Meinongian Semantics and Artificial Intelligence

William J. Rapaport[†]
rapaport@buffalo.edu

ABSTRACT

This essay describes computational semantic networks for a philosophical audience and surveys several approaches to semantic-network semantics. 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.

1. Meinong, Philosophy, and Artificial Intelligence

Philosophy was not kind to Meinong, the late-19th/early-20th-century cognitive scientist, until the 1970s renaissance in Meinong studies (Findlay, 1963; Grossmann, 1974; Rapaport, 1978; 1991b; Routley, 1979; Lambert, 1983; Schubert-Kalsi, 1987). Even so, his writings are often treated as curiosities (or worse) by mainstream philosophers. Meinong's contribution to philosophy can be characterized in terms of his thoroughgoing intensionalism. While some philosophers ridiculed or rejected this approach, some AI researchers – for largely independent, though closely related, reasons – argued for it. Here, I explore some of their arguments and show the relevance of Meinongian theories to research in AI.

2. Semantic Networks

Knowledge representation and reasoning (KRR) is an area of AI concerned with systems for representing, storing, retrieving, and inferring information in cognitively adequate and computationally efficient ways. The represented

[†] Department of Computer Science & Engineering and Center for Cognitive Science, University at Buffalo, Buffalo, NY 14260, USA.

information need not necessarily be true, so a better terminology is '*belief* representation' (Rapaport & Shapiro, 1984; Rapaport, 1986b; 1992; Rapaport et al. 1997).

A *semantic network* is a representational system consisting of a labeled, directed graph whose "nodes" (vertices) represent *objects* and whose "arcs" (edges, or "links", or "pointers") represent *binary relations* among them (Findler, 1979; Brachman & Levesque, 1985; Sowa, 1991; 1992; 2002; Lehmann, 1992). Woods (1975, p. 44) says, "The major characteristic of the semantic networks that distinguishes them from other candidates [for KR systems] is the characteristic notion of a link or pointer which connects individual facts into a total structure."

Quillian's (1967; 1968; 1969) early "semantic memory" 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 *w* was the entire network of nodes and arcs reachable by following directed arcs originating at the node representing *w*. *Inheritance* (or *hierarchical*) *networks* use such arc labels as "inst[ance]", "isa", and "property" to represent taxonomic structures (Bobrow &Winograd, 1977; Charniak & McDermott, 1985, pp. 22–27; Thomason, 1992; Brachman & Levesque, 2004, ch. 10; see Fig. 1). Schank's Conceptual Dependency representational scheme uses nodes to represent conceptual primitives, and arcs to represent dependencies and semantic case relations among them (Schank & Rieger, 1974; Brand, 1984, ch. 8; Rich & Knight, 1991, pp. 277–288; Hardt, 1992; Lytinen 1992). The idea is an old one: Networks like those of Quillian, and Bobrow & Winograd's KRL (1977), or Brachman's KL-ONE (Brachman, 1979; Brachman & Schmolze, 1985; Woods & Schmolze, 1992; and subsequent "description logics" – Brachman & Levesque, 2004, ch. 9) bear strong family resemblances to "Porphyry's Tree" (Fig. 2) – the mediaeval device used to illustrate the Aristotelian theory of definition by species and differentia.

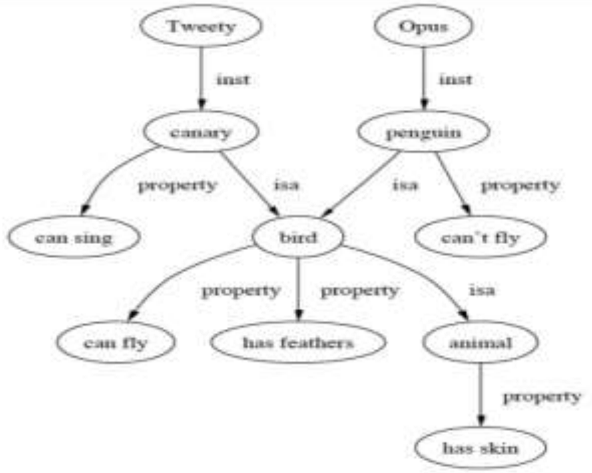

Figure 1: An inheritance network representing the propositions: Tweety is (an instance of) a canary; Opus is (an instance of) a penguin; A canary is a bird; A penguin is a bird; A canary can (i.e., has the property of being able to) sing; A penguin can't (i.e., has the property of not being able to) fly; A bird is an animal; A bird can fly; A bird has feathers; An animal has skin. However, the precise representations cannot be determined unambiguously from the network without a clearly specified syntax and semantics.

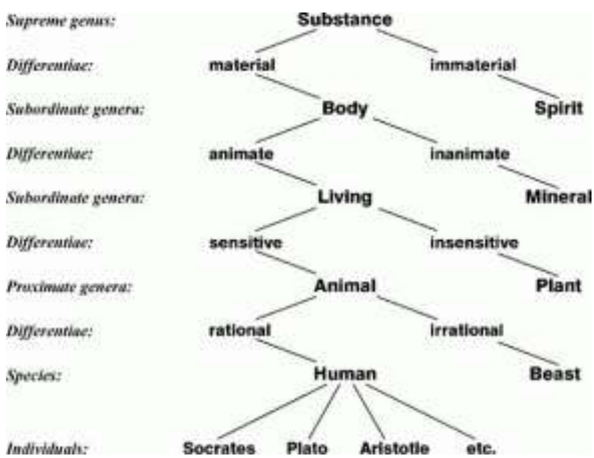

Figure 2: Porphyry's Tree: A mediaeval inheritance network (From Sowa 2002).

