Meinongian Semantics
and
Artificial Intelligence

William J. Rapaport

Department of Computer Science and Engineering,
Department of Philosophy, and Center for Cognitive Science
State University of New York at Buffalo, Buffalo, NY 14260-2000
`rapaport@cse.buffalo.edu`
`http://www.cse.buffalo.edu/∼rapaport/`

originally written in 1985;
updated 3 February 1995;
updated again, 28 June 2005

Abstract

This paper introduces semantic networks to a philosophical audience and surveys several approaches to semantic-network semantics that, at the original time of writing (1985) had not previously been treated in the AI or computational-linguistics literature (but see §12), though there is a large philosophical literature investigating them in some detail. In particular, propositional semantic networks (exemplified by SNePS) are discussed, it is argued that only a fully intensional, Meinongian semantics is appropriate for them, and several Meinongian systems are presented.

This paper was once allegedly forthcoming in: Simons, Peter (ed.), *Essays on Meinong* (Dettelbach: Roell Verlag) (or possibly Amsterdam: Rodopi Editions). A shorter version appeared as: Rapaport, William J. (1985), "Meinongian Semantics for Propositional Semantic Networks", *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics (University of Chicago)* (Morristown, NJ: Association for Computational Linguistics): 43–48, [http://acl.ldc.upenn.edu/P/P85/P85-1006.pdf].

1 Meinong, Philosophy, and Artificial Intelligence

Philosophy has not been kind to Alexius Meinong, a late-19th/early-20th-century cognitive scientist (see Findlay 1963; Grossmann 1974; Rapaport 1978, 1991b; Routley 1979; Lambert 1983; Schubert-Kalsi 1987). Only recently (ca. 1970s) has there been a renaissance in Meinong studies. Even so, his writings are often treated as curiosities (or worse) by mainstream philosophers. One way of characterizing Meinong's contribution to philosophy is in terms of his emphasis on what might now be called his thoroughgoing intensionalism. While some philosophers were ridiculing and rejecting this approach, some computer scientists working in the field of artificial intelligence (AI) have—for largely independent, though closely related, reasons—been arguing for it. In this essay, I explore some of their arguments and show the relevance of Meinongian theories to current research in AI.

2 Semantic Networks

The area of AI known as "knowledge representation" is concerned with systems for representing, storing, retrieving, and inferring information in cognitively adequate and computationally efficient ways. (A better terminology is "*belief*" *representation*; cf. Rapaport & Shapiro 1984; Rapaport 1986b, 1992; Rapaport et al. 1997.)

A *semantic network* is a representational system consisting of a graph (usually a labeled, directed graph), whose nodes represent objects and whose arcs represent relations among them. There have been several different kinds in the literature (see Sowa 1992, 2002 for a general introduction to them, and Findler 1979, Brachman & Levesque 1985, Sowa 1991, and Lehmann 1992 for technical surveys of some of them). As William A. Woods (1975: 44) puts it,

> The major characteristic of the semantic networks that distinguishes them from other candidates [for knowledge representation systems] is the characteristic notion of a link or pointer [an arc, in the above terminology] which connects individual facts into a total structure.

M. Ross Quillian's early "semantic memory" (1967, 1968, 1969) introduced semantic networks as a model of *associative memory*: Nodes represented words and meanings; arcs represented "associative links" among these. The "full concept" of a word was the entire network of nodes and arcs reachable from the node representing that word. *Inheritance* (or *hierarchical*) *networks* use such arc labels as "inst[ance]", "isa", and "prop[erty]" to represent taxonomic structures (see Fig. 1; cf. Bobrow & Winograd 1977, Charniak & McDermott 1985: 22–27). Roger Schank's Conceptual Dependency representational scheme uses nodes to represent conceptual primitives and arcs to represent dependencies and semantic case relations among them (cf. Schank & Rieger 1974; Brand 1984, Ch. 8; Rich & Knight 1991: 277–288; Hardt 1992). The idea is an old one: Networks like those of Quillian, Daniel G. Bobrow and Terry Winograd's KRL (1977), or Ronald J. Brachman's KL-ONE (Brachman 1979, Brachman & Schmolze 1985, Woods & Schmolze 1992) bear strong family

