Ontological assumptions in knowledge representation*

Graeme Hirst
Department of Computer Science
University of Toronto
Toronto, Canada M5S 1A4

Abstract

If knowledge representation formalisms are to be suitable for semantic interpretation of natural language, they must be more adept with representations of existence and non-existence than they presently are. I review the philosophical background, and exhibit some ontological problems for KR. I then look at the shortcomings of current approaches, including several intensional formalisms and the work of Hobbs. The Meinongian theory of Parsons is considered. Lastly, I present a naïve ontology for knowledge representation, identifying about nine distinct kinds of existence.

1 Introduction

Most contemporary logics implicitly or explicitly base the semantics of the quantifiers $\exists$ and $\forall$ on the widely-held ontological assumptions of Russell [1905, 1918] and Quine [1948]. A small but growing number of philosophers (*e.g.*, [Parsons 1980, Routley 1980, Lambert 1983] believe that these assumptions are mistaken,[1] and have proposed various alternatives. In this paper, I will discuss the consequences of the Russell–Quine assumptions for knowledge representation formalisms, and show that an adequate treatment requires a multi-faceted view of existence.

My motivation comes from the KR needs of natural language understanding. As I have argued elsewhere [Hirst 1988b], a KR formalism to be used in an NLU system for unrestricted text must have (at least) the expressive power of natural language (for otherwise it could not be a target language for semantic interpretation). Moreover, natural languages reflect genuine properties of the real world (with different languages possibly highlighting different properties or viewpoints). Thus, KR research may include exhibiting sentences of natural language and considering how their meaning, and the world it reflects, may be adequately represented—where 'adequately' means that the representation permits the same inferences to be drawn as the original sentence. Here, I am concerned with sentences that speak of existence, of non-existence, or of non-existent objects.

2 Three ontological slogans

2.1 Existence is not a predicate

Immanuel Kant, in his *Critique of pure reason* [1787, B.625ff], argued that existence is not a property that may be predicated of an entity the same way that properties like color and species can be.

Kant was responding to an argument by St Anselm of Canterbury [Anselm 1078, II] that purported to demonstrate the existence of God *a priori*: his 'ontological proof'. Anselm's argument was basically this: What we mean by God is, by definition, that entity that is right up the top end of the scale in all desirable properties: the entity that is most wise, most good, and so on. On the scale of existence, clearly actual or necessary existence is better than mere conceptual or possible existence; therefore existence is a defining property of God; therefore God exists.[2] Descartes [1641, V] later took much the same approach: God has all perfections; existence is a perfection; therefore God exists.[3]

*Parts of this paper were written while the author was at the Department of Artificial Intelligence, University of Edinburgh, with the support of a Visiting Fellowship from the U.K. Science and Engineering Research Council. The balance of the work was supported by the Natural Sciences and Engineering Research Council of Canada.

[1] Introducing his work, Parsons says of the Russell–Quine position that "clear progress is rare in philosophy, and I was pleased to have [it as] an example to cite. But as I thought about it more, I became increasingly dissatisfied" (p. xii).

[2] Compare Smullyan's proof [1978, p. 205–206] that unicorns (or anything else you like) exist: To prove that unicorns exist, it suffices to prove the stronger statement that existing unicorns exist. But for existing unicorns to not exist would be a contradiction; therefore existing unicorns exist; therefore unicorns exist.

[3] For the history of the argument, and a discussion of some of the ontological issues mentioned below, see Barnes

Now, being able to define things into existence like this is metaphysically disturbing, and doesn't really seem possible. Thus, Hume [1779, IX] tried to show that it is not possible that an entity exist of necessity, and Kant took the position described above, which is often characterized as "Existence is not a predicate". This position is now widely accepted in philosophy [Plantinga 1967, p. 38]. Nevertheless, while it may have the merit of keeping God off our backs, it raises difficulties in knowledge representation.

What I want to show in this paper is that existence *can* be predicated, but (lest God be found to be an emergent property of our knowledge representations; no *deus ex machina* here!) it is neither a single predicate nor an ordinary kind of predication.

2.2 Everything exists

An adequate treatment of existence in KR formalisms is complicated not only by the problem described above, but also by a related set of difficulties that derive from a position often summarized by the slogan "Everything exists" (*cf* [Quine 1948, p. 1]). That is, there is nothing that doesn't exist, for if it doesn't exist it isn't anything, and statements apparently about non-existents are either incoherent or can be explained away. The development of this approach is due mainly to Russell [1905, 1918] and, later, Quine [1948]. The Russell–Quine position has become so firmly entrenched in twentieth-century Anglo-American philosophy that it is usually accepted without question [Parsons 1980, p. 1–5]. If we take the slogan literally, then even if existence can be predicated of an entity, it is no more than a tautology; no entities don't exist. And to assert non-existence of something would be self-contradictory. As we will see, this position too is problematic for knowledge representation.