3. Semantics of Semantic Networks

Semantic networks are not essentially "semantic" (Hendrix, 1979; but cf. Woods, 1975; Brachman, 1979). Viewed as a data structure, a semantic network is a *language* (possibly with an associated logic or inference mechanism) for representing information about some domain. As such, it is a purely *syntactic* entity. They are called "semantic" primarily because of their *uses* as ways of representing the *meanings* of linguistic items. (However, this sort of syntax can be viewed as a kind of semantics, as in the so-called "Semantic Web"; cf. Rapaport 1988; 2000; 2003; 2012.)

As a notational device, a semantic network can itself be given a semantics. I.e., the arcs and nodes 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 (Woods, 1975; Brachman, 1977; 1983; and, especially, McDermott, 1981). E.g., in an inheritance network like that of Figure 1, how is the inheritance of properties to be represented or – more importantly – blocked? (If flying is a property inherited by the canary Tweety in virtue of its being a bird, what is to *prevent* the property of flying from being inherited by the flightless penguin Opus?) Do nodes represent *classes* of objects, *types* of objects, *individual* objects, or something else? Can *arcs* be treated as objects (perhaps with ("meta-")arcs linking *them* in some fashion)?

Providing a semantics for semantic networks is more akin to providing one for a *language* than for a *logic*. In the latter case, but not 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*; this would be given in terms of such a notion as *meaning*. 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 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.

Linguists and philosophers have argued for the importance of *intensional* semantics for natural languages (Montague, 1974; Parsons, 1980, Rapaport, 1981). At the same time, computational linguists and other AI researchers have recognized the importance of representing intensional *entities* (Woods, 1975; Brachman, 1979; McCarthy, 1979; Maida & Shapiro, 1982; Hirst, 1991). It seems reasonable that a semantics for such a representational system should itself be an intensional semantics.

In this essay, 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 the fully intensional, propositional Semantic Network Processing System (SNePS, [http://www.cse.buffalo.edu/sneps/]; Shapiro, 1979; 2000a; Shapiro & Rapaport, 1987; 1992; 1995), for which Israel (1983) offered a possible-worlds semantics. But possible-worlds semantics, while countenancing intensional entities, are not *fully* intensional: They treat intensional entities extensionally. Each L_W I discuss has 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 Maida & Shapiro (1982) argued that *only* intensional entities should be represented.

Woods (1975, pp. 38–40) 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 – and he characterizes *philosophical* semantics as the study of the relations between such a notation and truth conditions or meanings. 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 (1975, p. 45) proposes three "requirements of a good semantic representation": *logical adequacy* – it 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".

Logical adequacy 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, p. 36f). Some sentences *can* be interpreted as referring to nonexistents; so, a semantic network ought to be able to represent this, hence must be able to represent intensional entities. (The other criteria are discussed in §5.)

A second reason 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, p. 45). Therefore, 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 (see the argument in Woods, 1975, p. 46). Second, no computational efficiency would be gained by having them (Woods, 1975; p. 47). Third, it should not be done if one is interested in adequately representing *human* processing (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, p. 48).

A third argument for the need to represent intensional objects comes from consideration of question-answering programs (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. (Discourse Representation Theory uses a similar technique; Kamp & Reyle, 1993.)

Finally, intensional nodes are clearly needed for the representation of verbs of propositional attitude (Woods, 1975, p. 67; Rapaport & Shapiro, 1984; Rapaport, 1986b; 1992; Wiebe & Rapaport, 1986; Rapaport et al., 1997), and they can be used in quantificational contexts to represent "variable entities" (Woods, 1975, p. 68ff; Fine,1983; Shapiro, 1986; 2000b; 2004; Ali & Shapiro, 1993). Maida & Shapiro (1982) claims that, although semantic networks *can* represent real-world (extensional) entities or linguistic items, they *should*, for certain purposes, *only* represent intensional ones, especially when representing referentially opaque contexts (e.g., belief, knowledge), the concept of a truth value (as in 'John wondered whether *P*'), and 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 (and perhaps *only* can) 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 (inten*T*ional) object of thought, hence inten*S*ional. (Rapaport, 2012, §3.1, surveys arguments for this "narrow" or "internal" perspective.) In order to link the mind to the actual world (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, p. 300). The "link" would be forged by connections to other intensional nodes or by consistent input-output behavior that improves over time (Rapaport, 1985/1986, pp. 84–85; Rapaport, 1988; Srihari & Rapaport, 1989; Shapiro & Rapaport, 1991 surveys the wide variety of items that can be represented by intensional entities).

5. SNePS

A SNePS semantic network consists of labeled nodes and labeled, directed arcs satisfying the Uniqueness Condition (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, p. 179; Shapiro & Rapaport, 1991). When SNePS is used to model "the belief structure of a thinking, reasoning, language using being" (Maida & Shapiro, 1982, p. 296; cf. Shapiro, 1971b, p. 513), the concepts are the objects of mental (i.e., inten*T*ional) acts such as thinking, believing, wishing, etc. Such objects are inten*S*ional (cf. Rapaport, 1978).

