

Computational Logics and the Philosophy of Language: The Problem of Lexical Meaning in Formal Semantics

Marcello Frixione
University of Genova and IIASS “E.R. Caianiello”
Vietri sul Mare (SA). Italy
e-mail: frix@dist.unige.it

Abstract

This paper deals with the possible contributions that logical researches carried on in the field of artificial intelligence (AI) could give to formal theories of meaning developed by logically oriented philosophers of language within the tradition of analytic philosophy. In particular, I will take into account a topic which is problematic in many respects for traditional logical accounts of meaning, i.e., the problem of lexical semantics. My thesis is that AI logics could give useful instruments to face some aspects of the problem, whilst other aspects fall outside the scope of a logical treatment, and require other kinds of computational instruments.

1 Lexical semantics and formal theories of meaning

The dominant formal paradigm in the philosophical study of meaning is so-called model theoretic semantics. Based on formal tools of a set theoretic nature, model theoretic semantics originated from the researches on the meaning of expressions of formal languages. The simplest and historically prior form of model theoretic approach is extensional Tarskian semantics for first order predicate logic. According to Tarskian semantics, an *interpretation* I for a first order theory is a ordered pair $I = (\varphi, D)$. D is the *domain* of the interpretation, that is to say, a set of objects on which the constructs of the language are interpreted. φ is an *interpretation function*, associating a reference in D to the various constructs of the language (e.g., to every individual constant is associated an element of D , to one place predicative letters are associated subsets of D , and so on). Recursive compositional rules allow

to establish the reference of syntactically complex expressions, given the reference of their components. In particular, the truth value of a closed formula can be established starting from the reference of the atomic symbols occurring in it. An interpretation I that makes true a certain formula α is said to be a *model* of α (in symbols, $I \models \alpha$). A model of a set Γ of formulas is a model of all formulas in Γ . Model theoretic semantics allows to formally characterise such notions as the validity of a formula, or the relation of logical consequence between formulas. For example, a formula α is *valid* ($\models \alpha$) if and only if it is true in all the interpretations. A formula β is *logical consequence* of a formula α ($\alpha \models \beta$) if and only if β is true in all the models of α . Historically, after the development of Tarskian extensional theory, model theoretic semantics has been extended to deal with intensional aspects of meaning with the development of possible world semantics, due mainly to the researches of Rudolph Carnap and Saul Kripke. Originally developed to account for the meaning of formal languages, possible world semantics showed to be a very flexible and powerful tool, which allowed one to face many aspects of the meaning of natural languages. The application of formal intensional semantics to natural languages reached its full development with the work of Richard Montague [?]. Montague's intensional framework is, at the present time, the most outstanding approach to formal semantics of natural languages, and is adopted not only by philosophers of language, but also by many linguists (see e.g. [?]). This approach is based on the idea that the meaning of a sentence can be identified with its truth conditions: knowing what a given sentence means amounts to know how must the world be to make it true.

In spite of its successes, the model theoretic approach leaves some unsolved problems. One example is the study of lexical semantics. Model theoretic techniques are in different respects insufficient to satisfactorily treat word meaning. In the model theoretic tradition, a partial treatment of lexical semantics is obtained via the introduction of *meaning postulates* [?]. Meaning postulates are formulas that are assumed to hold in all the interpretations of a language, with the effect of constraining the set of the possible interpretations of extra-logical primitive symbols of the language itself. Simple examples of formulas that can be used as meaning postulates are the following:

$$(1) \quad \forall x(\text{poodle}(x) \rightarrow \text{dog}(x))$$

$$(2) \quad \forall x(\text{dog}(x) \rightarrow \neg \text{cat}(x)).$$

If we assume that these formulas are true in all interpretations, the number of admissible models is reduced. For example, (1) imposes that the interpretation of *poodle* is a subset of the interpretation of *dog*; (2) imposes that the interpretation of *dog* is disjointed from the interpretation of *cat*.

