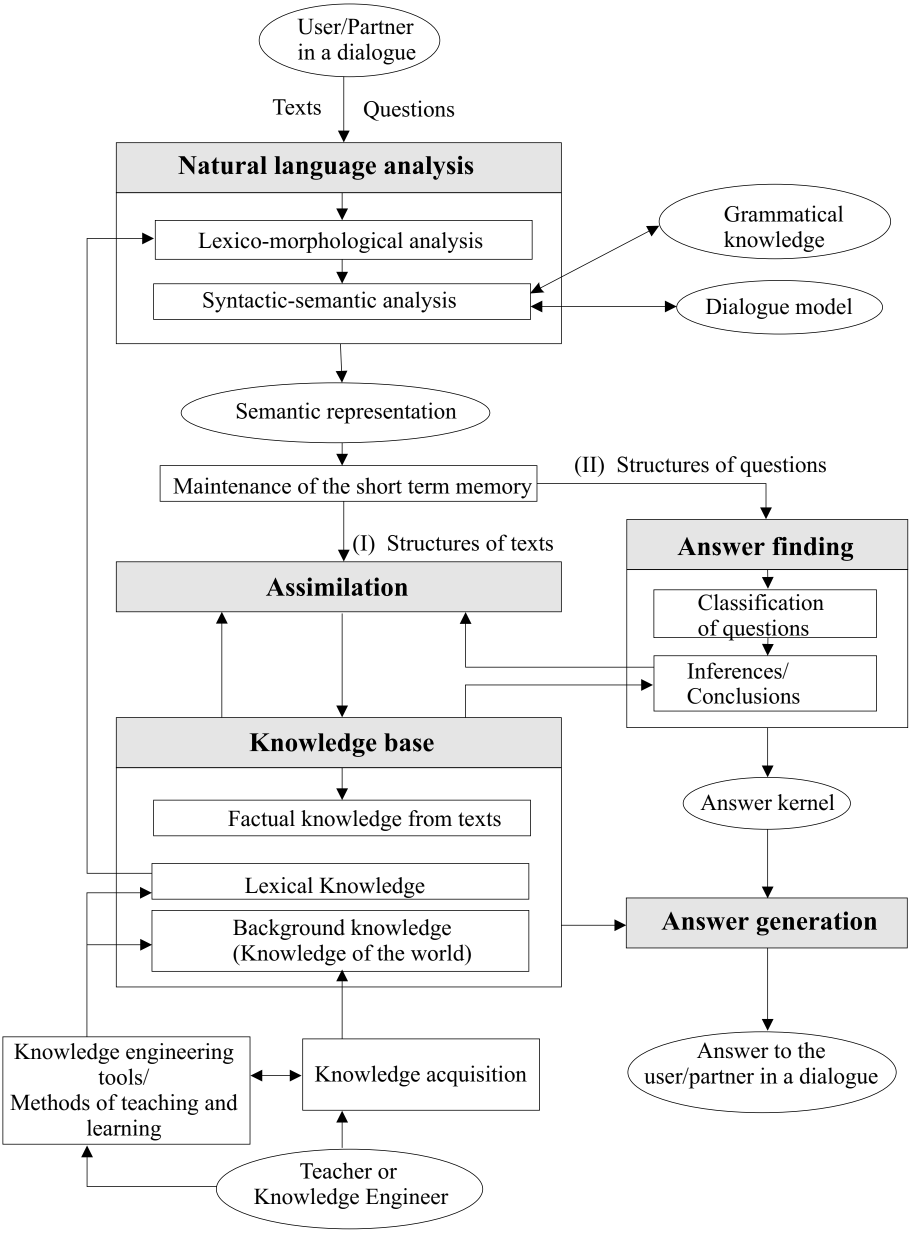

Figure 1.2. Overview of the most important functional components of a QAS

pressed in a convenient format for semantic representation. In this process, the **lexico-morphological analysis** is mainly based on lexical knowledge, while the **syntactic-semantic analysis** is mainly supported by grammatical knowledge and world knowledge. Furthermore, the interpretation of natural language expressions generally requires a **dialogue model** describing the situational embedding of the utterances. This is especially important for the understanding of deictic language elements (which, among other things, comprise deictic pronouns, like "*I*", "*you*", or deictic adverbs, like "*here*", "*there*", "*yesterday*").[4]