To a large degree, the question seems to be nothing more than what the word *exist* does or should mean, and what status is to be assigned to 'non-existent objects'. Quine grants two kinds of existence: concrete, physical existence in the world (the kind that Margaret Thatcher has), and abstract, non-physical existence (the kind that the number 27 has). "Idea[s] in men's heads" [1948, p. 2] are included in one or the other of these categories, and so too, I assume, are events and actions. Clearly, this is a wider definition of existence than the kind that Anselm and Descartes wished to attribute to God. Presumably they intended some divine equivalent of physical existence—able to have causal interaction with the physical world—and would be unhappy with the idea that God existed only in the way the number 27 does. Likewise, Hume and Kant were using the narrower definition, for many mathematical objects obviously exist of necessity (the number 27; the least prime greater than 27). So perhaps existence in this other sense, non-physical existence without causal connection to the world, *could* be a predicate.

2.3 There are things that don't exist

Quine's sense of the word *exist* may be wider than Anselm's and Descartes's, but it is still much narrower than that of Meinong [1904], who described his position in an oxymoron: "There are objects of which it is true that there are no such objects" [1904, Levi *et al* translation, p. 83]. For Meinong (like Brentano before him), every thought or idea, such as the idea of a gold mountain, must be 'directed toward' some object, and so all objects of thought have being in some sense, even if not real-world existence. Meinong therefore wanted to give status to objects such as the gold mountain, which is not real, and the round square, which is not even possible, arguing that the gold mountain is just as good an object as Mount Everest, and the fact that it is unreal makes no difference. Note that the question is not about the *concept* or *idea* of the gold mountain and whether that exists; clearly, it does. But when we say that the gold mountain is 1000 metres tall, we aren't just talking about an idea; it is not the idea that is 1000 metres tall but the alleged thing that the idea is of.

Russell pointed out that Meinong's approach got into trouble with objects like the gold mountain that exists—which isn't real even though existence is part of its definition (*cf* footnote 2). It also troubled him that there was any sense in which a contradiction like a round square could exist.[4]

Thus the question to be considered is what, exactly, do quantifiers like $\exists$ and $\forall$ quantify over? If an expression begins with '$\forall x$' or '$\exists x$', then what values may be used or considered for x? Do they include Margaret Thatcher, the number 23, World War II, my putting the cat out last night, the late Alan Turing, the possibility of rain tomorrow, suavity, fear, the set of round squares, the concept of round squares, or Sherlock Holmes? In other words, what is in the universe of quantification? What exists?

3 What exists?

The burden on the everything-exists position is to explain the apparent counterexamples—the entities that don't exist and yet seemingly form part of the population of everyday naïve ontology. In the next subsection, I will list some of the problematic examples from natural language, and in section 3.2 I will show how Russell tries to dissolve the problems.

[1972].

[4] Parsons [1980, p. 38–42] has argued that a round square is not a contradiction in the same way a non-square square is, and that the former is a good object but not the latter. Such distinctions need not concern us in this paper.

3.1 What doesn't exist?

Things that aren't there: Perhaps the simplest apparent counterexample (one that we will see Russell's answer to in section 3.2) is that we can explicitly speak of non-existence and non-existent things:

(1) There's no one in the bathroom.

(2) The car I need doesn't exist. [spoken after a long and fruitless search for a suitable car] [Williams 1981, p. 37]

(3) There's no such thing as the bogeyman; he doesn't exist, and neither does Margaret Thatcher.

(4) Nadia doesn't own a dog.

(5) Round squares are impossible, gold mountains merely unlikely.

We may also speak of events that don't occur and actions that are not taken:

(6) A complete lack of money has prevented renovation of the rectory.

(7) The workers threatened to hold a strike if their pay claims were not met. The company acceded to the demands, and the strike was averted.

(8) There are no trains to Saginaw on Sundays. [*i.e.*, the event of a train going to Saginaw on a Sunday does not occur.]

(9) Due to maintenance work on the line, the 6:06 to Saginaw will not run on Sunday.

(10) Today's lecture is cancelled.

(11) Nadia refrained from commenting on Ross's new hairstyle.

(12) Ross failed to notice that Nadia had failed to feed the newt.

Existence itself as an object: We can seemingly speak of existence as an object, one that need not exist:

(13) The existence of Pluto was predicted by mathematics and confirmed by observation.

(14) The existence of Vulcan was predicted by mathematics but disproved by observation.

(15) It's a good thing that carnivorous cows don't exist. [*i.e.*, the non-existence of carnivorous cows is a good thing.]

Claims of reality: We can even (untruly, but not incoherently) assert that unreal objects exist:

(16) I saw a gold mountain near the freeway this morning.

(17) Round squares make me seasick—especially the green ones.

(18) Unreal objects exist.

We can also report such beliefs of others without committing ourselves.

(19) Nadia believes that a unicorn named Old Ironsides has been intercepting her mail and stealing the fashion magazines.

Claims of possibility: We can speak of possible objects and events without committing ourselves either to their reality or unreality, and of objects and events whose existence is merely contingent upon other things.