However, traditional meaning postulates are unsatisfactory in at least two different respects. First of all, traditional meaning postulates can express only necessary and/or sufficient conditions for the application of a concept. The problem is that, for concepts corresponding to natural language words, necessary and/or sufficient

conditions are rather the exception than the norm. Usually, it is possible to identify only sets of features, that characterise concepts in “typical” cases, but that admit an undefined number of exceptions. For example, everybody that knows the meaning of the word “lemon”, almost surely must know that usually lemons are yellow. But this information cannot be captured by means of a traditional meaning postulate. Indeed, a formula like $\forall x(\textit{lemon}(x) \rightarrow \textit{yellow}(x))$ is immediately falsified by the existence of lemons that are not yellow (in that, for example, they are not ripe, or rotten). Similarly, the fact that birds usually fly has intuitively something to do with the meaning of the word “bird” (somebody that fails to know that birds usually fly, surely does not know what the word “bird” means). But the possible examples of birds that cannot fly are numberless. Analogous problems arise with verbal expressions. The meaning of “running” seems to involve that what runs usually moves, but in atypical cases this may not happen (e.g., consider people running on a tapis roulant). In general, for almost all common sense concepts (including lexical concepts) it is not possible to individuate sets of necessary and/or sufficient conditions, that are rich enough to characterise them. This kind of problems has been independently put in evidence by researchers from different fields. Consider, for example, in the field of cognitive psychology, the notion of prototype proposed by Eleanor Rosch [?], or, in AI, Marvin Minsky’s concept of frame [?]. In philosophy, partly similar conclusions have been achieved, in very different philosophical contexts, by Ludwig Wittgenstein [?] with the notion of family resemblance, and by Hilary Putnam [?] with the notion of stereotype.

There is a second aspect, largely independent from the above considerations, with respect to which model theoretic semantics is inadequate to face the problem of lexical meaning. It concerns the relations between words and extra linguistic reality. The point is that model theoretic semantics cannot capture the “intended interpretation” of lexical items. The aim of model theoretic semantics is that of determining the meaning of syntactically complex expressions, starting from the meaning of their components. However, as far as the atomic extra-logical elements of a language are concerned (that is, in a linguistic perspective, the elements of the lexicon), their meaning is simply assumed as given. At the extensional level, this is evident in the case of Tarskian semantics for first order logic. Extensional Tarskian semantics allows to establish the reference (i.e., the truth value) of the closed formulas of a first order language, assuming as known the behaviour of the interpretation function for primitive extra logical symbols (i.e., individual constants, predicative and functional letters). This aspect of Tarski’s theory has been clearly put in evidence by Field [?], who stressed how Tarskian semantics does not give a reduction of semantic notions (e.g., the notion of truth) to non semantic notions. Tarski simply reduces the semantic notion of truth to other semantic notions, as, for example, the notion of denotation of a singular term or of a predicate, which are assumed as primitive. *Mutatis mutandis*, the situation is unchanged in possible world intentional semantics. The intension of a lexical item is modelled as a function from possible worlds to extensions of the suitable semantic type. However, model theoretic semantics has nothing to say on which extension must be associ-

ated to a lexical item in the different possible worlds. In other words, the theory does not tell anything about the difference between the extensions of two lexical items of the same syntactical type. In this sense, we can say that model theoretic semantic does not help in individuating the truth conditions of the sentences of a language; rather, it allows to calculate the truth conditions of sentences given the intensions of the primitive symbols occurring in them.

A different way to formulate the problem is the following. As seen before, the introduction of meaning postulates has the effect of limiting the number of admissible models. However, this is not sufficient to fix the interpretation of primitive extra logical symbols. For example, once that (1) and (2) are introduced as meaning postulates, nothing tells us that the interpretation of *poodle* is the set of poodles and the interpretation of *dog* is the set of dogs. (1) and (2) are compatible with an interpretation in which, for example, *poodle* is interpreted on the set of maples, *dog* is interpreted on the set of trees, and *cat* is interpreted on the set of cars. In general, no set of meaning postulates can assure that the interpretation of primitive symbols is the “intended” one. Given any set of meaning postulates, there exists always an infinite number of not intended, non isomorphic interpretations satisfying them (see, for example, [?]). Also in this case, analogous considerations hold both for extensional Tarskian semantics and for intensional possible world semantics.