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* relations between concepts, then those relations must be represented by nodes, since they have then become objects of thought, i.e., concepts. If "to be is to be the value of a [bound] variable" (Quine, 1980, p. 15; cf. Shapiro 1971a, pp. 79–80), then nodes represent such values; arcs do not. I.e., 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 '→' and '+' are to the symbols '*S*', '*NP*', and '*VP*' in the rewrite rule:

$S \rightarrow NP + VP.$

It is because propositions are represented by nodes and never by arcs that SNePS is a "propositional" semantic network (cf. Maida & Shapiro, 1982, p. 292). It can also be used to represent the inheritability of properties, either by explicit rules or by *path-based inference* (Shapiro, 1978; Srihari, 1981).

Figure 3 shows a sample SNePS network. Node m1 represents the *proposition* that [[b1]] (i.e., the thing represented by node 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* (rather than the world), 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 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 chosenby 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.

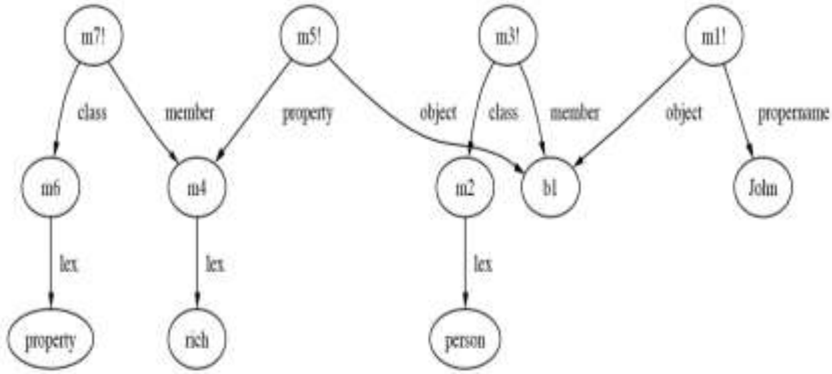

Figure 3: A SNePS representation for "A person named 'John' has the property of being rich". Where [[n]] represents the meaning of node n, [[m1]] = [[b1]] is named 'John'; [[m3]] = [[b1]] is a person; [[m5]] = [[b1]] is rich; [[m7]] = Being rich is a property.

SNePS satisfies Woods's three criteria (§4). Clearly, it is "logically" (better: representationally) adequate. Shapiro (1982) developed a generalized augmented-transition-network grammar 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" (better: inferential) processing (Shapiro, 1979; 1989; 1991; 2000a; Rapaport & Shapiro, 1984; Rapaport, 1986; 1991a; Shapiro & Rapaport, 1987; 1992; 1995; Martins & Shapiro, 1988; Martins & Cravo, 1991; Johnson & Shapiro, 2005a; Johnson & Shapiro, 2005b, Johnson, 2006, Fogel & Shapiro, 2011).

6. Israel's Possible-Worlds Semantics for SNePS

Israel's semantics for SNePS assumed "the general framework of Kripke-Montague style model-theoretic accounts" (Israel, 1983, p. 3), presumably because he took it as "quite clear that [Maida and Shapiro]...view their formalism as a Montague-type type-theoretic, intensional system" (Israel, 1983, p. 2). He introduced "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, p. 3). An *individual concept* is a function $ic : I \rightarrow D$. 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 are also represented in SNePS by constant individual nodes (§5) – were modeled as functions "from I into the power set of the set of individual concepts." Propositional nodes were modeled by "functions from I into $\{T,F\}$," although Israel felt that a "hyperintensional" logic would be needed in order to handle propositional attitudes.

Israel had difficulty interpreting `member`, `class`, and `isa` arcs in this framework. This is to be expected: 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 rejected it; cf. Maida & Shapiro, 1982, p. 297). Woods argues that a possible-worlds semantics is not psychologically valid, that the semantic representation must be finite (Woods, 1975, p. 50). Israel (1983, p. 5) himself hinted at the inappropriateness of this approach: "[I]f 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". 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). E.g., the object of my act of thinking of a unicorn is the objectum: *a unicorn*; the object of my act of believing that the Earth is flat is the objective: *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").

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 (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.

Meinong's theory and a fully intensional KRR system like SNePS are closely related. SNePS itself is much like *Aussersein*: All nodes are implicitly in the network all the time (Shapiro, personal communication). A SNePS base node (i.e., an atomic constant) represents an objectum; 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 (Maida & Shapiro, 1982; Rapaport & Shapiro, 1984, Rapaport, 1986b; Shapiro & Rapaport, 1991; Rapaport et al., 1997). When SNePS is used in a natural-language processing system (Shapiro, 1982; Rapaport & Shapiro, 1984; Rapaport, 1986; 1988; 1991a; Shapiro & Rapaport, 1995), 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 Russell on grounds of inconsistency: First, 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. Second, 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.)

Several formalizationsof Meinongiantheoriesovercome these problems.In §§8–10, 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's, 1979 AI theory), 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. E.g., 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 *Bill's pen* can be represented as: <belonging to Bill, being a pen> (or, for short, as: $<B,P>$). *Being a pen* is also a *constituent* of this M-object: $P \text{ c} <B,P>$; and 'Bill's pen is a pen' is true in virtue of this objective. In addition, $<B,P>$ *exemplifies* (ex) the property of *being constituted by two properties.* There might be an actual object, say, α, corresponding to $<B,P>$, that *exemplifies* the property of *being a pen* ($\alpha \text{ ex } P$) as well as (say) the property of *being 6 inches long.* But $\neg$(*being 6 inches long* c $<B,P>$).

The M-object *the round square*, $<R,S>$, 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 $<R,S>$ exemplifies neither of those properties, and 'The round square is not square' is *true* in virtue of *that.* I.e., 'is' is ambiguous.