(20) There may be someone in room 23 who can help you.

(21) If you assemble the parts correctly, you will have created a handsome two-metre model of the CN Tower.

(22) It might rain tomorrow.

Existence at other times: We can refer to things that don't now exist, but did or will. We can speak of things now gone:

(23) Alan Turing was a brilliant mathematician.

(24) Last night's dinner was disastrous.

Sometimes, we may or even *must* even use the present tense for things of the past, suggesting that they have some kind of continuing existence:

(25) *(a)* Alan Turing is a celebrated mathematician. [after Barnes [1972, p. 48]]
(b) *Alan Turing was a celebrated mathematician. [in the sense that he continues to be celebrated]

(26) *(a)* Alan Turing is dead.
(b) *Alan Turing was dead.

And we can talk of things to come:

(27) Tomorrow's dinner is going to be delicious.

(28) The baby that Diane is planning to have will surely interfere with her violin lessons.

Fictional and imaginary characters: We can speak of fictional entities and classes as if they really existed:

(29) Sherlock Holmes was the protagonist of many stories by Conan Doyle.

(30) Sherlock Holmes lived in London with his friend, Dr Watson.

(31) Nadia models herself upon Sherlock Holmes.

(32) Dragons don't have fur. [Plantinga 1967, p. 40]

and possibly even

(33) Sherlock Holmes is no longer alive.

3.2 The Russell–Quine ontology

3.2.1 Paraphrases and the theory of descriptions

Russell's approach, his *theory of descriptions* [1905, 1918], was to regard apparent assertions of existence and non-existence as merely paraphrases—in logic or a literal English rendering thereof—of other forms in which the assertion is not actually made. Instead, the offending bits are expressed as variables and quantifiers, and the resulting expression is something that can legitimately be true or false. Thus, *Dragons exist* is a paraphrase of *There is at least one thing that is a dragon*:

(34) $\exists x(dragon(x))$

Since no such x exists, the sentence is false. Likewise, *Dragons don't exist* is a paraphrase of the negation of (34):

(35) $\forall x(\neg dragon(x))$
'For any x, it is not the case that x is a dragon.'

Attempts to assert properties of non-existent objects may be handled in a similar manner:

(36) Dragons like baklava.
$\forall x(dragon(x) \rightarrow likes\text{-}baklava(x))$

This is vacuously true if there are no dragons [Russell 1918, p. 229]; but statements about particular dragons would be false:

(37) My dragon likes baklava.
$\exists x(my\text{-}dragon(x) \wedge likes\text{-}baklava(x))$

This is false because there is no x for which the left-hand side of the conjunction is true. One might instead have used a vacuously true form like that of (36), but the form of (37) reflects Russell's belief that such sentences were false, and also his concerns with definite descriptions (see below).

In the natural language versions of these statements, we have the apparent problem that even to mention dragons seems to give them some sort of existence; to say that *Dragons like baklava* seems to presuppose the existence of the class of dragons. Russell's claim was that on the 'correct' reading—the representations above, or literal English glosses of them—the problem dissolves. The logical forms contain no assertion of the existence of a non-empty class of dragons. Moreover, the predicate *dragon* is itself a complex term, and may be regarded as simply an abbreviation for a description such as

(38) $fire\text{-}breathing(x) \wedge leather\text{-}winged(x) \wedge \ldots$

Definite references may also be paraphrased. Thus:

(39) The builder of Waverley station was a Scot.

$$\exists x(built(Waverley, x) \wedge \forall y(built(Waverley, y) \rightarrow y = x) \wedge Scot(x))$$

'One and only one entity built Waverley station, and that one was a Scot.' [Russell 1905, p. 113–114]

(If the noun phrase being interpreted does not contain sufficient information to uniquely identify the individual, information from context may be added. Thus (39) might also be the representation of simply *The builder is a Scot* if the context made it clear that the builder was that of Waverley station.) A similar treatment upon *The present king of France is bald* shows the sentence to be false, like (37), because there is no entity denoted by *the present king of France*.[5] Quine [1948, p. 7] showed how the method can be extended to include proper names, so that sentences about named fictional entities might be paraphrased:

(40) Sherlock Holmes is smart.

$\exists x(isHolmes(x) \wedge smart(x))$

'There is an x that has the property of being Sherlock Holmes, and x has the further property of being smart.'

Again, the result is a sentence that is false, for there is no x that has the property of being Sherlock Holmes.

3.2.2 Problems with the theory

Paraphrasing in this manner immediately disposes of some of the problems mentioned in section 3.1, but it does so at some cost.