Following Diego Marconi (see e.g. [?]), I distinguish an *inferential* component and a *referential* component of lexical semantics (note that Marconi formulates this distinction with respect to lexical *competence*). The *inferential* component concerns the complex whole of connections linking together the elements of the lexicon, and expressing the non logical relations existing among the words of a language. Examples of relations of this kind are that dogs are mammals; that birds usually fly; that a father is a human male who has sons; that if somebody runs, then he (usually) moves, and so on. Usually, such information is used by speakers to draw inferences. The *referential* component concerns the mapping of words on the one hand, and objects, events and situations in the world on the other. For example, referential competence has to do with the ability of classifying a given animal as a dog, or of distinguishing it from a cat, of describing somebody as running, and so on. It could be described as the ability to determine the values of the interpretation function for primitive, extra logical symbols of a language. So, the referential component of lexical semantics has to do with the identification of the real truth conditions of sentences. In view of this distinction, the above considerations could be stated by saying that both inferential and referential aspects of lexicon are, in different respects, problematic for model theoretic semantics. The former because traditional meaning postulates are insufficient to capture all the inferentially relevant connections between lexical items (as in the case of prototypical features). The latter because model theoretic tools do not allow to capture the “intended” interpretation of primitive extra logical symbols of a language.

2 Non monotonicity, lexical meaning, and formal semantics

The logical formalisms developed in AI have been seldom adopted or proposed to face aspects of lexical semantics that fall outside the possibilities of traditional techniques. For example, fuzzy logic has been proposed to formalise the fuzziness of many lexical concepts, and to try to capture some prototypical aspects of word meaning ([?, ?]). Richmond Thomason indicated non monotonic logics as suitable candidates to face many aspects of lexicon [?]. Also knowledge representation systems as frames and semantic network can be considered in this perspective. They have often been used to represent lexical meaning in AI natural language processing systems, and in this perspective have sometimes raised the interest of some philosophers of language. Here I assume that frames and semantic nets can be considered in good approximation as notational variants of logical formalisms, seldom endowed with non classical features, such as, for example, non monotonicity to treat defaults and exceptions. In our perspective, all these formalisms can be considered ways to work out more flexible systems of meaning postulates, allowing to capture inferential aspects of lexicon, that fall outside the expressive possibilities of traditional meaning postulates.

In the following I shall consider in particular the contribution of non monotonic logics to the problem of lexical meaning representation. I shall adopt as a starting point for my argument the formalism of circumscription. *Circumscription* is an approach to non monotonic reasoning originally proposed by John McCarthy [?, ?, ?]. It is based on the following intuition. Some predicates apply only to exceptional cases, and do not hold in typical situations. In common sense reasoning, people usually implicitly assume that such predicates do not hold, unless it is explicitly known that it is not so. An example could be the “albino” predicate. Albino people are a very small subset of human beings. Usually, nobody, in reasoning about a person he does not know, takes into account the possibility that he is albino, unless some explicit information on the matter is available. The intuitive idea at the basis of circumscription is that, in a logical theory, the predicates like albino must be “circumscribed”. Circumscribing a predicate P in a theory T amounts to assume that P has the smallest extension compatible with the information in T . For example, consider the following theory (written in the language of first order calculus with identity):

$$T = \{albino(Gigi), Piero \neq Gigi, Pippo \neq Gigi, Pippo \neq Piero\}.$$

Circumscribing the predicate *albino* in T amounts to assume that Gigi is the only albino individual in the theory domain. That is to say, from the circumscription of *albino* in T , it must follow that:

$$\forall x(albino(x) \rightarrow x = Gigi).$$

With circumscription we can formalise rules expressing defeasible regularities, such as, for example, that birds usually fly. To represent non monotonic rules of this

kind with circumscription, a predicate ab is introduced, whose intuitive meaning is *abnormal* [?]. For example, that birds usually fly is represented by the following formula:

$$(\dagger)\forall x(bird(x) \wedge \neg ab(x) \rightarrow flies(x)).$$

That is to say, if something is a bird, and it is not an abnormal bird, it flies. To obtain the desired consequences, the ab predicate must be circumscribed. In other words, circumscribing ab amounts to assume that are abnormal all and only the individuals that are explicitly known to be abnormal. I.e., it is assumed that abnormal birds are as few as possible. Consider for example the theory:

$$T = \{\forall x(bird(x) \wedge \neg ab(x) \rightarrow flies(x)), bird(Tweety)\}.$$