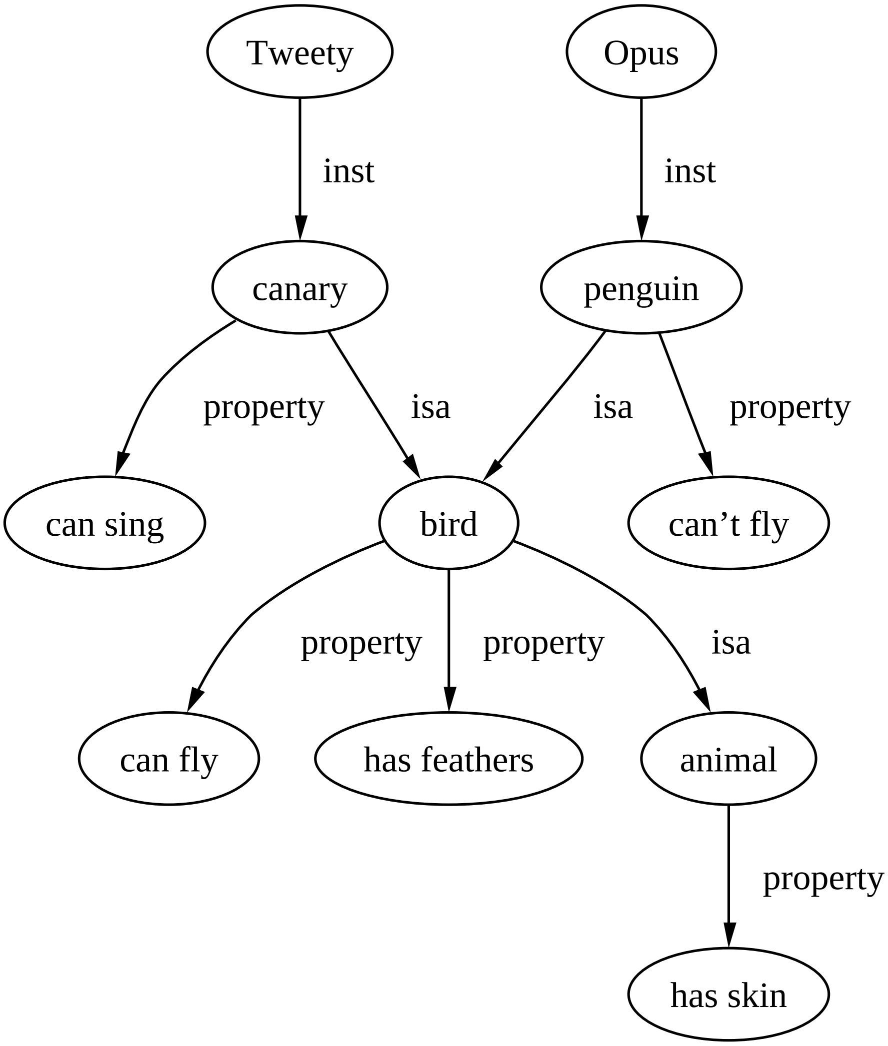

Figure 1: An inheritance network

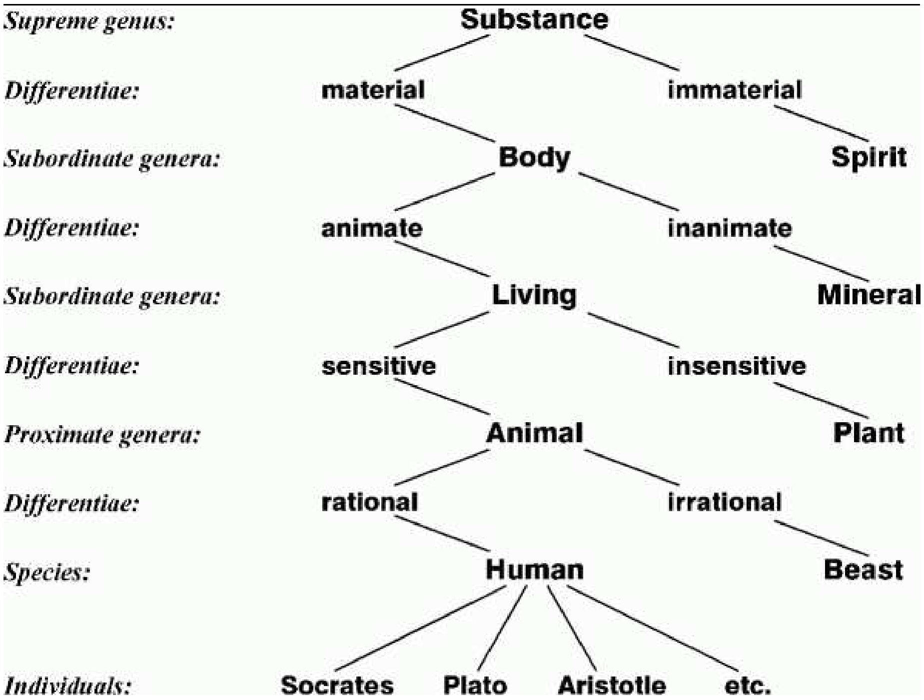

Figure 2: Porphyry's Tree:A mediaeval inheritance network. (From Sowa 2002.)

resemblances to "Porphyry's Tree" (Fig. 2)—a mediaeval device used to illustrate the Aristotelian theory of definition by species and differentia.

3 Semantics of Semantic Networks

It has been pointed out that there is nothing essentially "semantic" about semantic networks (Hendrix 1979; but cf. Woods 1975, Brachman 1979). Indeed, viewed as a data structure, it is arguable that a semantic network is a *language* (possibly with an associated logic or inference mechanism) for representing information about some domain, and, as such, is a purely *syntactic* entity. They have come to be called "semantic" primarily because of their *uses* as ways of representing the *meanings* of linguistic items. (On the other hand, I have argued that this sort of syntax can be viewed as a kind of semantics; cf. Rapaport 1988, 2000.)

As a notational device, a semantic network can itself be given a semantics. That is, the arcs, nodes, and rules of a semantic-network representational system can be given interpretations in terms of the entities they are used to represent. Without such a semantics, a semantic network is an arbitrary notational device liable to misinterpretation (cf. Woods 1975; Brachman 1977, 1983; and, especially, McDermott

1981). For instance, in an inheritance network like that of Figure 1, how is the inheritance of properties to be blocked: If flying is a property inherited by the canary Tweety in virtue of its being a bird, what is to prevent it from being inherited by the penguin Opus? What do nodes represent: classes of objects, types of objects, individual objects? Can arcs be treated as objects, perhaps with arcs linking them in some fashion?

The task of providing a semantics for semantic networks is more akin to the task of providing a semantics for a *language* than for a *logic*, since, in the latter case, but not in the former, notions like *argument validity* must be established and connections must be made with axioms and rules of inference, culminating ideally in soundness and completeness theorems. But underlying the logic's semantics there must be a semantics for the logic's underlying language, and this would be given in terms of such a notion as *meaning*. Here, typically, an interpretation function is established between syntactical items from the language L and ontological items from the "world" W that the language is to describe. This, in turn, is usually accomplished by describing the world in another language, L_W, and showing that L and L_W are notational variants by showing (ideally) that they are isomorphic.

Recently, linguists and philosophers have argued for the importance of *intensional* semantics for natural languages (cf. Montague 1974, Parsons 1980, Rapaport 1981). At the same time, computational linguists and other AI researchers have begun to recognize the importance of representing intensional entities (cf. Woods 1975, Brachman 1979, McCarthy 1979, Maida & Shapiro 1982, and, for a survey, Hirst 1989, 1991). It seems reasonable that a semantics for such a representational system should itself be an intensional semantics. In this paper, I discuss the arguments of Woods and others and outline several fully intensional semantics for intensional semantic networks by discussing the relations between a semantic-network "language" L and several candidates for L_W. For L, I focus on Stuart C. Shapiro's fully intensional, propositional Semantic Network Processing System (SNePS, [http://www.cse.buffalo.edu/sneps/]; Shapiro 1979, 2000a; Shapiro & Rapaport 1987, 1992, 1995), for which David Israel (1983) has offered a possible-worlds semantics. But possible-worlds semantics, while countenancing intensional entities, are not *fully* intensional, since they treat intensional entities extensionally. The L_Ws I discuss all have fully intensional components.

4 Arguments for Intensions

The first major proponent of the need to represent intensional objects in semantic networks was Woods (1975). Brachman (1977) showed a way to do this. And Anthony S. Maida and Stuart C. Shapiro (1982) have argued that *only* intensional entities should be represented.

Woods characterizes *linguistic* semantics as the study of the relations between (a) such linguistic items as sentences and (b) meanings expressed in an unambiguous notation—an internal representation (Woods 1975: 40)—and he characterizes *philosophical* semantics as the study of the relations between such a notation and truth conditions or meanings (Woods 1975: 38f). Thus, he takes semantic networks as examples of the "range" of linguistic semantics and the "domain" of philosophical semantics. Semantic networks, then, are models of the realm of objects of thought (or, perhaps, of the "contents" of psychological acts)—i.e., of Meinong's *Aussersein*.

Woods proposes three "requirements of a good semantic representation": *logical adequacy*—the semantic representation must "precisely, formally, and unambiguously represent any particular interpretation that a human listener may place on a sentence"; *translatability*—"there must be an algorithm or procedure for translating the original sentence into this representation"; and *intelligent processing*—"there must be algorithms which can make use of this representation for the subsequent inferences and deductions that the human or machine must perform on them" (Woods 1975: 45).

The logical-adequacy criterion constitutes one reason why semantic networks

> must include mechanisms for representing propositions without commitment to asserting their truth or belief . . . [and why] they must be able to represent various types of intensional objects without commitment to their existence in the external world, their external distinctness, or their completeness in covering all of the objects which are presumed to exist. (Woods 1975: 36f.)

Note that since some sentences *can* be interpreted as referring to nonexistents, a semantic network ought to be able to represent this, hence must be able to represent intensional entities. (See the next section for a discussion of the other criteria.)

A second argument is that

> semantic networks should not . . . provide a "canonical form" in which all paraphrases of a given proposition are reduced to a single standard (or canonical) form . . .. (Woods 1975: 45.)

If so, then they should not represent *extensional* entities, which would be such canonical forms. There are three reasons why canonical forms are to be avoided. First, there aren't any (the interested reader is referred to the argument in Woods 1975: 46). Second, no computational efficiency would be gained by having them (cf. Woods 1975: 47). Third, it should not be done if one is interested in adequately representing human processing (cf. Rapaport 1981). Sometimes redundant information must be stored: Even though an uncle is extensionally equivalent to a father's-brother-or-mother's-brother, it can be useful to be able to represent uncles directly; thus, it is not an extension, but, rather, an intension, that must be represented (cf. Woods 1975: 48).

Another argument for the need to represent intensional objects comes from consideration of question-answering programs (cf. Woods 1975: 60ff). Suppose that a "knowledge base" has been told that

> The dog that bit the man had rabies.

How would the question

> Was the man bitten by a dog that had rabies?

be represented? Should a *new* node be created for "the dog that bit the man"? The solution is to create such a new node and then decide if it is co-referential with an already existing one.

Finally, intensional nodes are clearly needed for the representation of verbs of propositional attitude (cf.: Woods 1975: 67; Rapaport & Shapiro 1984; Rapaport 1986b, 1992; Wiebe & Rapaport 1986; Rapaport et al. 1997; and the references therein), and they can be used in quantificational contexts to represent "variable entities" (Woods 1975: 68f; cf. Fine 1983; Shapiro 1986, 2000b; Ali & Shapiro 1993; [http://www.cse.buffalo.edu/sneps/Projects/sneps3.html]).