An M-object o exists iff there is an actual object α that is "Sein-correlated" with it: $o \text{ exists iff } \exists\alpha[\alpha \text{SC} o]$ iff $\exists\alpha\forall F[F \text{ c } o \rightarrow \alpha \text{ ex } F]$. Note that incomplete objects, such as $<B,P>$, can exist. However, the M-object *the existing golden mountain*, $<E,G,M>$, has the property of existing (because $E \text{ c} <E,G,M>$) but does not exist (because $\neg\exists\alpha[\alpha \text{SC} <E,G,M>]$, 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). (Strict adherence to Fodorian methodological solipsism (Fodor, 1980) would seem to require that the Fodorian language of thought (LOT; Fodor, 1975) have syntax but no semantics. More recently, Fodor (2008, p. 16) suggests that LOT needs a purely *referential* semantics. Instead, I am proposing a *Meinongian* semantics for LOT, on the grounds that "non-existent" objects are best construed as internal mental entities.) SNePS base nodes can be taken to represent M-objectaand properties; SNePSpropositional nodes can be taken to represent M-objectives. Two alternatives for networks representing the three M-objectives: R c $<R,S>$, S c $<R,S>$, and $<R,S>$ ex *being impossible* are shown in Figures 4 and 5.

In Figure 4, 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.

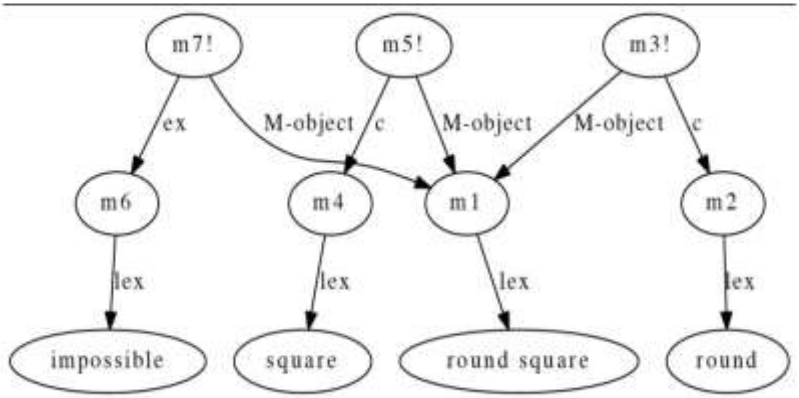

Figure 4: A SNePS representation of"The round square is round" (m3!), "The round square is square" (m5!),and "The round square is impossible" (m7!), on Rapaport's theory.

In Figure 5, 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.

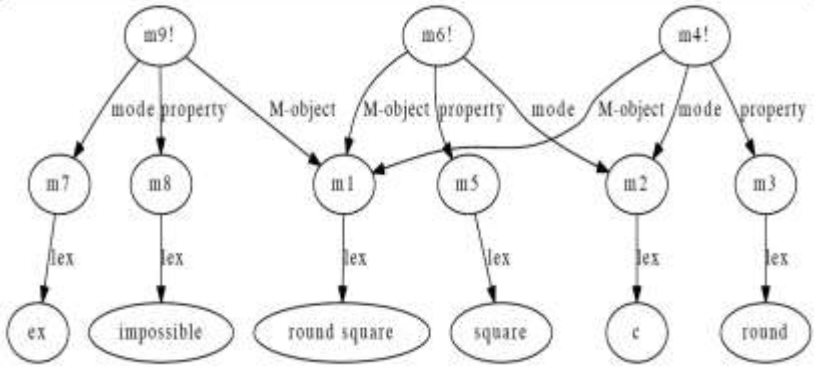

Figure 5: An alternative SNePS representation of "The round square is round" (m4!), "The round square is square" (m6!),and "The round square is impossible" (m9!), on Rapaport's theory.

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. (Also, the second can be used to avoid "Clark's paradox": Rapaport, 1978; 1983; Clark, 1983; Landini, 1985; Poli, 1998.)

Actual (i.e., extensional) objects, however, should *not* be represented (Maida & Shapiro, 1982, pp. 296–298). To the extent to which such objects are essential to this Meinongian theory, the present theory is perhaps an inappropriate one. (And similarly for McCarthy, 1979.)

The distinction between two modes of predication, shared by my theory and Castañeda's (§10), has its advantages. Consider the problem of relative clauses (Woods 1975, p. 60ff): How should sentence (1) (in §4) be represented? A Meinongian solution along the lines of Rapaport (1981) is:

<being a dog, having bit a man> ex having rabies.

The fact that,

for every Meinongian object *o*, if *o* = <...*F*...>, then *F* c *o*,

can then be used to infer the sentence:

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

I.e., 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, p. 13f). In the absence of prior context, this analysis is correct for (1). 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

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 blurry: 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. E.g., *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".

However, as Woods points out, it is important to distinguish between defining and asserted properties of a node (Woods, 1975, p. 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 (§10, 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. (This is *not* the same as the structural/assertional distinction among types of links; cf. Woods, 1975, p. 58f.)