First, all sentences that assert properties of non-existents are false if talking about a single thing and true if talking about a class, so negating such sentences doesn't change their truth value! For example, the negation of (37) is:

(41) My dragon doesn't like baklava.
$\exists x(my\text{-}dragon(x) \wedge \neg likes\text{-}baklava(x))$

This is false for the same reason that (37) is. Likewise, the negation of (36), *Dragons don't like baklava*, is still

[5] The problem here is, of course, presupposition failure—the sentence tries to talk about something that doesn't exist, and does so without any of the "redeeming" characteristics of the sentences about non-existents that were exhibited in section 3.1. Russell's position on presupposition was famously disputed by Strawson [1950], and is no longer generally accepted. Strawson's position is that the presuppositions of a sentence (or, more precisely, of a particular utterance of a sentence) are distinct from its main assertion, and, unlike the main assertion, are unchanged by sentence negation. If a presupposition is false, then the main assertion, or the sentence itself, can be neither true nor false; rather, it has no truth value at all. For a review of current approaches to presupposition, see [Levinson 1983] or [Horton 1987].

A treatment of presupposition *per se* is beyond the scope of the present paper; for that, see [Horton 1987, Horton and Hirst 1988]. I am concerned here rather with the treatment of the entities that may be felicitously presupposed.

true. The underlying problem here, of course, is that English negation and logical negation aren't the same. If we put a '¬' in front of the logical form of (37), we do change its truth value, but that's not what the English word *not* does. In particular, negation in English (and probably in all natural languages) preserves the presuppositions of the original sentence. In the case of (41), alas, it also preserves Russell's erroneous approach to presuppositions (see footnote 5).

A second problem is a technical one in the nature of the paraphrasing task itself: it destroys, quite deliberately, the similarity between the surface form of the sentence and the representation of its meaning. I have argued elsewhere [Hirst 1987, Hirst 1988a] for the virtues of compositional semantic representations in which each element is a direct reflection of a surface constituent of the sentence. While it is not always possible to maintain this, the advantages to be gained from it are such that it is not to be given up lightly.

Third, and most seriously, there are, as we saw earlier, sentences about non-existents for which one's intuition strongly contradicts the theory of descriptions. These include sentences about defining properties of non-existents and sentences in which non-existents seem to have some actual interaction with the real world.

In the first of these classes, we have sentences such as this:

(42) Dragons have a horn in the middle of their foreheads.

For Russell, this is true, though in any ordinary conversation it would be thought of as false. Likewise, we all agree with Russell and Quine about the falsity of (43):

(43) Sherlock Holmes was stupid.

but we disagree about the reason: in ordinary conversation this sentence is taken as false exactly because (40) is taken as true (*cf* [Parsons 1980, p. 37]).

In the second class are sentences like the following assertions of non-existence themselves. While we might accept representations like (35) for the denial of classes, the denial of the existence of specific entities is trickier. Consider again:

(44) Ross cancelled the lecture.

(45) The [threatened] strike was averted by last-minute negotiations.

On Russell's theory, sentences like these must invariably be false, which is clearly wrong. Notice that paraphrase, in the style of sentence (35), doesn't help here, because these sentences are asserting more than just non-existence; they are asserting a causal relationship. The expression *The strike was averted* means that the strike never occurred—it did not exist—and that some specific action by someone prevented its occurrence. And which strike was averted? The particular strike that the workers threatened to hold, which has specific properties of time, cause, participants, and so on that differentiate it from all other real or potential strikes, all properties that could be used when constructing the description in a Russellian paraphrase. But under Russell's view, we cannot truthfully talk about this strike at all, for it does not exist; any sentence that attempts to refer to it will be false. (Note, as before, that we can't get out of this by saying that the reference is to the idea of the strike; it is not the idea that is averted.)

It might be objected here that to say *The strike was averted* is a looseness of the English language, for one can also use an indefinite reference, saying *A strike was averted*; perhaps this is the basic form that should be interpreted [Barry Richards, personal communication]:

(46) Someone caused that there be no strike.
$\exists y(cause(y, \neg\exists x(strike(x))))$

(We shall allow *cause* as a predicate that takes a proposition in its second argument, and which asserts that the entity in the first argument caused the second argument to be true.) The problem with this tack is the need to say exactly what didn't happen. After all, there are a lot of strikes that were not averted; but (46) says there were no strikes at all. Clearly, some identification from the context is necessary: what was averted was a strike by some particular set of workers at some particular time over some particular claim—so we must identify *the* strike in context, bringing us back to where we started.

Another objection could be that the proper paraphrase is *The strike that was planned was averted*, the claim being that the strike does exist, non-physically, like mathematical objects, by virtue of its having been planned. (This would explain why it sounds a bit funny to say *The accident was averted* instead of *An accident was averted* (*cf* above), as accidents aren't planned.) The problem with this is that one cannot avert mathematical objects any more than one can avert ideas. Perhaps what was averted was the physical realization of this non-physical object—in effect, the instantiation of a concept. I will pursue this line in section 4.2 below.

One could also claim that if the strike was planned, it exists as a 'future object'. To examine this, we must consider the role of time. Unfortunately, Russell provides no treatment of existence at times other than the present, but we can speculate on how he would extend his theory to do so.