Maida and Shapiro's argument is that although semantic networks *can* represent real-world (extensional) entities or linguistic items, they *should*, for certain purposes, *only* represent intensional ones. The purposes they single out are: representing referentially opaque contexts (e.g., belief, knowledge), representing the concept of a truth value (as in 'John wondered whether *P*'), and representing questions.

In general, intensional entities are needed if one is representing a mind. Why would one *need* extensional entities if one is representing a mind? To represent co-referentiality? No; as we shall see, this can be done (and perhaps can only be done) using only intensional items. To talk about extensional entities? But why would one want to? Everything that a mind thinks or talks about is an object of thought, hence intensional. In order to link the mind to the actual world—i.e., to avoid solipsistic representationalism? But consider the case of perception: There are internal representations of external objects, yet these "need not extensionally represent" those objects (Maida & Shapiro 1982: 300). The "link" would be forged by connections to other intensional nodes or by consistent input-output behavior that improves over time (cf. Rapaport 1985/1986: 84–85, Rapaport 1988, Srihari & Rapaport 1989). (For a survey of the wide variety of items that can be represented by intensional entities, cf. Shapiro & Rapaport 1991.)

5 SNePS

A SNePS semantic network is primarily a *propositional* network (see below). It can, however, also be used to represent the inheritability of properties, either by explicit rules or by *path-based inference* (Shapiro 1978, Srihari 1981). It consists of labeled nodes and labeled, directed arcs satisfying (*inter alia*) the following uniqueness condition (cf. Maida & Shapiro 1982):

(U) There is a 1-1 correspondence between nodes and represented concepts.

A *concept* is "anything about which information can be stored and/or transmitted" (Shapiro 1979: 179). When a semantic network such as SNePS is used to model "the belief structure of a thinking, reasoning, language using being" (Maida & Shapiro 1982: 296; cf. Shapiro 1971b: 513), the concepts are the objects of mental (i.e., intentional) acts such as thinking, believing, wishing, etc. Such objects are intensional (cf. Rapaport 1978). (For further elucidation of this concept of "concept", cf. Shapiro & Rapaport 1991.)

It follows from (U) that the arcs do not represent concepts. Rather, they represent binary, structural relations between concepts. If it is desired to talk *about* certain relations between concepts, then those relations must be represented by nodes, since they have then become objects of thought, i.e., concepts. In terms of Quine's dictum

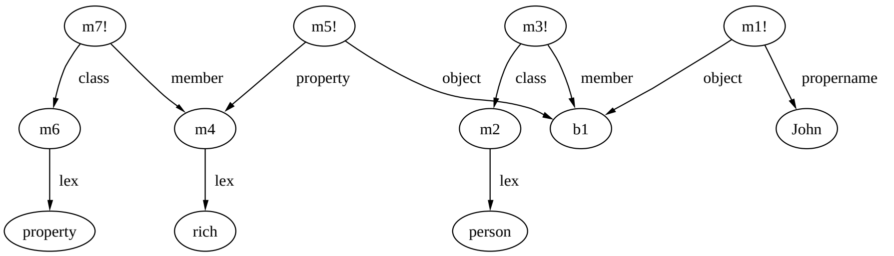

Figure 3: A SNePS representation for "A person named 'John' has the property of being rich".

that "to be is to be the value of a [bound] variable" (Quine 1980: 15; cf. Shapiro 1971a: 79–80), nodes represent such values, arcs do not. That is, given a domain of discourse—including items, *n*-ary relations among them, and propositions—SNePS nodes would be used to represent all members of the domain. The arcs are used to structure the items, relations, and propositions of the domain into (other) propositions. As an analogy, SNePS arcs are to SNePS nodes as the symbols '$\rightarrow$' and '+' are to the symbols 'S', 'NP', and 'VP' in the rewrite rule: $S \rightarrow NP + VP$. It is because no propositions are represented by arcs that SNePS is a "propositional" semantic network (cf. Maida & Shapiro 1982: 292).

Figure 3 shows a sample SNePS network. Node m1 represents the proposition that b1 has the name represented by the node labeled 'John', which is expressed in English by the lexical item 'John'. Node m3 represents the proposition that b1 is a member of the class represented by m2, which is expressed in English by 'person'. Node m5 represents the proposition that b1 (i.e., the person John) is rich (and m4 represents the property expressed by the adjective 'rich'). Finally, node m7 represents the proposition that being rich is a member of the class of things called 'property'. (Nodes whose labels are followed by an exclamation mark, e.g., m1!, are "asserted" nodes, i.e., nodes that are believed by the system; see Shapiro 2000a for details.)

When a semantic network such as SNePS is used to model a mind, the nodes represent only intensional items (Maida & Shapiro 1982; cf. Rapaport 1978). Similarly, if such a network were to be used as a notation for a fully intensional natural-language semantics (such as the semantics presented in Rapaport 1981; cf. Rapaport 1988), the nodes would represent only intensional items. Thus, a semantics for such a network ought itself to be fully intensional.

There are two pairs of types of nodes in SNePS: constant and variable nodes, and atomic (or individual) and molecular (typically, propositional) nodes. (For a discussion of the semantics of variable nodes, see Shapiro 1986.) Except for a few pre-defined arcs for use by an inference package, all arc labels are chosen by the user; such labels are completely arbitrary (albeit often mnemonic) and depend on the domain being represented. The "meanings" of the labels are provided (by the user) only by means of explicit rule nodes, which allow the retrieval or construction (by inferencing) of propositional nodes.

SNePS satisfies Woods's three criteria. It should be clear from the above discussion that it is logically adequate. Shapiro has developed a generalized augmented-transition-network grammar (Shapiro 1982) for automatically translating sentences into SNePS networks and for automatically expressing SNePS networks in sentences of a natural language, thus making SNePS translatable. And the SNePS inference package (supplemented with the SNeBR Belief Revision system) together with user-supplied rules, render it capable of intelligent processing. (See Shapiro 1979, 1989, 1991, 2000a; Rapaport & Shapiro 1984; Rapaport 1986, 1991a; Shapiro & Rapaport 1987, 1992, 1995; Martins & Shapiro 1988; Martins & Cravo 1991; Johnson & Shapiro 2005ab for examples.)

6 Israel's Possible-Worlds Semantics for SNePS

David Israel's semantics for SNePS assumes "the general framework of Kripke-Montague style model-theoretic accounts" (Israel 1983: 3), presumably because he takes it as "quite clear that [Maida and Shapiro] ... view their formalism as a Montague-type type-theoretic, intensional system" (Israel 1983: 2). He introduces "a domain D of possible entities, a non-empty set I (... of possible worlds), and ... a distinguished element w of I to represent the real world" (Israel 1983: 3). An *individual concept* is a function $\mathrm{ic} : I \to D$. Each constant individual SNePS node is modeled by an ic; variable individual nodes are handled by "assignments relative to such a model". However, predicates—which, the reader should recall, are also represented in SNePS by constant individual nodes—are modeled as functions "from I into the power set of the set of individual concepts." Propositional nodes are modeled by "functions from I into $\{T,F\}$," although Israel feels that a "hyperintensional" logic would be needed in order to handle propositional attitudes.

Israel has difficulty interpreting `member`, `class`, and `isa` arcs in this framework. This is to be expected for two reasons. First, it is arguably a mistake to *interpret* them (rather than giving rules for them), since they are arcs, hence arbitrary and non-conceptual. Second, a possible-worlds semantics is *not* the best approach (nor is it "clear" that this is what Maida and Shapiro had in mind—indeed, they explicitly reject it; cf. Maida & Shapiro 1982: 297). Woods argues that a possible-worlds semantics is not psychologically valid, that the semantic representation must be finite (Woods 1975: 50). Israel himself hints at the inappropriateness of this approach:

> if one is focussing on propositional attitude[s] ... it can seem like a waste of time to introduce model-theoretic accounts of intensionality at all. Thus the air of desperation about the foregoing attempt .... (Israel 1983: 5.)