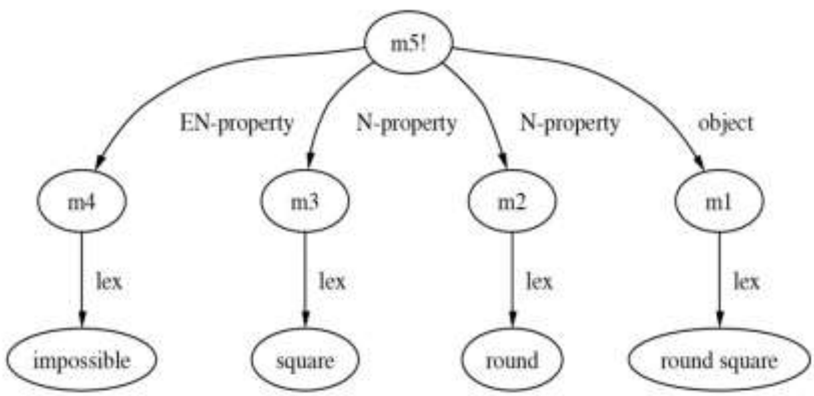

Figure 6: A SNePS representation of "The round square is round, square, and impossible" on Parsons's theory.

10. Castañeda's Theory

Castañeda's theory of "guises" (1972; 1975a; 1975b; 1975c; 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 (e.g., *being round*, *being square*, *being blue*), sets of these (called *guise cores*; e.g., {*being round*, *being square*}), and an ontic counterpart, *c*, of the definite-description operator, which is used to form *guises* from guise cores; e.g., *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*{...*F*...} is-internally *F*. E.g., 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 anequivalence relation that is used in the analyses of (1) externalpredication, (2) co-reference, and (3) existence:Let $a = c\{...F...\}$ be a guise, and let $a[G] =_{df} c(\{...F...\} \cup \{G\})$. Then (1) *a* is-externally *G* (in one sense) if $C^*(a, a[G])$. E.g., 'the Morning Star is a planet' is true because $C^*(c\{M,S\}, c\{M,S,P\})$; i.e., *the Morning Star* and *the*

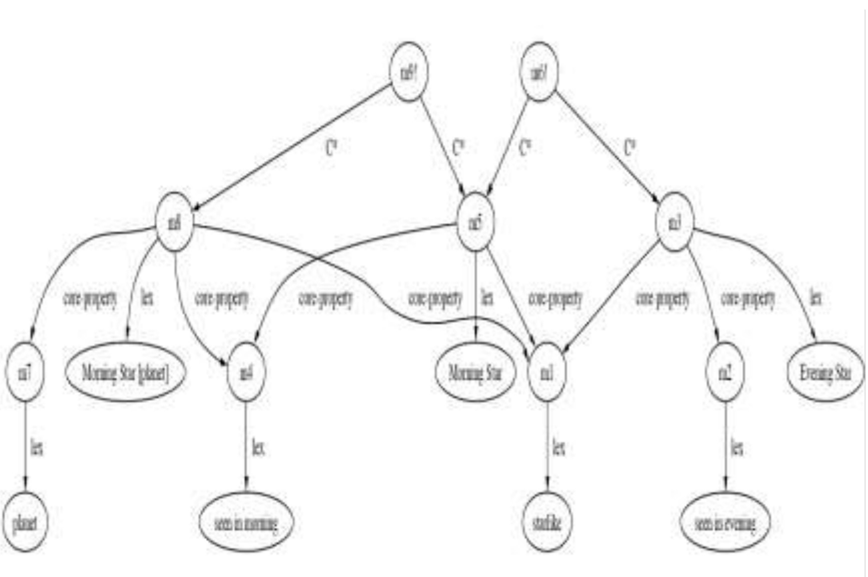

Figure 7: A SNePS representation of: "The Morning Star is the Evening Star" (m6) and "The Morning Star is a planet" (m9), on Castañeda's theory.

Morning Star that is a planet are consubstantiated. (2) Guise *a* "is the same as" guise *b* iff C^*ab. E.g., 'the Morning Star is the same as the Evening Star' is true because $C^*(c\{M,S\},c\{E,S\})$. And (3) *a* exists iff 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", i.e., between guises in a "belief space". E.g., C^{**}(Hamlet, the Prince of Denmark).

C^* and C^{**} correspond almost exactly to the use of the EQUIV arc in SNePS. Maida & Shapiro (1982, p. 303f) use the EQUIV-EQUIV case-frame to represent co-reference (which is what C^* is), but EQUIV-EQUIV more properly represents *believed* co-reference – which is what C^{**} is (Rapaport, 1986b). It should be clear how guise theory can provide a semantics for SNePS. In Figure 7, m3 represents the guise *the evening star*, whose "core-properties" are *being seen in the evening* and *being starlike*. Node m5 represents the guise *the morning star*, whose "core-properties" are *being seen in the morning* and *being starlike*. Node m6 represents the proposition that [[m3]] and [[m5]] are consubstantiated. Similarly, node m8 represents 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 represents the proposition that [[m5]] and [[m8]] are consubstantiated.

A remaining problem is 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? Woods said:

> 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, p. 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 $\exists\alpha[\alpha \text{ SC } o]$ iff *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, p. 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, p. 299).

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.

Two lines of research are possible: (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. (For more recent investigations into an intensional semantics for SNePS, see Wiebe & Rapaport, 1986; Shapiro & Rapaport, 1987; 1991; 1995; Rapaport & Shapiro, 1995; Shapiro et al., 1996; Rapaport et al., 1997; Rapaport, 2003; and Shapiro, 2003.) And a new version of SNePS (SNePS-3) is being designed that has several advantages, such as being able to represent "donkey" sentences and branching quantifiers (Ali 1993, 1994, 1995; Ali & Shapiro, 1993; Shapiro, 2000b).

Philosophy may not have been kind to Meinong; perhaps AI will be.