Let's consider the simpler case first: the past. It is unclear from Russell's account how he would paraphrase, say, *Alan Turing was smart* and *Alan Turing is dead.* That is, would he allow the scope of quantification to include past entities? Doing so would let the first of these sentences be paraphrased like any other,

and the past-tense verb would just be an artifact of the past-ness of Alan Turing himself, not included in the paraphrase:

(47) Alan Turing was smart.
$\exists x(isTuring(x) \wedge smart(x))$

This would then be a true sentence, unlike *Sherlock Holmes was smart.* But trying this for the second sentence:

(48) Alan Turing is dead.
$\exists x(isTuring(x) \wedge dead(x))$

doesn't work, because Turing wasn't dead when he existed, and the verb tense hasn't behaved. At a minimum, we need to add some notion of time points or intervals such that propositions can be true at some times and not others; thus, (48) would be true today, but false in 1945 and 1862—false in 1945 because Turing was still alive, and false in 1862 because he hadn't yet come within the scope of the existential quantifier.

Thus the universe is seen as travelling through time, collecting up entities into its ontology as it proceeds. Once a thing has started to exist, it never stops. This helps represent sentences (47) and (48), but I don't think this view can be pleasing for the everything-exists gang, for the fact remains that Alan Turing does not now exist in the world any more so than the gold mountain does, nor does he seem to exist as a mathematical object. (The idea of Turing continues to exist, but it's not that that's dead.) There doesn't seem to be any good reason why his brief time on earth should give Turing any subsequent ontological advantage over the gold mountain.[6]

These problems may be seen even more clearly if we now consider future entities, such as the strike that the faculty are threatening to hold. We can talk about this just as easily as we can about Alan Turing (albeit with less certainty)—it will be long and nasty, it will cause the university president to resign, it may never happen (!). For Quine, certainly (and presumably for Russell—guilt by association), the strike is merely a 'possible object', to be kept out of one's ontology at all costs (*cf* his arguments against the existence of the 'possible man in the doorway' [Quine 1948]). So now the averted strike is out on two separate counts, each fatal on its own. When it was still a planned strike, it was merely a possible object; after it was averted, it became a past object as well.

But for KR and NLU, this is simply not acceptable. I have shown above that objects like Alan Turing and the averted strike must be able to be represented, quantified over, and reasoned about just as much as Margaret Thatcher. So the Russell–Quine position is inadequate, and we must look for alternatives. This I will do in sections 5 and 6, after first examining the degree to which KR formalisms share the Russell–Quine deficiencies.

4 Existence assumptions in KR formalisms

To what extent are knowledge representation formalisms able to deal adequately with existence and non-existence? The universe of discourse of a system is, of course, circumscribed by what's in its knowledge base; but given that non-existent entities may have to be included (and, in a full NLU system, *must* be included), how does the average formalism behave?

[6] A rejoinder that I shall not take very seriously: Alan Turing does in fact still exist, or at least his soul does, in Heaven or Hell or somewhere like that. On this view, one might say that the best paraphrase for *Alan Turing is dead* is one of these:

(i) Alan Turing's body doesn't exist (or no longer exists).
$\neg\exists x(bodyOfTuring(x))$

(ii) Alan Turing is in the afterlife.
$\exists x \exists y(isTuring(x) \wedge afterlife(y) \wedge in(x, y))$

Form (i) is undoubtedly true, and the truth of form (ii) depends on whether there is an afterlife and if so who's there (issues that I will not solve in this paper).

The value of this particular objection is to draw attention to the cultural bias in the expression of the problem; perhaps we say that Alan Turing *is* dead just because English reflects our long cultural history of belief in a soul and an afterlife. If we are careful to avoid such bias in our language, we will be able to analyze the problem correctly (or so said a large twentieth-century school of philosophy). Notice, for example, that English offers no analogous expressions for the past existence of objects to which we do not (culturally) attribute an afterlife; if my wristwatch has ceased to be, I can say *My wristwatch was destroyed* but not *My wristwatch is destroyed* (and only as a joke or metaphor, *My wristwatch is dead*). Thus when we say that Turing is dead, our paraphrase should be no more than that there is no x such that $isTuring(x)$; and that this statement was false at an earlier time is an implicature of the word *dead.*

I don't think that this argument goes through. There are too many other things we can say about entities of the past that seem to presume their continued existence:

(iii) Alan Turing {is | *was} a celebrated mathematician.

(iv) Nadia models herself upon Alan Turing.

(v) Nadia knows more about NP-completeness than Alan Turing ever did. [Although Turing is referred to in the past tense, the entity *Alan Turing's knowledge of NP-completeness* is available for comparison with an entity, Nadia's knowledge, that exists in the present and did not exist at the time of Turing.]

(vi) Nadia modelled her new sculpture upon my old wristwatch (which was destroyed last year).

(vii) The Flat Earth Society is now disbanded.