Moreover—and significantly—a *possible*-worlds approach is misguided if one wants to be able to represent *impossible* objects, as one *should* want to if one is doing natural-language semantics (Rapaport 1978, 1981, 1991a; Routley 1979). A fully intensional semantic network demands a fully intensional semantics. The main rival to Montague-style, possible-worlds semantics (as well as to its close kin, situation semantics (Barwise & Perry 1983)) is *Meinongian semantics.*

7 Meinong's Theory of Objects

Meinong's (1904) theory of the objects of psychological acts is a more appropriate foundation for a semantics of propositional semantic networks as well as for a natural-language semantics. In brief, Meinong's theory consists of the following theses (cf. Rapaport 1976, 1978, 1991b):

(M1) Thesis of Intentionality:
Every mental act (e.g., thinking, believing, judging, etc.) is "directed" towards an "object".

There are two kinds of Meinongian objects: (1) *objecta*, the individual-like objects of such a mental act as thinking-of, and (2) *objectives*, the proposition-like objects of such mental acts as believing(-that) or knowing(-that). For example, the object of my act of thinking of a unicorn is: *a unicorn*; the object of my act of believing that the Earth is flat is: *the Earth is flat.*

(M2) Not every object of thought exists (technically, "has being").

(M3) It is not self-contradictory to deny, nor tautologous to affirm, existence of an object of thought.

(M4) Thesis of Aussersein:
All objects of thought are *ausserseiend* ("beyond being and non-being").

For present purposes, *Aussersein* is most easily explicated as a domain of quantification for non-existentially-loaded quantifiers, required by (M2) and (M3).

(M5) Every object of thought has properties (technically, "*Sosein*").

(M6) Principle of Independence:
(M2) and (M5) are not inconsistent. (For more discussion, cf. Rapaport 1986a.)

Corollary:
Even objects of thought that do not exist have properties.

(M7) Principle of Freedom of Assumption: (a) Every set of properties (*Sosein*) corresponds to an object of thought.

(b) Every object of thought can be thought of (relative to certain "performance" limitations).

(M8) Some objects of thought are incomplete (i.e., undetermined with respect to some properties).

(M9) The meaning of every sentence and noun phrase is an object of thought.

It should be obvious that there is a close relationship between Meinong's theory and a fully intensional knowledge-representation and reasoning system like SNePS. SNePS itself is much like *Aussersein*; Shapiro (personal communication) has said that all nodes are implicitly in the network all the time. In particular, a SNePS base (i.e., atomic-constant) node represents an objectum, and a SNePS propositional node represents an objective. Thus, when SNePS is used as a model of a mind, propositional nodes represent the objectives of beliefs (cf. Maida & Shapiro 1982, Rapaport & Shapiro 1984, Rapaport 1986b, Shapiro & Rapaport 1991, Rapaport et al. 1997); and when SNePS is used in a natural-language processing system (cf. Shapiro 1982; Rapaport & Shapiro 1984; Rapaport 1986, 1988, 1991a), individual nodes represent the meanings of noun phrases and verb phrases, and propositional nodes represent the meanings of sentences.

Meinong's theory was attacked by Bertrand Russell on grounds of inconsistency: (1) According to Meinong, the round square is both round and square (indeed, this is a tautology); yet, according to Russell, if it is round, then it is *not* square. (2) Similarly, the existing golden mountain must have all three of its defining properties: being a mountain, being golden, and existing; but, as Russell noted, it *doesn't* exist. (Cf. Rapaport 1976, 1978 for references.)

There have been several formalizations of Meinongian theories in recent philosophical literature, each of which overcomes these problems. In subsequent sections, I briefly describe three of these and show their relationships to SNePS. (Others, not described here, include Routley 1979—cf. Rapaport 1984—and Zalta 1983.)

8 Rapaport's Theory

On my own reconstruction of Meinong's theory (Rapaport 1976, 1978, 1981, 1983, 1985/1986—which bears a coincidental resemblance to McCarthy 1979), there are two types of objects: *M-objects* (i.e., the objects of thought, which are intensional) and *actual objects* (which are extensional). There are two modes of predication of properties to these: M-objects are *constituted* by properties, and both M-objects and actual objects can *exemplify* properties. For instance, the pen with which I wrote the manuscript of this paper is an actual object that *exemplifies* the property of *being white.* Right now, when I think about that pen, the object of my thought is an M-object that is *constituted* (in part) by that property. The M-object *Jan's pen* can be represented as: ⟨belonging to Jan, being a pen⟩ (or, for short, as: $\langle J,P \rangle$). *Being a pen* is also a *constituent* of this M-object: $P \mathrel{\mathrm{c}} \langle J,P \rangle$; and 'Jan's pen is a pen' is true in virtue of this objective. In addition, $\langle J,P \rangle$ *exemplifies* (ex) the property of *being constituted by two properties.* There might be an actual object, say, α, corresponding to $\langle J,P \rangle$, that *exemplifies* the property of *being a pen* (α ex P) as well as (say) the property of *being 6 inches long.* But *being 6 inches long* $\not\mathrm{c} \langle J,P \rangle$.

The M-object *the round square,* $\langle R,S \rangle$, is constituted by precisely two properties: being round (R) and being square (S); 'The round square is round' is true in virtue of this, and 'The round square is not square' is false in virtue of it. But $\langle R,S \rangle$ exemplifies neither of those properties, and 'The round square is not square' is *true* in virtue of *that.* That is, 'is' is ambiguous.

An M-object o exists if and only if there is an actual object α that is "Sein-correlated" with it: o exists if and only if $\exists\alpha[\alpha \mathrm{SC} o]$ if and only if $\exists\alpha\forall F[F \mathrel{\mathrm{c}} o \rightarrow \alpha \mathrel{\mathrm{ex}} F]$. Note that incomplete objects, such as $\langle J,P \rangle$, can exist. However, the M-object *the existing golden mountain,* $\langle E,G,M \rangle$, has the property of

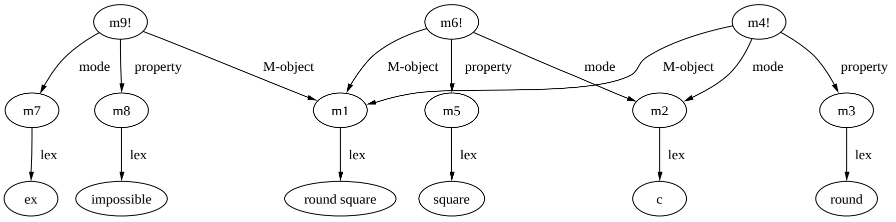

Figure 4: A SNePS representation of
“The round square is round”, “The round square is square”,
and “The round square is impossible”, on Rapaport’s theory.

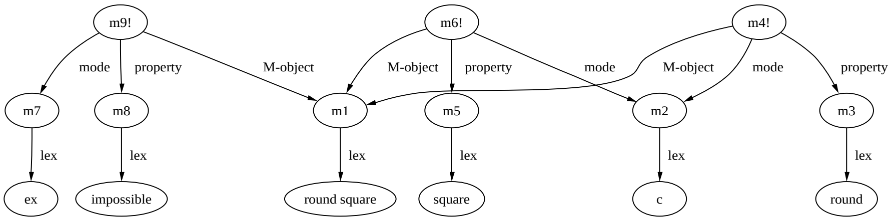

Figure 5: An alternative SNePS representation of
“The round square is round”, “The round square is square”,
and “The round square is impossible”, on Rapaport’s theory.

existing (because E c $\langle E,G,M\rangle$) but does not exist (because $\neg\exists\alpha[\alpha$ SC $\langle E,G,M\rangle]$, as an empirical fact).

The intensional fragment of this theory can be used to provide a semantics for SNePS in much the same way that it can been used to provide a semantics for natural language (Rapaport 1981, 1988). SNePS base nodes can be taken to represent M-objecta and properties; SNePS propositional nodes can be taken to represent M-objectives. Two alternatives for networks representing the three M-objectives: R c $\langle R,S\rangle$, S c $\langle R,S\rangle$, and $\langle R,S\rangle$ ex *being impossible* are shown in Figures 4 and 5. (Also, the second can be used to avoid “Clark’s paradox”; see Rapaport 1978, 1983; Poli 1998.)