Given the information in T , no individual can be deduced to be abnormal. Therefore, circumscribing ab in T amounts to assume that $\forall x\neg ab(x)$. As a consequence, we have that $\forall x(bird(x) \rightarrow flies(x))$ and $flies(Tweety)$. Let us suppose now that T is extended to the following theory T' :

$$T' = T \cup \{bird(Fred), \neg flies(Fred), Fred \neq Tweety\}.$$

From T' it can be derived that $ab(Fred)$. So, circumscribing ab in T' , we will obtain $\forall x(ab(x) \rightarrow x = Fred)$, and, as a consequence, $\forall x(bird(x) \wedge x \neq Fred \rightarrow flies(x))$. The situation is analogous, even if more complex, when, in a theory there are more abnormality predicates to be circumscribed, as in the following example:

$$(i) \quad \forall x(bird(x) \wedge \neg ab_2(x) \rightarrow flies(x))$$

$$(ii) \quad \forall x(ostrich(x) \wedge \neg ab_1(x) \rightarrow \neg flies(x))$$

Adopting circumscription, formulas like (†) or (i-ii) can be used as meaning postulates, to express defeasible regularities in the characterisation of the inferential component of lexical meaning.

The relevance of non monotonic formalisms in the logical treatment of lexical meaning is stressed in the following passage by Richmond Thomason:

Logician work in AI has generally recognized the need for augmenting the Fregean logical framework in order to deal with problems of common sense reasoning. The most generally accepted line of development is the incorporation of non monotonicity into the logic. And this feature, it turns out, is precisely what is needed to accommodate many of the problems that emerged in Montague-style lexical semantics. [...] It is no surprise that lexical semantics is full of defeasible generalizations, and a general technique for expressing such generalizations would greatly extend the coverage of logicist theories of world meaning. ([?], p. 4)

Formally, circumscribing some predicate in a first order theory T amounts to adding to T suitable second order axioms. Different kinds of axioms give rise to different kinds of circumscriptive theories, apt to formalise different types of non monotonic reasoning. I do not focus on these aspects here. From our point of view, it is worth noting that circumscription admits a semantic treatment of a

model theoretic nature. Model theoretic semantics for circumscriptive logics was proposed by McCarthy himself, and it is based on the concept of *minimal model* of a theory with respect to a certain set of predicates. The minimal model semantics captures the intuition that the models of the circumscription of P in a theory T are those models of T in which the predicate P has the smallest extension compatible with available information.

As an example, I describe the minimal model treatment of the simplest form of circumscription, i.e. *predicate circumscription* ([?]). In terms of minimal model semantics, predicate circumscription is characterised as follows (for sake of simplicity, I consider here the case in which only one predicate is circumscribed).

Let T be a first order theory, and P a predicate of the language of T , and let $M = (\varphi, D)$ and $M' = (\varphi', D')$ be two models of T . We shall say that M' is a *P-submodel* of M for the theory T if and only if:

1. $D' = D$;
2. $\varphi'[P] \subseteq \varphi[P]$;
3. $\varphi'[K] = \varphi[K]$ for every constant K of the language of T different from P .

We shall say that M is a *P-minimal* model of T if and only if every model of T which is a *P-submodel* of M is identical to M .

For every theory T , and for every predicative constant P of the language of T , the models of the predicate circumscription of P in T are all and only the *P-minimal* models of T . In other words, a formula α follows from the predicate circumscription of P in T if and only if α is true in all *P-minimal* models of T . *P-minimal* models of T are those models of T in which the predicate P has the smallest extension compatible with T itself.