For the most part, KR formalisms are Russellian in their approach to ontology. It is a general characteristic of KR formalisms that even when they are declarative they are assertional—that is, to state something is to assert its truth; one cannot say that something is false. One can, of course, assert the negation of falsehoods on those occasions when this yields truth, but there is no concept of the truth value of a statement being independent of the expression of the statement. Likewise, it is a usual assumption in formalisms that to use a term is to assert that it denotes, and, in particular, that it denotes an extant entity. To assert, for example, *cancelled*(*lecture23*, *Ross*), implies for most systems (*e.g.*, KRYPTON [Brachman *et al* 1983] and conceptual graphs [Sowa 1984]) that *lecture23* exists just as much as *Ross* does, even if the expression says that it doesn't.

4.1 Hobbs: Ontological promiscuity

Not all KR formalisms impute existence to denotations of their terms. A simple first-order system in which (ignoring all the philosophical wisdom discussed above) *exists* is a predicate like any other has been proposed by Hobbs [1985] in his paper entitled "Ontological promiscuity". The 'promiscuity' of the title refers not to the Meinong-like inclusion of all non-existent objects, but rather to reification of events and properties as objects;[7] Hobbs's set of objects, over which quantifiers range and in which all variables are assumed to denote, is a Platonic universe, "highly constrained by the way the ... material world is" (p. 63). The formalism is deliberately simple and 'flat', without modals, intensions, or even negation.

In this approach, no object mentioned in a representation is assumed to exist in the real world unless such existence is either explicitly stated or axiomatically derivable. For example, *Ross worships Zeus* is represented as:

(49) $Exist(E) \wedge worship'(E, Ross, Zeus)$

This says that E is a worshipping by Ross of Zeus, and E exists. The predicate *worship'* is transparent in its second argument but not its third. This means that the existence of E implies the existence of Ross, but not that of Zeus. Hobbs shows that with an adaptation of Zalta's system of abstract objects [Zalta 1983], this approach is able to deal with several problems of opaque contexts that are usually thought to require higher-order representations, while at the same time remaining (moderately) faithful to the surface form of the English sentence.

[7] Treating events as objects, in the style of Davidson [1980], is a position that I have adopted in this paper and assumed to be relatively uncontroversial even for supporters of the Quine–Russell position. Treating properties as objects is a separate question somewhat orthogonal to the concerns of the present paper; suffice it to say here that Quine and Russell would not, I think, approve.

Although Hobbs mentions non-existence only briefly, it is clear that by extending his approach we can account for some of the problems mentioned above. Just as transparent argument positions entail existence, we will allow an argument position to be *anti-transparent*, entailing that the object in that position does *not* exist.[8] (Anti-transparent positions are not to be confused with Hobbs's opaque positions, which entail nothing.) We can then represent the prevention of the occurrence of the strike:

(50) The strike was averted.
$strike(s) \wedge \exists x(Exist(E) \wedge avert'(E, x, s))$

It would be stipulated that *avert'* is transparent in its second argument and anti-transparent in its third—that is, the existence of E implies the non-existence of s.

The existence of existence also seems representable. Hobbs has a 'nominalization operator', ', which turns an n-ary predicate into an $(n + 1)$-ary predicate whose first argument is the condition that holds when the base predicate is true of the other arguments. We saw this above with the ternary predicate *worship'*$(E, Ross, Zeus)$, derived from the binary predicate *worship*$(Ross, Zeus)$. Since *Exist* is just another predicate, there is nothing to stop us nominalizing it:

(51) The existence of carnivorous cows is predicted by GB theory.

$Exist'(E_1, carnivorous\text{-}cows)$
$\wedge predict'(E_2, GB\text{-}theory, E_1)$
$\wedge Exist(E_2)$

'E_1 is the existence of carnivorous cows, E_2 is the prediction of E_1 by GB theory, and E_2 exists (but E_1 might not).'

On the other hand, there is no treatment of fictional objects. Non-existent objects can be mentioned, as we saw in the assertion of *Ross worships Zeus*, but there is nothing that lets us say that Zeus exists in fiction whereas the Giant Cosmic Groundhog (which I just made up) and the averted strike do not. An obvious move is simply to add a predicate *Fictional* to the formalism. Then *worship'* would have the property that its third argument must exist either in the real world (like Nadia, whom Ross also worships) or in fiction (even if only a small fiction in Ross's mind). Hobbs's Platonic universe would now have a tripartite division into the existent, the fictional, and all the rest.[9]

[8] Hobbs explicitly excludes negation from his formalism, but I shall assume it to be added in the usual way.

[9] I will resist the temptation to be side-tracked onto the question of characterizing more precisely what it means to be fictional. However, there is no principled reason I can see for limiting the property to the imaginary entities in published or oral literature; those in lies and untrue thoughts are just as fictional as Sherlock Holmes ever

But so far, this approach doesn't give an adequate treatment of objects like Alan Turing—for simplicity, Hobbs did not include any notion of time in his formalism, so we can't talk about Turing's different statuses at different times. In addition, it seems that Anselm's fallacy is valid in the system. Although Hobbs gives no examples of definitions, it would seem that *Exist* could be used directly or indirectly as a defining characteristic, it being just another predicate. Its direct use in a definition could be prohibited by stipulation; but preventing its indirect use is not possible, as it is a deliberate feature of the system that existence can be axiomatically derived from various assertions—one has to be allowed to define predicates with transparent arguments. Thus, following Descartes's version of the fallacy, one could define the predicate *perfect* to be transparent in its (sole) argument, and then assert that God is, by definition, *perfect*.