Figure 4 can be read as follows: Node m4 represents the M-objective that *round* is a “c”onstituent of the “M-object” *the round square*. Node m6 represents the M-objective that *square* is a “c”onstituent of the “M-object” *the round square*. And node m9 represents the M-objective that the “M-object” *the round square* “ex”emplifies being impossible.

Figure 5 can be read as follows: Node m4 represents the M-objective that *round* is a “property” that the “M-object” *the round square* has under the “c” (constituency) "mode" of predication. Node m6 represents the M-objective that *square* is a "property" that the "M-object" *the round square* has under the "c" (constituency) "mode" of predication. And node m9 represents the M-objective that the "M-object" *the round square* has the "property" *being impossible* under the "ex" (exemplification) "mode" of predication.

The difference between the representations in the two figures is that in Figure 5, but not in Figure 4, it is possible to talk about constituency and exemplification.

Actual (i.e., extensional) objects, however, should not be represented (cf. Maida & Shapiro 1982: 296–298). To the extent to which such objects are essential to this Meinongian theory, the present theory is perhaps an inappropriate one. (A similar remark holds, of course, for McCarthy 1979.)

The distinction between two modes of predication, shared by my theory and Castañeda's (see below) has its advantages, however. Consider the problem of relative clauses (Woods 1975: 60ff): How should

> (*) The dog that bit the man has rabies

be represented? Here is a Meinongian solution along the lines of Rapaport 1981:

> ⟨being a dog, having bit a man⟩ ex having rabies.

The fact that

> for every Meinongian object o, if $o = \langle \ldots\ F\ \ldots \rangle$, then F c o

can then be used to infer the sentence:

> The dog bit the man. (Or: A dog bit the man.)

That is, the difference between information in the relative clause and the information in the main clause is (or can be represented by) the difference between internal and external predication; it is the difference between defining and asserted properties (see §9, below). This analysis is related to the semantic Principles of Minimization of Ambiguity and of Maximization of Truth advocated in Rapaport 1981: 13f: In the absence of prior context, this analysis is correct for (*). But a full computational account would include something like the following:

```
If there is a unique dog that bit a (specified) man,
    then use the representation of that dog as subject
    else build: ⟨being a dog, having bit a man⟩ ex having rabies.
```

9 Parsons's Theory

Terence Parsons's theory of nonexistent objects (1980; cf. Rapaport 1976, 1978, 1985a) recognizes only one type of object—intensional ones—and only one mode of predication. But it has two types of properties: *nuclear* and *extranuclear.* The former includes all "ordinary" properties such as: being red, being round, etc.; the latter includes such properties as: existing, being impossible, etc. But the distinction is

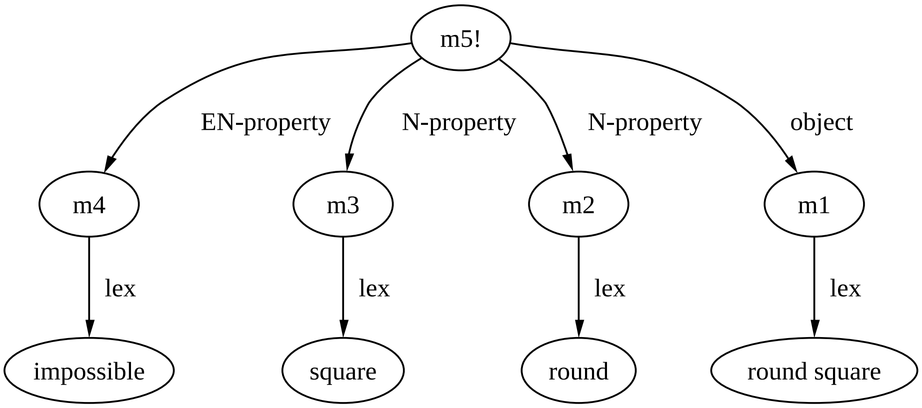

Figure 6: A SNePS representation of
“The round square is round, square, and impossible”,
on Parsons’s theory.

blurry, since for each extranuclear property, there is a corresponding nuclear one. For every set of nuclear properties, there is a unique object that has only those properties. Existing objects must be complete (and, of course, consistent), though not all such objects exist. For instance, *the Morning Star* and *the Evening Star* don’t exist (if these are taken to consist, roughly, of only two properties each). *The round square*, of course, is (and only is) both round and square and, so, isn’t non-square; though it is, for that reason, impossible, hence not real. As for *the existing golden mountain*, *existence* is extranuclear, so the set of these three properties doesn’t have a corresponding object. There is, however, a “watered-down”, nuclear version of existence, and there *is* an existing golden mountain that has *that* property; but *it* doesn’t have the extranuclear property of existence, so it doesn’t *exist*.

Parsons’s theory could provide a semantics for SNePS, though the use of two types of properties places restrictions on the possible uses of SNePS. On the other hand, SNePS could be used to represent Parsons’s theory (though a device would be needed for marking the distinction between nuclear and extranuclear properties) and, hence, together with Parsons’s natural-language semantics, to provide a tool for computational linguistics. Figure 6 suggests one way that this might be done. (Node m5 represents the proposition that the Meinongian “object” *the round square* has *round* and has *square* as “N”uclear“-properties” and has *being impossible* as an “E”xtra“N”uclear“-property”.)

There is a potential problem with Parsons’s theory. As Woods points out, it is important to distinguish between defining and asserted properties of a node (Woods 1975: 53). Suppose there is a node representing John’s height and suppose that John’s height is greater than Sally’s height. We need to represent that the former *defines* the node and that the latter *asserts* something non-defining of it. This is best done by means of a distinction between internal and external predication, as on my theory or Castañeda's (see below). It could perhaps be done with the nuclear/extranuclear distinction, but less suitably, since *being John's height* and *being greater than Sally's height* are both *nuclear* properties. (Note that this is *not* the same as the structural/assertional distinction among types of links; cf. Woods 1975: 58f.)

10 Castañeda's Theory

Hector-Neri Castañeda's theory of "guises" (1972, 1975abc, 1977, 1979, 1980, 1989; cf. Rapaport 1976, 1978, 2005) is a better candidate. It is a fully intensional theory with one type of object: *guises* (intensional items corresponding to sets of properties), and one type of property. More precisely, there are properties (for example, *being round*, *being square*, *being blue*, ...), sets of these (called *guise cores*; for example, $\{$*being round*, *being square*$\}$), and an ontic counterpart, *c*, of the definite-description operator, which is used to form *guises*; for example, $c\{$*being round*, *being square*$\}$ is the round square. Guises can be understood, roughly, as things-under-a-description, as "facets" of (physical and non-physical) objects, as "roles" that objects play, or, in general, as objects of thought.

Guise theory has two modes of predication: *internal* and *external.* In general, the guise $c\{\ldots F \ldots\}$ is-internally F. For example, the guise (named by) *the round square* is-internally only round and square. The two guises *the tallest mountain* and *Mt. Everest* are related by an external mode of predication called *consubstantiation* (C^*). Consubstantiation is an equivalence relation that is used in the analyses of (1) external predication, (2) co-reference, and (3) existence: Let $a = c\{\ldots F \ldots\}$ be a guise and let $a[G] =_{df} \mathrm{c}(\{\ldots F \ldots\} \cup \{G\})$. Then (1) a is-externally G (in one sense) if $C^*(a,\ a[G])$. For instance, 'the Morning Star is a planet' is true because $C^*(c\{M,S\},c\{M,S,P\})$; i.e., *the Morning Star* and *the Morning Star that is a planet* are consubstantiated. (2) Guise a "is the same as" guise b if and only if C^*ab. For instance, 'the Morning Star is the same as the Evening Star' is true because $C^*(c\{M,S\},c\{E,S\})$. And (3) a exists if and only if there is a guise b such that C^*ab.