Predicate circumscription can be used to formalise cases like the example of the albino predicate above. However, it cannot account for non monotonic rules of the kind of (\dagger). In such cases also the predicate *flies* must be allowed to change its extension during the process of minimisation (intuitively, the less are the abnormal birds, the more are the flying ones). This can be obtained by means of another, more powerful, kind of circumscription, namely *formula circumscription* [?]. Even more complex is the case of set of rules like (i-ii), where more then one abnormality predicate is minimised. These cases require other forms of circumscription, such as *prioritized circumscription* [?, ?]. However, the central idea of minimal model approach to circumscription is already evident in the simplest case of predicate circumscription.

The possibility of a model theoretic account of circumscription witnesses for the possibility of its integration in the framework of intensional semantics. The homogeneity of circumscription with Montague paradigm is stressed by Thomason:

Among the available theories of defeasible reasoning that could be applied in lexical semantics, circumscription is the one that perhaps can most easily and naturally blended in the existing work in Montague's framework, because it provides an integrated logical package that is based on higher-order logic. [...] The project of

developing a broadly successful logic-based account of semantic interrelationships among the lexical items of a natural language is roughly comparable in scope with the project of developing a high-level theory of common-sense knowledge. ([?], p. 457)

Besides circumscription, other logical formalisations of non monotonic reasoning have been proposed. Among them, Reiter’s default logic and non monotonic modal logics (e.g., autoepistemic logic). Various relations exist between these different types of non monotonic formalisms, and various equivalence results have been proved. I do not address these issues here. For our purposes, it is interesting to shortly consider the general semantic approach proposed by Yohav Shoham ([?, ?]), in order to offer a unitary framework for non monotonic reasoning, that could encompass and generalise the existing non monotonic systems. Shoham’s proposal is based on the concept of *preferred model*. The basic, intuitive, idea is that, in non monotonic reasoning, only a subset of the models of a theory is taken into consideration. These models enjoy some special characteristic, they are “preferred” from some point of view. This is evident in the case of circumscription, where models are preferred, in which the circumscribed predicates have a smaller extension. Shoham’s theory originates exactly as a generalisation of McCarthy’s minimal model semantics for circumscription. Different preference criteria give rise to different non monotonic logics. Shoham called his approach *preferential semantics*.

Shoham formulation of preferential semantics is very comprehensive. Let \mathbf{L} be a standard logic. Shoham calls “standard logic” a classical monotonic logic, for example propositional logic or first order predicate logic, or a modal, propositional or predicative, logic. A non monotonic logic can be obtained from \mathbf{L} introducing a preferential ordering on the *interpretations* of \mathbf{L} . According to Shoham, an *interpretation* is everything can stay on the left of \models , where \models is the usual (monotonic) relation of semantic consequence of \mathbf{L} . Hence, in this context, an interpretation can be an interpretation in the strict sense (e.g., a Tarskian interpretation for first order logic), or, in the case of a modal logic, a pair (I, w) , where I is a Kripke structure, and w is a possible world. This allows, for example, to encompass in the framework of preferential semantics also autoepistemic logics, whose semantics is given in terms of possible worlds. Let $<$ be a strict partial ordering on the interpretations of \mathbf{L} , and M_1 and M_2 interpretations of \mathbf{L} . $M_1 < M_2$ means that M_2 is *preferred* to M_1 . \mathbf{L} and $<$ define a new *preferential logic* $L_{<}$. Let α and β be formulas of the language of \mathbf{L} . The concepts of *satisfaction* and of *logical consequence* for $L_{<}$ are defined as follows.

Definition: *an interpretation M preferentially satisfies α (in symbols, $M \models_{<} \alpha$) if $M \models \alpha$, and if no other interpretation M' exists, such that $M < M'$ and $M' \models \alpha$. In this case, M is called a preferred model of α .*

Definition: *β is a preferential consequence of α (in symbols, $\alpha \models_{<} \beta$) if, for every interpretation M , if $M \models_{<} \alpha$, then $M \models \beta$ (in other words, β is a preferential*

consequence of α iff β is true in all the preferred models of α).

The ordering relations between models of the minimal model approach can be seen as particular cases of preference relations in preferential semantics. Besides circumscription, preferential semantics has been applied to other non monotonic systems, including autoepistemic logic and default logic.