4.2 Intensional approaches

Although it was important for Meinong that thoughts and ideas could be directed to non-existent objects, I have said little up to now, except in passing, about ideas, intensions, and concepts. Indeed, both Russell and Hobbs were at pains to avoid the standard Fregean distinction [Frege 1892] between intension and extension (*Sinn* and *Bedeutung*). But even Quine grants ideas a place in his universe (see section 2.2 above); so we now turn to this topic. I will use the terms *concept*, *idea*, and *intension* interchangeably below; the technical differences between them will be unimportant. Likewise, I will conflate *extension* with the *denotation*, *realization*, or *instance* of an idea.

An adequate treatment of concepts as 'first-class objects' has often eluded knowledge representation systems. By a first-class object, I mean here an object that can be referred to as an individual in its own right, used in inference, be a component of other objects, and so on. This would be necessary if we were to act on the suggestion (section 3.2.2 above) that sentence (45) be represented as the prevention of the realization of an instance of the concept of strikes. Now, because concepts are used to define other objects, many systems accord them a special status that precludes their simultaneously acting as ordinary objects or individuals. A typical example is Frail [Charniak *et al* 1983], a language in which concepts are *generic frames*, but inference can be carried out only on *instances* of those frames; it is not possible for a frame to be simultaneously generic and an instance. In Krypton [Brachman *et al* 1983], which makes a careful separation of 'terminological' knowledge (which goes in its 'T-box') and assertions about the world (in its 'A-box'), it is possible to reason with the terminological knowledge, which can be thought of as statements about concepts, but concepts *per se* can still not be reified as first-class individuals.

Languages in which concepts *are* first-class objects include McCarthy's first-order language [McCarthy 1977, McCarthy 1979], Shapiro and colleagues' SNePS [Maida and Shapiro 1982, Shapiro and Rapaport 1987], and Sowa's Conceptual Graphs [Sowa 1984]. Such languages must provide a mechanism to relate objects to the concepts of which they are instances. For example, Sowa's conceptual graphs tie concepts and their extensions together by notational means. Thus `[CAT:*]` represents the concept of cats, and `[CAT:#234]` represents some particular cat (namely, cat number 234). The notation `[CAT:*x]` represents the individual concept of a cat: a single cat, but not any particular known one; the `x` may be thought of as a variable, so that all occurrences of `[CAT:*x]` must refer to the same (unknown) cat, but `[CAT:*y]` may be a different one. The different types may be used interchangeably (with different meaning, of course) in the graph representations that can be built. However, all graphs are implicitly existentially quantified; that is, the ontology is implicitly Russellian.

Likewise, McCarthy's language has both concepts and extensions as entities, not formally distinguished from one another.[10] A function called *denot* maps concepts to the entities, if any, that they denote. (Thus individual concepts such as *John* are mapped to an individual, and presumably generic concepts like *Dog*, not explicitly mentioned by McCarthy, would be mapped to an appropriate set of individuals.) A predicate *Exists* is true of those concepts for which there is a denotation.[11] 'Parallel' predicates may be defined for denotations and concepts. For example, if *ishorse* is a predicate true of horses, then *Ishorse* can be defined as a predicate true of concepts for which *ishorse* is true of their denotations, and possibly also true of some concepts that don't have denotations, such as *Pegasus*.

The SNePS network formalism is of special interest, as Rapaport [1985b] has suggested that Parsons's theory (section 6 below) could give it a formal semantics. In SNePS, *all* entities are intensions, and extensions *per se* are not used. This is because SNePS takes representations to be those of the knowledge of an

was, even if not as widely known. In the strictest interpretation, this means that just mentioning examples of non-existent entities, such as the Giant Cosmic Groundhog above, gives them status as fictional—which is not really what one wants, as then *all* non-existent entities would be fictional.

[10] The typographical distinctions in McCarthy's formulas are for the reader's convenience, and are not part of the theory.

[11] McCarthy's *Exists* is not to be confused with Hobbs's predicate of the same name (section 4.1 above). McCarthy's *Exists* is a predicate true of concepts that have real-world denotations; Hobbs's *Exists* is true of the real-world objects themselves.

agent, rather than of the world directly. The intensions are connected to reality only through the agent's perception.[12] Like McCarthy, Shapiro and colleagues show only individual concepts, such as the node *John* representing the idea of John;[13] I assume that if the agent is to think about the idea of John, it will need a node that represents the idea of the idea, with a *denot*-like arc relating them.