Another external mode of predication is *consociation* (C^{**}). This is also an equivalence relation, but one that holds between guises that a mind has "put together", that is, between guises in a "belief space". For instance, C^{**}(Hamlet, the Prince of Denmark).

C^* and C^{**} correspond almost exactly to the use of the EQUIV arc in SNePS. Maida and Shapiro (1982: 303f) use the EQUIV-EQUIV case-frame to represent co-reference (which is what C^* is), but, as I have suggested in Rapaport 1986b, EQUIV-EQUIV more properly represents *believed* co-reference—which is what C^{**} is. It should be clear how guise theory can provide a semantics for SNePS. Figure 7 suggests how this might be done. (Node m3 is the guise *the evening star*, whose "core-properties" are *being seen in the evening* and *being starlike*. Node m5 is the guise *the morning star*, whose "core-properties" are *being seen in the morning* and *being starlike*. Node m6 expresses the proposition that m3 and m5 are consubstantiated. Similarly, node m8 is the guise whose "core-properties" are *being starlike*, *being seen* *in the morning*, and *being a planet* (the “planet-protraction of the morning star”, in Castañeda’s terminology), and node m9 expresses the proposition that m5 and m8 are consubstantiated.)

Some problems remain, however: in particular, the need to provide a SNePS correlate for internal predication and the requirement of explicating external predication in terms of relations like C^*. Note, too, that nodes m3, m5, and m8 in Figure 7 are “structured individuals”—a sort of molecular base node.

11 Conclusion

How should we decide among these theories? Consider another comment by Woods:

> Whereas previously we construed our nodes to correspond to real existing objects, now we have introduced a new type of node which does not have this assumption. Either we now have two very different types of nodes (in which case we must have some explicit . . . mechanism in the notation to indicate the type of every node) or else we must impose a unifying interpretation. . . . One possible unifying interpretation is to interpret every node as an intensional description and assert an explicit predicate of existence for those nodes which are intended to correspond to real objects. (Woods 1975: 66f.)

The two-types-of-nodes solution is represented by my theory (and by McCarthy’s); the unified theory is Castañeda’s (with self-consubstantiation as the existence predicate). Thus, Woods’s ideal as well as SNePS are closer to Castañeda’s theory. Or, one could take the intensional fragment of my theory and state that $\exists\alpha[\alpha \text{ SC } o]$ if and only if *o* ex Existence.

Or consider Maida and Shapiro again:

> [W]e should be able to describe within a semantic network any conceivable concept, independently of whether it is realized in the actual world, and we should also be able to describe whether in fact it is realized. (Maida & Shapiro 1982: 297.)

The latter is harder. We would need either (a) to represent extensional entities (as could be done on my theory, using SC), or (b) to represent a special existence predicate (as on Parsons’s theory, using extranuclear existence), or (c) to use some co-referentiality mechanism (as in SNePS *and* in Castañeda’s theory), or (d) to conflate two such nodes into one (which brings us back to the first solution but doesn’t eliminate the need for intensional entities; cf. Maida & Shapiro 1982: 299).

In any case, I hope to have provided evidence that it is possible to provide a fully intensional, non-possible-worlds semantics for SNePS and similar semantic-network formalisms. The most straightforward way is to use Meinong’s theory of objects, though his original theory has the disadvantage of not being formalized. As we have seen, there are several extant formal Meinongian theories that can be used, though each has certain disadvantages or problems.

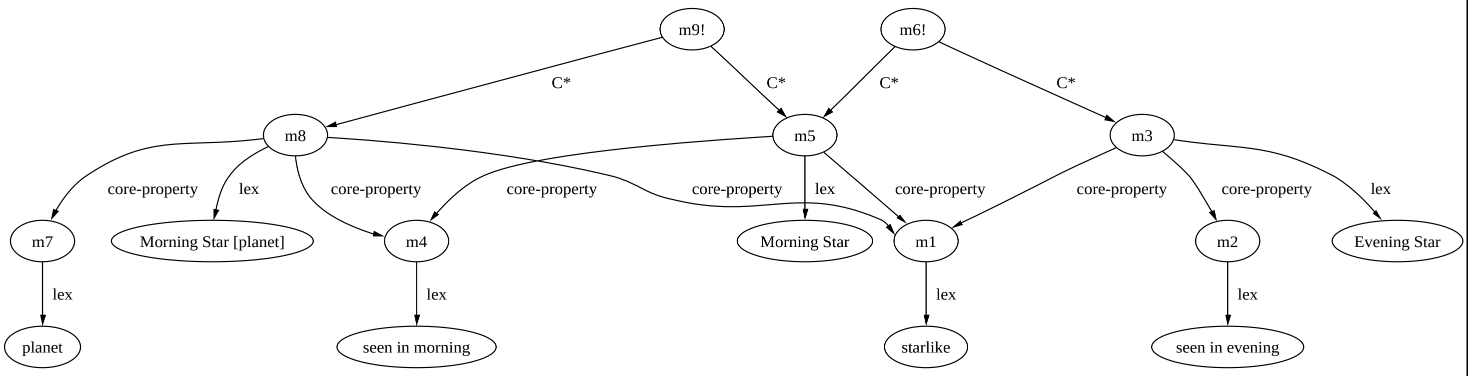

Figure 7: A SNePS representation of "The Morning Star is the Evening Star" (m6) and "The Morning Star is a planet" (m9) on Castañeda's theory.

Two lines of research are currently being investigated: (1) Take SNePS as is, and provide a new, formal Meinongian theory for its semantic foundation. This has not been discussed here, but the way to do this should be clear from the possibilities examined above. My own theory (stripped of its extensional fragment) or a modification of Castañeda's theory seem the most promising approaches. (2) Modify SNePS so that one of the extant formal Meinongian theories can be so used. SNePS is, in fact, currently being modified by the SNePS Research Group—for independent reasons—in ways that make it closer to Castañeda's guise theory, by the introduction of structured individuals—"base nodes" with descending arcs for indicating their "internal structure".

Thus, although philosophy may not have been kind to Meinong, perhaps AI will be.

12 Postscript 2005

This essay was written a long time ago, in March 1985. In the intervening two decades, much progress has been made that is not reflected in the essay you have just read. I have, however, updated some of the references and notation, and the promissory notes with respect to an intensional semantics for SNePS have since been cashed, in part, in Shapiro & Rapaport 1987, 1991, and Shapiro et al. 1996.

Acknowledgments

This research was supported in part by SUNY Buffalo Research Development Fund grant #150-9216-F and National Science Foundation Grant No. IRI-8610517. An earlier version appeared as Rapaport 1985b. I am grateful to Stuart C. Shapiro, Hector-Neri Castañeda, and the members of the SNePS Research Group for comments and discussion.

References

Ali, Syed, & Shapiro, Stuart C. (1993), "Natural Language Processing Using a Propositional Semantic Network with Structured Variables", *Minds and Machines* 3: 421–451.

Barwise, Jon, & Perry, John (1983), *Situations and Attitudes* (Cambridge, MA: MIT Press).

Bobrow, Daniel G., & Winograd, Terry (1977), "An Overview of KRL, a Knowledge Representation Language", *Cognitive Science* 1: 3–46.

Brachman, Ronald J. (1977), "What's in a Concept: Structural Foundations for Semantic Networks", *International Jouranl of Man-Machine Studies* 9: 127–152.

Brachman, Ronald J. (1979), "On the Epistemological Status of Semantic Networks", in Findler 1979: 3–50.

Brachman, Ronald J. (1983), "What IS-A Is and Isn't: An Analysis of Taxonomic Links in Semantic Networks", *Computer* 16 (October 1983) 30–36.

Brachman, Ronald J., & Levesque, Hector J. (eds.) (1985), *Readings in Knowledge Representation* (Los Altos, CA: Morgan Kaufmann).

Brachman, Ronald J., & Schmolze, James G. (1985), "An Overview of the KL-ONE Knowledge Representation System", *Cognitive Science* 9: 171–216.