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. The essay has benefited from discussions with Shapiro, Castañeda, and members of the SNePS Research Group.

REFERENCES

Ali, S.S. (1993). A Structured Representation for Noun Phrases and Anaphora. In *Proc. 15th Annual Conf. Cog.Sci. Soc.*. Hillsdale, NJ: Erlbaum, 197–202.

Ali, S.S. (1994). A 'Natural Logic' for Natural Language Processing and Knowledge Representation. Buffalo, NY: SUNY Buffalo Dept. Comp. Sci., Tech. Rept. 94-01. PhD dissertation. Retrieved 10 April 2012 from http://www.cse.buffalo.edu/sneps/Bibliography/ali94-01.pdf

Ali, S.S. (1995). ANALOG: A Knowledge Representation System for Natural Language Processing. In Yfantis (Ed.), *Intelligent Systems*. Dordrecht: Kluwer, 327–332.

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

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

Bobrow, D.G., & Winograd, T. (1977). An Overview of KRL, a Knowledge Representation Language. *Cogn. Sci.*, *1*, 3–46.

Brachman, R.J. (1977). What's in a Concept: Structural Foundations for Semantic Networks.*Int'l. J. Man-Machine Studies*, *9*, 127–152.

Brachman, R.J. (1979). On the Epistemological Status of Semantic Networks. In Findler(1979), 3–50.

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

Brachman, R.J., & Levesque, H.J. (Eds.) (1985). *Readings in Knowledge Representation.* Los Altos, CA: Morgan Kaufmann.

Brachman, R.J., & Levesque, H.J. (2004). *Knowledge Representation and Reasoning*. Amsterdam: Elsevier.

Brachman, R.J., & Schmolze, J.G. (1985). An Overview of the KL-ONE Knowledge Representation System. *Cogn. Sci.*, *9*, 171–216.

Brand, M.(1984). *Intending and Acting*. Cambridge, MA: MIT Press.

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

Castañeda, H.-N. (1975a). Individuals and Non-Identity: A New Look. *Amer. Philosophical Quarterly*, *12*, 131–140.

Castañeda, H.-N. (1975b). Identity and Sameness. *Philosophia*, *5*, 121–150.

Castañeda, H.-N. (1975c). *Thinking and Doing*: Dordrecht: Reidel.

Castañeda, H.-N. (1977). Perception, Belief, and the Structure of Physical Objects and Consciousness. *Synthese*, *35*, 285–351.

Castañeda, H.-N. (1979). Fiction and Reality. *Poetica*, *8*, 31–62.

Castañeda, H.-N. (1980). Reference, Reality, and Perceptual Fields. *Proc.&Addr. Amer. Phil. Assn.*, *53*, 763–823.

Castañeda, H.-N. (1989). *Thinking, Language, and Experience*. Minneapolis: Univ. Minnesota Press.

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

Clark, R.L.(1983). Predication Theory: Guised and Disguised. In Tomberlin (Ed.), *Agent, Language, and the Structure of the World*. Indianapolis: Hackett, 111–129.

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

Findler, N.V. (1979). *Associative Networks*. New York: Academic.

Fine, K. (1983). A Defence of Arbitrary Objects. *Proc. Aristotelian Soc.,* Supp. Vol. *58*, 55–77.

Fodor, J.A. (1975). *The Language of Thought.* New York: Crowell.

Fodor, J.A. (1980). Methodological Solipsism Considered as a Research Strategy in Cognitive Psychology. *Behavior and Brain Sciences.*, 3, 63–109.

Fodor, J.A. (2008). *LOT 2: The Language of Thought Revisited.* Oxford: Clarendon.

Fogel, A.I., & Shapiro, S.C. (2011).On the Use of Epistemic Ordering Functions as Decision Criteria for Automated and Assisted Belief Revision in SNePS. In Sardina & Vassos (Eds.), *Proc. 9th Int'l. Workshop, Non-Monotonic Reasoning, Action, and Change (NRAC'11).* Melbourne: RMIT Univ., School of Comp. Sci. & Info. Tech., Tech. Rept.RMIT-TR-11-02, 31–38.

Grossmann, R. (1974). *Meinong*. London: Routledge.

Hardt, S.L. (1992). Conceptual Dependency. In Shapiro (Ed.), *Encyclopedia of Artificial Intelligence, 2nd ed.*. New York: Wiley, 259–265.

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

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

Israel, D.J. (1983). Interpreting Network Formalisms. In Cercone (Ed.), *Computational Linguistics.* Oxford: Pergamon, 1–13.

Johnson, F.L. (2006). Dependency-Directed Reconsideration. Buffalo: SUNY Buffalo Dept. Comp. Sci. & Eng'g.). PhD dissertation. Retrieved 17 April 2012 from http: //www.cse.buffalo.edu/sneps/Bibliography/joh06a.pdf

Johnson, F.L., & Shapiro, S.C. (2005a). Dependency-Directed Reconsideration. *Proc. 20th Nat'l. Conf. Artif.Intel. (AAAI-05).* Menlo Park, CA: AAAI Press.

Johnson, F.L., & Shapiro, S.C. (2005b). Improving Recovery for Belief Bases. in Morgenstern & Pagnucco (Eds.), *Proc. 6th Workshop, Nonmonotonic Reasoning, Action, and Change (IJCAI-05, Edinburgh).*

Kamp, H., & Reyle, U. (1993). *From Discourse to Logic.* Dordrecht: Kluwer.