In conclusion, the present section can be summarised as follows. Non monotonic logics could be useful to face many inferential aspects of lexical semantics. Moreover, they are homogeneous with formal tools traditionally adopted in the logico-philosophical tradition, and, in this respect, they could be integrated within an intensional semantic framework. However, as we shall see in the next section, these techniques leave referential aspects of lexicon completely unaffected.

3 Are AI logics enough?

The fact itself that non monotonic logics admit a model theoretic semantic treatment suggests that the above considerations concerning the impossibility of determining intended interpretations (see § 1) apply also in this case. The situation with non monotonic meaning postulates is even worse than in the traditional case. In fact, non monotonic meaning postulates are “looser” than classical ones in constraining the models of a theory. This point can be efficaciously illustrated with a joke by Ettore Petrolini:

Two friends are telling each other riddles.

“Can you tell me what’s green, sits on the top of the piano, and goes ‘chirp-chirp-chirp’?”

The other thinks and thinks, and finally responds:

“I don’t know, you tell me.”

“A herring”

“But herrings aren’t green!”

“So just paint it.”

“There’s no way it’s on the piano!”

“So just put it there.”

“And on the top of all that, herrings don’t go chirp.”

“Right! But if I didn’t say that, it would have been too easy to guess!” ([?], p. 172)

That is to say, typical herrings are greyish, dumb, and stay in the depths of the sea. However, if these are prototypical features, that can be violated by atypical specimens, then also a green object on the top of a piano could belong to the interpretation of the predicate *herring*. As an extreme case, if any feature represented in a set of meaning postulates is defeasible, then any interpretation would be a model of that theory. We saw that two meaning postulates as $\forall x(poodle(x) \rightarrow dog(x))$ and $\forall x(dog(x) \rightarrow \neg cat(x))$ would admit as their models interpretations in which the predicate *poodle* is interpreted on the set of maples, *dog* is interpreted on the set of trees, and *cat* is interpreted on the set of cars. In the case of circumscriptive theories, it can happen that a theory including the following meaning postulates:

$$\forall x(fledged(x) \wedge \neg ab_1(x) \rightarrow flies(x))$$

$$\forall x(\text{bird}(x) \wedge \neg \text{ab}_2(x) \rightarrow \text{fledged}(x))$$

cannot exclude from its models interpretations in which, for example, *fledged* is interpreted on the set of dogs, *flies* is interpreted on the set of baboons, and *bird* is interpreted on the set of bicycles. According to preferential semantic terminology, the problem is that non monotonic meaning postulates do not constrain *all* the models of a theory, but only their *preferred* models. For example, the second of the above circumscriptive meaning postulates imposes that those models are preferred, in which the members of the interpretation of *bird* that are also members of the interpretation of *fledged* are as many as possible. But this does not exclude that, in certain models, it could happen that no birds are fledged.

Approaching the referential aspects of lexicon amounts, in a computational perspective, to face the problem of *symbol grounding* in the sense introduced by Steven Harnad [?]. Primitive extra-logical symbols at the linguistic level must be grounded on perception, on motor activities, and, in general, on non linguistic capabilities of an autonomous computational agent. In a model theoretic perspective, symbol grounding can be considered as akin, in a certain sense, to the problem of devising how the interpretation function (and its inverse) can be calculated for (some of) the primitive symbols of a language. I maintain that, in order to calculate the values of such functions, the primitive elements on which computations are defined must be chosen at a “lower” level of analysis with respect to lexical concepts. They cannot be semantically homogeneous to natural language words, and must ultimately be connected to measurements and data coming from sensors. It is improbable that logic could offer the most suitable paradigm to face this kind of problems, and computational devices such as “subsymbolic” neural networks could prove to be more adequate. Much work has been done in the perspective of integrating symbolic systems and neural networks (see for example Massimo De Gregorio’s contribution to this same volume [?]). However, a general formal framework putting together in a principled way logic and “subsymbolic” forms of computation is still needed (for a promising proposal in this direction, see Peter Gärdenfors’ theory of conceptual spaces, presented in his contribution to this volume [?]).