Brand, Myles (1984), *Intending and Acting: Toward a Naturalized Action Theory* (Cambridge, MA: MIT Press).

Castañeda, Hector-Neri (1972), "Thinking and the Structure of the World", *Philosophia* 4 (1974) 3–40; reprinted in 1974 in *Critica* 6 (1972) 43–86.

Castañeda, Hector-Neri (1975a), "Individuals and Non-Identity: A New Look", *American Philosophical Quarterly* 12: 131–140.

Castañeda, Hector-Neri (1975b), "Identity and Sameness", *Philosophia* 5: 121–150.

Castañeda, Hector-Neri (1975c), *Thinking and Doing: The Philosophical Foundations of Institutions* (Dordrecht: D. Reidel).

Castañeda, Hector-Neri (1977), "Perception, Belief, and the Structure of Physical Objects and Consciousness", *Synthese* 35: 285–351.

Castañeda, Hector-Neri (1979), "Fiction and Reality: Their Basic Connections", *Poetica* 8: 31–62.

Castañeda, Hector–Neri (1980), "Reference, Reality, and Perceptual Fields", *Proceedings and Addresses of the American Philosophical Association* 53: 763–823.

Castañeda, Hector-Neri (1989), *Thinking, Language, and Experience* (Minneapolis: University of Minnesota Press).

Charniak, Eugene, & McDermott, Drew (1985), *Introduction to Artificial Intelligence* (Reading, MA: Addison-Wesley).

Findlay, J.N. (1963), *Meinong's Theory of Objects and Values, 2nd edition* (Oxford: Clarendon Press).

Findler, Nicholas V. (1979), *Associative Networks: The Representation and Use of Knowledge by Computers* (New York: Academic Press).

Fine, Kit (1983), "A Defence of Arbitrary Objects", *Proceedings of the Aristotelian Society*, Supp. Vol. 58: 55–77.

Grossmann, Reinhardt (1974), *Meinong* (London: Routledge & Kegan Paul).

Hardt, Shoshana L. (1992), "Conceptual Dependency", in S.C. Shapiro (ed.), *Encyclopedia of Artificial Intelligence*, 2nd edition (New York: John Wiley): 259–265.

Hendrix, Gary G. (1979), "Encoding Knowledge in Partitioned Networks", in Findler 1979: 51–92.

Hirst, Graeme (1989), "Ontological Assumptions in Knowledge Representation", *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning (KR-89, Toronto)* (San Mateo, CA: Morgan Kaufmann): 157–169.

Hirst, Graeme (1991), "Existence Assumptions in Knowledge Representation", *Artificial Intelligence* 49: 199–242.

Israel, David J. (1983), "Interpreting Network Formalisms", in N. Cercone (ed.), *Computational Linguistics* (Oxford: Pergamon Press): 1–13.

Johnson, Frances L., & Shapiro, Stuart C. (2005a), "Dependency-Directed Reconsideration: Belief Base Optimization for Truth Maintenance Systems", in *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI-05)* (Menlo Park, CA: AAAI Press).

Johnson, Frances L., & Shapiro, Stuart C. (2005b), "Improving Recovery for Belief Bases", in Leora Morgenstern & Maurice Pagnucco (eds.), *Proceedings of the 6th Workshop on Nonmonotonic Reasoning, Action, and Change (IJCAI-05, Edinburgh)*.

Lambert, Karel (1983), *Meinong and the Principle of Independence* (Cambridge, UK: Cambridge University Press).

Lehmann, Fritz (ed.) (1992), *Semantic Networks in Artificial Intelligence* (Oxford: Pergamon Press).

Maida, Anthony S., & Shapiro, Stuart C. (1982), "Intensional Concepts in Propositional

Semantic Networks", *Cognitive Science* 6: 291–330; reprinted in Brachman & Levesque 1985: 170–189.

Martins, João & Shapiro, Stuart C. (1988), "A Model for Belief Revision", *Artificial Intelligence* 35: 25–79.

Martins, João, & Cravo, Maria R. (1991), "How to Change Your Mind", *Noûs* 25: 537–551.

McCarthy, John (1979), "First Order Theories of Individual Concepts and Propositions", in J.E. Hayes, D. Michie, & L. Mikulich (eds.), *Machine Intelligence* 9 (Chichester, UK: Ellis Horwood): 129–147; reprinted in Brachman & Levesque 1985: 523–533.

McDermott, Drew (1981), "Artificial Intelligence Meets Natural Stupidity", in John Haugeland (ed.), *Mind Design: Philosophy, Psychology, Artificial Intelligence* (Cambridge, MA: MIT Press): 143–160.

Meinong, Alexius (1904), "Über Gegenstandstheorie", in Rudolf Haller (ed.), *Alexius Meinong Gesamtausgabe*, Vol. II (Graz, Austria: Akademische Druck- u. Verlagsanstalt, 1971): 481–535; English translation ("The Theory of Objects") by Isaac Levi et al., in Roderick M. Chisholm (ed.), *Realism and the Background of Phenomenology* (New York: Free Press, 1960): 76–117.

Montague, Richard (1974), *Formal Philosophy*, Richmond H. Thomason (ed.) (New Haven: Yale University Press).

Parsons, Terence (1980), *Nonexistent Objects* (New Haven: Yale University Press).

Poli, Roberto (1998), Review of Dale Jacquette's *Meinongian Logic*, in *History and Philosophy of Logic* 2: 117–119.

Quillian, M. Ross (1967), "Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities", *Behavioral Science* 12: 410–430; reprinted in Brachman & Levesque 1985: 97–118.

Quillian, M. Ross (1968), "Semantic Memory", in Marvin Minsky (ed.), *Semantic Information Processing* (Cambridge, MA: MIT Press): 227–270.

Quillian, M. Ross (1969), "The Teachable Language Comprehender: A Simulation Program and Theory of Language", *Communications of the ACM* 12(8): 459–476.

Quine, Willard Van Orman (1980), "On What There Is", in W.V. Quine, *From a Logical Point of View* (Cambridge, MA: Harvard University Press, 2nd ed.): 1–19.

Rapaport, William J. (1976), *Intentionality and the Structure of Existence*, Ph.D. dissertation (Bloomington: Indiana University Department of Philosophy).

Rapaport, William J. (1978), "Meinongian Theories and a Russellian Paradox", *Noûs* 12: 153–80; errata, *Noûs* 13 (1979) 125.

Rapaport, William J. (1981), "How to Make the World Fit Our Language: An Essay in Meinongian Semantics", *Grazer Philosophische Studien* 14: 1–21.

Rapaport, William J. (1983), "Meinong, Defective Objects, and (Psycho-)Logical Paradox", *Grazer Philosophische Studien* 18: 17–39.

Rapaport, William J. (1984), Critical study of Routley's *Exploring Meinong's Jungle and Beyond*, in *Philosophy and Phenomenological Research* 44: 539–552.

Rapaport, William J. (1985a), "To Be and Not To Be", *Noûs* 19: 255–271.

Rapaport, William J. (1985b), "Meinongian Semantics for Propositional Semantic Networks", *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics (University of Chicago)* (Morristown, NJ: Association for Computational Linguistics): 43–48.

Rapaport, William J. (1985/1986), "Non-Existent Objects and Epistemological Ontology", *Grazer Philosophische Studien* 25/26: 61–95.

Rapaport, William J. (1986a), Review of Lambert's *Meinong and the Principle of Independence*, in *Journal of Symbolic Logic* 51: 248–252.

Rapaport, William J. (1986b), "Logical Foundations for Belief Representation", *Cognitive*

Science 10: 371–422.

Rapaport, William J. (1988), "Syntactic Semantics: Foundations of Computational Natural-Language Understanding", in James H. Fetzer (ed.) *Aspects of Artificial Intelligence* (Dordrecht, Holland: Kluwer Academic Publishers): 81–131; reprinted (with numerous errors) in Eric Dietrich (ed.) (1994), *Thinking Computers and Virtual Persons: Essays on the Intentionality of Machines* (San Diego: Academic Press): 225–273.