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

Landini, G.(1985). Salvaging 'The F-er Is F': The Lesson of Clark's Paradox. *Philosophical Studies*, *48*, 129–136.

Lehmann, F. (Ed.) (1992). *Semantic Networks in Artificial Intelligence.* Oxford: Pergamon.

Lytinen, S.L.(1992). Conceptual Dependency and Its Descendants. *Computers Math. Applic.*, 23, 51–73. Reprinted in Lehmann (1992).

Maida, A.S., & Shapiro, S.C. (1982). Intensional Concepts in Propositional Semantic Networks. *Cogn. Sci.*, 6, 291–330.

Martins, J., & Shapiro, S.C. (1988). A Model for Belief Revision. *Artificial Intelligence., 35*, 25–79.

Martins, J., & Cravo, M.R. (1991). How to Change Your Mind. *Noûs*, *25*, 537–551.

McCarthy, J.(1979). First Order Theories of Individual Concepts and Propositions. In J.E. Hayes et al. (Eds.), *Machine Intelligence9.* Chichester, UK: Ellis Horwood, 129–147.

McDermott, D.(1981). Artificial Intelligence Meets Natural Stupidity. In J. Haugeland (Ed.), *Mind Design.* Cambridge, MA: MIT Press, 143–160.

Meinong, A. (1904). The Theory of Objects. Trans. I. Levi et al., in R.M. Chisholm (Ed.) (1960), *Realism and the Background of Phenomenology.* New York: Free Press, 76–117.

Montague, R. (1974). *Formal Philosophy*, R.H. Thomason (Ed.). New Haven: Yale Univ. Press.

Parsons, T.(1980).*Nonexistent Objects.* New Haven: Yale Univ. Press.

Poli, R. (1998). Review of D. Jacquette's *Meinongian Logic. Hist. Phil. Logic*, *19*, 117–119.

Quillian, M.R. (1967). Word Concepts. *Behav. Sci., 12*, 410–430.

Quillian, M.R.(1968). Semantic Memory. In M. Minsky (Ed.), *Semantic Information Processing*: Cambridge, MA: MIT Press, 227–270.

Quillian, M.R. (1969). The Teachable Language Comprehender. *Communications of the ACM, 12*, 459–476.

Quine, W.V.O. (1980). On What There Is. In W.V.O. Quine, *From a Logical Point of View*, 2nd ed.. Cambridge, MA: Harvard Univ. Press, 1–19.

Rapaport, W.J. (1976). *Intentionality and the Structure of Existence*. Indiana Univ. Dept. Phil., Bloomington. PhD dissertation.

Rapaport, W.J. (1978). Meinongian Theories and a Russellian Paradox. *Noûs 12*, pp. 153–180; errata, *Noûs, 13* (1979), 125.

Rapaport, W.J. (1981). How to Make the World Fit Our Language. *Grazer Phil. Studien, 14*, 1–21.

Rapaport, W.J. (1983). Meinong, Defective Objects, and (Psycho-)Logical Paradox. *Grazer Phil. Studien, 18*, 17–39.

Rapaport, W.J. (1984). Critical Notice of Routley's *Exploring Meinong's Jungle and Beyond. Philos. Phenomen. Res., 44*, 539–552.

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

Rapaport, W.J. (1985b). Meinongian Semantics for Propositional Semantic Networks. *Proc. 23rd Annual Meeting, Ass. Computational Ling.* (Univ. Chicago). Morristown, NJ: Assn. Computational Ling., 43–48.

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

Rapaport, W.J. (1986a). Review of Lambert's *Meinong and the Principle of Independence. J. Symbolic Log., 51*, 248–252

Rapaport, W.J. (1986b). Logical Foundations for Belief Representation. *Cogn. Sci., 10*, 371–422.

Rapaport, W.J.(1988). Syntactic Semantics. In Fetzer (Ed.), *Aspects of Artificial Intelligence*. Dordrecht: Kluwer, pp. 81–131. Errata retrieved 6 May 2011 fromhttp: //www.cse.buffalo.edu/~ rapaport/Papers/synsem.original.errata.pdf.

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

Rapaport, W.J. (1991b). Meinong, Alexius; I: Meinongian Semantics. In Burkhardt & Smith (Eds.), *Handbook of Metaphysics and Ontology.* Munich: PhilosophiaVerlag, 516–519.

Rapaport, W.J. (1992). Belief Representation Systems. In Shapiro (Ed.), *Encyclopedia of Artificial Intelligence, 2nd ed.*. New York: Wiley, 98–110.

Rapaport, W.J. (2000). How to Pass a Turing Test. *J.Logic, Lang., Info., 9*, 467–490.

Rapaport, W.J.(2003). What Did You Mean by That? Misunderstanding, Negotiation, and Syntactic Semantics. *Minds and Machines, 13*, 397–427.

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

Rapaport, W.J. (2012). Semiotic Systems, Computers, and the Mind: How Cognition Could Be Computing. *Int'l. J. Signs & Semiotic Syst., 2,*32–71.

Rapaport, W.J., & Shapiro, S.C. (1984).Quasi-Indexical Reference in Propositional Semantic Networks. *Proc. 10th Int'l. Conf. Comp.Ling. (COLING-84, Stanford Univ.).* Morristown, NJ: Assn. Comp. Ling., 65–70.

Rapaport, W.J., & Shapiro, S.C.(1995).Cognition and Fiction. In Duchan et al. (Eds.), *Deixis in Narrative*. Hillsdale, NJ: Erlbaum, 107–128.