To ensure the interaction of all components, the same representation formalism should be used for the **lexical information** (see Chap. 12) and the **background knowledge** needed for the language understanding process, as well as for the **dialogue model** (MultiNet has been used successfully in all three fields). MultiNet can also be used to a certain extent for the formalization of **grammatical knowledge**, which plays a role in the **word-class-controlled functional analysis** [116], especially in the semantic interpretation of prepositions and conjunctions [258].

The semantic structures of single sentences are stored at first in a short-term memory so that intersentential references (especially pronoun references between sentences) can be resolved. Afterward, questions and propositions (in general texts) are processed differently. While questions are subjected to logical answer finding, the information contained in texts (propositions) has to be assimilated into the knowledge base. The **assimilation process** connects incoming meaning structures with knowledge already available in the knowledge base or possibly identifies them with equivalent pieces of information to avoid double storage. In addition, the assimilation has to close apparent "semantic gaps" in texts by using background knowledge available in the knowledge base.[5]

Finding an appropriate answer to a given question is based on a process of **question classification** (see Sect. 3.2.4). The type of query does not only deter-

---

[4] The term **deixis** denotes the phenomenon that certain language expressions are related to elements of the situational context of an utterance ("*here*" denotes the location of the speaker/writer; "*today*" denotes the day comprising the moment of speaking or writing the expression, etc.).

[5] Let us take the following sentences: "*The firm NN developed a new car. The motor needs only 3 litres of gas per 100 km.*" These sentences lack a semantic connection, if there is no knowledge available that a car has a motor as its part (this kind of information is called **world knowledge** instead of **linguistic knowledge**). Especially the reference induced by the definite article in the phrase "*the motor*" cannot be resolved without this knowledge. In the present case, the assimilation should be able to find the correct subordination of concepts and to supply a corresponding part-whole relation PARS between the concepts ⟨new car⟩ and motor and add it properly to the conceptual structures already stored in the knowledge base.

mine the **inference method** to be applied, but also the type of knowledge asked for (situational vs. immanent, definitional vs. assertional, etc.; see Sect. 3.2.3). Additionally, the classification of questions is relevant to the answer finding and is used as a basis for **answer generation**. Answers to decision questions are of another type (namely "*Yes*" or "*No*") than answers to supplementary questions (also called "WH-questions"). In the latter case, it is typical that a single node of the semantic network, the so-called **answer kernel**, found during the process of answer finding, has to be reformulated in natural language. With so-called "essay questions" ("*What is a Y?*", "*What do you know about X?*", etc.), a whole text (essay) has to be generated, stemming mainly from the immanent knowledge of the answer kernel. While the aspect of deduction, or of logical inferences in general, plays a prominent part in answer finding for decision questions and supplementary questions, the aspect of information retrieval is dominant for essay questions. In the latter case, it is the retrieval of the immanent knowledge connected with the answer kernel and its reformulation which is predominant (see Sects. 3.2.3 and 13.2).

Concluding this chapter we want to state that MultiNet and its predecessors have proved their usefulness in many applications; among them we mention the following:

- Knowledge representation in question-answering systems [107, 98]
- Semantic interlingua in natural language interfaces to databases and to the Internet [109, 159, 114]
- Semantic annotation of large text corpora and automatic knowledge acquisition [98, 79]
- Backbone for building large semantically based computational lexica [100, 192]
- Central knowledge representation formalism in the virtual electronic laboratory VILAB [166, 164].