Rapaport, William J. (1991a), "Predication, Fiction, and Artificial Intelligence", *Topoi* 10: 79–111.

Rapaport, William J. (1991b), "Meinong, Alexius; I: Meinongian Semantics", in Hans Burkhardt & Barry Smith (eds.), *Handbook of Metaphysics and Ontology* (Munich: Philosophia Verlag): 516–519.

Rapaport, William J. (1992), "Belief Representation Systems", in Stuart C. Shapiro (ed.), *Encyclopedia of Artificial Intelligence, 2nd edition* (New York: John Wiley): 98–110.

Rapaport, William J. (2000), "How to Pass a Turing Test: Syntactic Semantics, Natural-Language Understanding, and First-Person Cognition", *Journal of Logic, Language, and Information*, 9(4): 467–490; reprinted in James H. Moor (ed.), *The Turing Test: The Elusive Standard of Artificial Intelligence* (Dordrecht: Kluwer, 2003): 161–184.

Rapaport, William J. (2005), "Castañeda, Hector Neri", in *Dictionary of Modern American Philosophers* (Bristol, UK: Thoemmes Continuum).

Rapaport, William J., & Shapiro, Stuart C. (1984), "Quasi-Indexical Reference in Propositional Semantic Networks", *Proceedings of the 10th International Conference on Computational Linguistics (COLING–84; Stanford University)* (Morristown, NJ: Association for Computational Linguistics): 65–70.

Rapaport, William J.; Shapiro, Stuart C.; & Wiebe, Janyce M. (1997), "Quasi-Indexicals and Knowledge Reports", *Cognitive Science* 21: 63–107; reprinted in Francesco Orilia & William J. Rapaport (eds.), *Thought, Language, and Ontology: Essays in Memory of Hector-Neri Castañeda* (Dordrecht: Kluwer Academic Publishers, 1998): 235–294.

Rich, Elaine, & Knight, Kevin (1991), *Artificial Intelligence, 2nd edition* (New York: McGraw-Hill).

Routley, Richard (1979), *Exploring Meinong's Jungle and Beyond* (Canberra: Australian National University, Research School of Social Sciences, Department of Philosophy).

Schank, Roger C., & Rieger, Charles J., III (1974), "Inference and the Computer Understanding of Natural Language", *Artificial Intelligence* 5: 373–412; reprinted in Brachman & Levesque 1985: 119–139.

Schubert Kalsi, Marie-Luise (1987), *Meinong's Theory of Knowledge* (Dordrecht: Martinus Nijhoff).

Shapiro, Stuart C. (1971a), "The MIND System: A Data Structure for Semantic Information Processing", *Report R-837-PR* (Santa Monica: Rand Corporation).

Shapiro, Stuart C. (1971b), "A Net Structure for Semantic Information Storage, Deduction and Retrieval", *Proceedings of the 2nd International Joint Conference on Artificial Intelligence (IJCAI–71; Imperial College, London)* (Los Altos, CA: Morgan Kaufmann), Vol. 2, pp. 512–523.

Shapiro, Stuart C. (1978), "Path-Based and Node-Based Inference in Semantic Networks", in David Waltz (ed.), *Theoretical Issues in Natural Language Processing (TINLAP-2)* (New York: Association for Computing Machinery): 219–225.

Shapiro, Stuart C. (1979), "The SNePS Semantic Network Processing System", in Findler 1979: 179–203.

Shapiro, Stuart C. (1982), "Generalized Augmented Transition Network Grammars For Generation From Semantic Networks", *American Journal of Computational Linguistics* 8: 12–25.

Shapiro, Stuart C. (1986), "Symmetric Relations, Intensional Individuals, and Variable Binding", *Proceedings of the IEEE* 74: 1354–1363.

Shapiro, Stuart C. (2000a), "SNePS: A Logic for Natural Language Understanding and Commonsense Reasoning", in Łucja M. Iwańska & Stuart C. Shapiro (eds.), *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language* (Menlo Park, CA/Cambridge, MA: AAAI Press/MIT Press): 175–195.

Shapiro, Stuart C. (2000b), "An Introduction to SNePS 3", in Bernhard Ganter & Guy W. Mineau (eds.), *Conceptual Structures: Logical, Linguistic, and Computational Issues, Lecture Notes in Artificial Intelligence 1867* (Berlin: Springer-Verlag): 510–524.

Shapiro, Stuart C., & Rapaport, William J. (1987), "SNePS Considered as a Fully Intensional Propositional Semantic Network", in Nick Cercone & Gordon McCalla (eds.), *The Knowledge Frontier: Essays in the Representation of Knowledge* (New York: Springer–Verlag): 262–315; shorter version appeared in *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI–86, Philadelphia)* (Los Altos, CA: Morgan Kaufmann, 1986): 278–283, a revised version of which appeared as "A Fully Intensional Propositional Semantic Network", in Leslie Burkholder (ed.), *Philosophy and the Computer* (Boulder, CO: Westview Press): 75–91.

Shapiro, Stuart C., & Rapaport, William J. (1991), "Models and Minds: Knowledge Representation for Natural-Language Competence", in Robert Cummins & John Pollock (eds.), *Philosophy and AI* (Cambridge, MA: MIT Press): 215–259.

Shapiro, Stuart C., & Rapaport, William J. (1992), "The SNePS Family", *Computers and Mathematics with Applications* 23: 243–275; reprinted in Fritz Lehmann (ed.), *Semantic Networks in Artificial Intelligence* (Oxford: Pergamon Press, 1992): 243–275.

Shapiro, Stuart C., & Rapaport, William J. (1995), "An Introduction to a Computational Reader of Narratives", in Judith Felson Duchan, Gail A. Bruder, & Lynne E. Hewitt (eds.), *Deixis in Narrative: A Cognitive Science Perspective* (Hillsdale, NJ: Lawrence Erlbaum Associates): 79–105.

Shapiro, Stuart C.; Rapaport, William J.; Cho, Sung-Hye; Choi, Joongmin; Feit, Elissa; Haller, Susan; Kankiewicz, Jason; & Kumar, Deepak (1996), "A Dictionary of SNePS Case Frames" [http://www.cse.buffalo.edu/sneps/Manuals/dictionary.ps].

Sowa, John (1991), *Principles of Semantic Networks: Explorations in the Representation of Knowledge* (San Mateo, CA: Morgan Kaufmann Publishers).

Sowa, John (1992), "Semantic Networks", in Stuart C. Shapiro (ed.), *Encyclopedia of Artificial Intelligence, 2nd edition* (New York: John Wiley): 1493–1511.

Sowa, John F. (2002), "Semantic Networks" [http://www.jfsowa.com/pubs/semnet.htm].

Srihari, Rohini K. (1981), "Combining Path-Based and Node-Based Inference in SNePS", *Technical Report No. 183* (Buffalo: SUNY Buffalo Department of Computer Science).

Srihari, Rohini K., & Rapaport, William J. (1989), "Extracting Visual Information From Text: Using Captions to Label Human Faces in Newspaper Photographs", *Proceedings of the 11th Annual Conference of the Cognitive Science Society (Ann Arbor, MI)* (Hillsdale, NJ: Lawrence Erlbaum Associates): 364–371.

Wiebe, Janyce M., & Rapaport, William J. (1986), "Representing *De Re* and *De Dicto* Belief Reports in Discourse and Narrative", *Proceedings of the IEEE* 74: 1405–1413.

Woods, William A. (1975), "What's in a Link: Foundations for Semantic Networks", in Daniel G. Bobrow & Allan M. Collins (eds.), *Representation and Understanding* (New York: Academic Press): 35–82; reprinted in Brachman & Levesque 1985: 217–241.

Woods, William A., & Schmolze, James G. (1992), "The KL-ONE Family", in Fritz Lehmann (ed.), *Semantic Networks in Artificial Intelligence* (Oxford: Pergamon Press): 133–177.

Zalta, Edward N. (1983), *Abstract Objects* (Dordrecht: D. Reidel).