Acknowledgements

I wish to thank Guglielmo Tamburrini for reading and discussing an earlier version of this paper.

References

- [Carnap52] Carnap, R., 1952, Meaning postulates. *Philosophical Studies*, 3, 65-73.
- [Chierchia90] Chierchia, G. e S. McConnell-Ginet, 1990, *Meaning and Grammar. An Introduction to Semantics*. The MIT Press, Cambridge, MA.

- [*DeGregorio96*] De Gregorio, M., 1996, Integrating inference and neural classification in a hybrid system for recognition tasks. This volume.
- [*Field72*] Field, H., 1972, Tarski's theory of truth. *The Journal of Philosophy*, 69 (13), 347-375.
- [*Gärdenfors96*] Gärdenfors, P., 1996, Concept representations and nonmonotonic inferences. This volume.
- [*Harnad90*] Harnad, S., 1990, The symbol grounding problem. *Physica D*, 42(1/3).
- [*Lakoff72*] Lakoff, G., 1972, Hedges: a study in meaning criteria and the logic of fuzzy concepts. In *Papers from the Eight Regional Meeting of the Chicago Linguistic Society*, Chicago Linguistic Society, Chicago, 183-228.
- [*Lifschitz85*] Lifschitz, V., 1985, Computing circumscription. *Proc. 9th IJCAI*, Los Angeles, CA, 121-127.
- [*Lifschitz86*] Lifschitz, V., 1986, On the satisfiability of circumscription. *Artificial Intelligence*, 28, 12-27.
- [*Marconi95*] Marconi, D., 1995, On the structure of lexical competence, *Proc. Aristotelian Society*, 95, 131-150.
- [*McCarthy77*] McCarthy, J., 1977, Epistemological problems of artificial intelligence. *Proc. 5th IJCAI*, Cambridge, MA, 1038-1044.
- [*McCarthy80*] McCarthy, J., 1980, Circumscription - a form of nonmonotonic reasoning. *Artificial Intelligence*, 13, 27-39.
- [*McCarthy84*] McCarthy, J., 1984, Applications of circumscription to formalizing common sense knowledge. *Proc. AAAI Workshop on Non-Monotonic Reasoning*, New Paltz, NY, 295-324. Also in *Artificial Intelligence*, 1986, 28, 89-116.
- [*Minsky75*] Minsky, M., 1975, A framework for representing knowledge. In P. Winston (ed.), *The Psychology of Computer Vision*, McGraw-Hill, New York.
- [*Montague74*] Montague, R., 1974, *Formal Philosophy*. Yale University Press, New Haven and London.
- [*Oden77*] Oden, G.C., 1977, Fuzziness in semantic memory: Choosing examples of subjective categories. *Memory and Cognition*, 5, 198-204.
- [*Petrolini71*] Petrolini, E., 1971, *Teatro*. Garzanti, Milano.
- [*Putnam70*] Putnam, H., 1970, Is semantics possible? In H. Putnam, *Mind, Language and Reality*. Cambridge University Press, Cambridge, 1975.
- [*Putnam81*] Putnam, H., 1981, *Reason, Truth and History*. Cambridge University Press, Cambridge.

- [Rosch75] Rosch, E., 1975, Cognitive representation of semantic categories. *Journal of Experimental Psychology: General*, 104, 192-233.
- [Shoham87] Shoham, Y., 1987, Non-monotonic logics, meaning and utility. *Proc. 10th IJCAI*, Milano, 388-93.
- [Shoham88] Shoham, Y., 1988, *Reasoning about Change*. MIT Press, Boston, MA.
- [Thomason91a] Thomason, R., 1991, Knowledge representation and knowlegde of words. *Proc. of Workshop on Lexical Semantics and Knowledge Representation*, University of California, Berkley, 1-8
- [Thomason91b] Thomason, R., 1991, Logicism, AI, and common sense: John McCarthy's program in philosophical perspective. In V. Lifschitz (ed.), *Artificial Intelligence and Mathematical Theory of Computation*, Academic Press, Boston.
- [Wittgenstein53] Wittgenstein, L., 1953, *Philosophische Untersuchungen*. Blackwell, Oxford.