Rapaport, W.J., Shapiro, S.C., & Wiebe, J.M.(1997). Quasi-Indexicals and Knowledge Reports.*Cogn. Sci., 21*, 63–107.

Rich, E., & Knight, K. (1991). *Artificial Intelligence, 2nd ed.*. New York: McGraw-Hill.

Routley, R.(1979). *Exploring Meinong's Jungle and Beyond.* Canberra: Australian National University., RSSS Dept. Phil.

Schank, R.C., & Rieger, C.J. (1974). Inference and the Computer Understanding of Natural Language. *Artif.Intel., 5,* 373–412.

Schubert Kalsi, M.-L.(1987). *Meinong's Theory of Knowledge*. Dordrecht: Nijhoff.

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

Shapiro, S.C. (1971b). A Net Structure for Semantic Information Storage, Deduction and Retrieval. *Proc. 2nd Int'l. Joint Conf. AI (IJCAI-71, Imperial College, London).* Los Altos, CA: Morgan Kaufmann, 2, 512–523.

Shapiro, S.C. (1978). Path-Based and Node-Based Inference in Semantic Networks. In Waltz (Ed.), *Theoretical Issues in Natural Language Processing (TINLAP-2).* New York: ACM, 219–225.

Shapiro, S.C. (1979). The SNePS Semantic Network Processing System. In Findler(1979), 179–203.

Shapiro, S.C. (1982). Generalized Augmented Transition Network Grammars for Generation from Semantic Networks. *Amer. J. Comp. Ling., 8,*12–25.

Shapiro, S.C. (1986). Symmetric Relations, Intensional Individuals, and Variable Binding. *Proc. IEEE, 74,*1354–1363.

Shapiro, S.C. (2000a). SNePS: A Logic for Natural Language Understanding and Commonsense Reasoning. In Iwanska& Shapiro (Eds.), *Natural Language Processing and Knowledge Representation.* Menlo Park, CA/Cambridge, MA: AAAI/MIT Press, 175–195.

Shapiro, S.C. (2000b). An Introduction to SNePS 3. In Ganter&Mineau (Eds.), *Conceptual Structures.* Berlin: Springer-Verlag LNAI 1867, 510–524.

Shapiro, S.C. (2003). MRS and SNePS: A Comparison. Retrieved 10 April 2012 from http: //www.cse.buffalo.edu/~ shapiro/Papers/sha03.pdf

Shapiro, S.C.(2004). A Logic of Arbitrary and Indefinite Objects. In D. Dubois et al. (Eds.), *Principles of Knowledge Representation and Reasoning: Proc. 9th Int'l. Conf. (KR2004),* Menlo Park, CA: AAAI Press, 565–575.

Shapiro, S.C., & Rapaport, W.J. (1987). SNePS Considered as a Fully Intensional Propositional Semantic Network. In Cercone & McCalla (Eds.), *The Knowledge Frontier.* New York: Springer–Verlag, 262–315.

Shapiro, S.C., & Rapaport, W.J. (1991). Models and Minds: Knowledge Representation for Natural-Language Competence. In Cummins & Pollock (Eds.), *Philosophy and AI.* Cambridge, MA: MIT Press, 215–259.

Shapiro, S.C., & Rapaport, W.J. (1992). The SNePS Family. *Computers Math. Applic., 23,* 243–275; reprinted in Lehmann (1992).

Shapiro, S.C., & Rapaport, W.J. (1995). An Introduction to a Computational Reader of Narratives. In Duchan et al. (Eds.), *Deixis in Narrative*. Hillsdale, NJ: Erlbaum, 79–105.

Shapiro, S.C., et al. (1996). A Dictionary of SNePS Case Frames. Retrieved 11 April 2012 from http: //www.cse.buffalo.edu/sneps/Manuals/dictionary.pdf

Sowa, J.F. (1991). *Principles of Semantic Networks*. San Mateo, CA: Morgan Kaufmann.

Sowa, J.F.(1992). Semantic Networks. In S.C. Shapiro (Ed.), *Encyclopedia of Artificial Intelligence, 2nd ed.*. New York: Wiley, 1493–1511.

Sowa, J.F. (2002). Semantic Networks. Retrieved 11 April 2012 from http: //www.jfsowa.com/pubs/semnet.htm

Srihari, R.K. (1981). Combining Path-Based and Node-Based Inference in SNePS. *Tech. Rep. 183*. Buffalo: SUNY Buffalo Dept. Comp. Sci. Retrieved 11 April 2012 from http: //www.cse.buffalo.edu/sneps/Bibliography/1981-9.pdf

Srihari, R.K., & Rapaport, W.J. (1989). Extracting Visual Information From Text. *Proc. 11th Annual Conf. Cogn. Sci.Soc.*.Hillsdale, NJ: Erlbaum, 364–371.

Thomason, R.H. (1992). NETL and Subsequent Path-Based Inheritance Theories. *Computers Math. Applic., 23*, 179–204.

Wiebe, J.M., & Rapaport, W.J. (1986). Representing De Re and De Dicto Belief Reports in Discourse and Narrative. *Proc. IEEE 74*, 1405–1413.

Woods, W.A. (1975). What's in a Link: Foundations for Semantic Networks. In D.G. Bobrow & A.M. Collins (Eds.), *Representation and Understanding*. New York: Academic, 35–82.

Woods, W.A., & Schmolze, J.G. (1992). The KL-ONE Family. *Computers Math. Applic., 23*,pp. 133–177; reprinted in Lehmann (1992).

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