It is interesting to note that, generally speaking, KR formalisms that treat concepts as first-class objects do not formally distinguish them from individuals. (Those that don't do—they have to, in order to discriminate against them.) I don't know of any principled reason for this. Such systems are *weakly intensional* systems, countenancing intensions but not making anything special of them. In contrast, *strongly intensional* systems take intensions to be not just first-class objects but objects of a distinct kind.[14] Montague semantics [Montague 1973] is a good (non-computational) example of a strongly intensional system.

I suspect that a strongly intensional system will be necessary for an ontologically adequate treatment of intensions. McCarthy could use his *denot* function to map intensions to their extensions, but going in the opposite direction requires an operator, as in Montague semantics. The examples of section 3.1 show such operations to be frequently necessary, and the modes of existence to be discussed in section 7 below suggest that a diverse set of operators may be required.

5 Free logics

One solution that has been suggested to the problems of the Russell–Quine approach is the use of *free logics*. A free logic is a logic that makes no assumptions about existence—specifically, a logic that tolerates terms that have no denotation in its universe but never quantifies over such terms.[15] For example, Woodruff's system UE [Woodruff 1970] is a free logic with truth-value gaps (*i.e.*, the truth values t, f, and u) and a distinction between assertions of truth and assertions of non-falsity. Non-denoting terms have no interpretation at all, and a predicate need only have truth value t or f if all its arguments denote. Thus the system is explicitly Strawsonian. In contrast, Schock's free logic [Schock 1968] has only two truth values, and (in the style of Frege) uses the empty set as the 'denotation' of non-denoting terms. Both systems have an 'existence' predicate, which is true just of those terms that denote.

Free logics seem to be an attractive solution in KR to the problems of Russellianism. From an NLU perspective, free logics help avoid Russellian paraphrases, thereby leading to a more compositional semantics. From a KR viewpoint, they are a conceptually easy extension of classical systems; deduction systems already exist for them; and truth-value gaps are already a focus of research in the field (*e.g.*, Patel-Schneider's four-valued logic [Patel-Schneider 1986]). But alas, free logics turn out to have most of the same problems for NLU as Russell's standard logic. Sentences about non-existents need not be false (at least in Woodruff's logic), but (except in a trivial, unhelpful way) they still can't be true.

6 Parsons: Non-existent objects

Hobbs's scheme implicitly countenanced nonexistent objects, but, as we saw, found itself limited because it tried not to make anything special of the notion of existence. Free logics also accept non-existent objects, but try their best to ignore them. We now turn to an approach that doesn't just accept such objects, but whole-heartedly embraces them. The approach is that of Parsons [1980]; it is explicitly motivated by Meinong's ideas (see section 2.3 above).[16]

Parsons defines the set of *nuclear properties* as the set of properties such as being green, being in New Zealand, or being Nadia. Such properties are "ordinary properties" [Parsons 1980, p. 24] that we regularly attribute to individuals, and corresponding to each is a *nuclear predicate* true of individuals that have that property. Nuclear predicates are in contrast to *extra-nuclear predicates*, of which the prime example is *Exists*. Thus, existence is taken as a predicate, but one of a special kind.

Parsons's universe contains only, for any set of nuclear properties, the unique object that has exactly that set of properties. There is an object that is green (and has no other property but that); there is an object that is both green and Nadia; there is even an object

[12] Thus SNePS is free of extensions only for an external observer of the system. The SNePS objects used by a computational agent that employs the formalism (such as Rapaport's CASSIE [Shapiro and Rapaport 1987]) are the concepts in that agent's 'mind', so to the observer they are intensions. To the agent itself, however, they are subjective extensions, identified with its perceptions of reality.

[13] For simplicity, I will ignore Shapiro's careful distinction between nodes and their names.

[14] Both the distinction and the terminology are due to Graeme Ritchie [personal communication].

[15] Hobbs's system (section 4.1 above) is not a free logic. While it makes no assumptions about real-world existence, it does assume that all terms denote something in the Platonic universe, and it quantifies over them.

[16] Rapaport [1981, 1985a] has also presented a Meinong-inspired theory of non-existent objects. Space does not permit discussion of both theories. The main differences between the two are that *(a)* Parsons has two types of predicate, whereas Rapaport has one type that can be applied in two different ways; and *(b)* Parsons has only one type of object, which may or may not exist, whereas Rapaport distinguishes Meinongian objects ('M-objects') from actual objects ('*sein*-correlates' of M-objects).

that is green and Nadia and Margaret Thatcher. But not all these objects exist in the real world—in some cases because they just happen not to, and in other cases because they are not possible. Being possible is another extra-nuclear predicate.

The tricky part is what to do with non-existent objects like the existent golden mountain. It's an object with the properties of goldenness, mountainhood, and existence, but it's not included in the universe as defined above because existence is not a nuclear property. Nonetheless, it must be accounted for, as we can still talk about it, and the account must not entail its existence. So following Meinong, Parsons introduces the concept of *watering down* extra-nuclear properties to a nuclear ones. Thus for Parsons, there is also a nuclear existence property, call it $Exist_n$, and that's what the existent golden mountain has. Watered-down existence says nothing about real, genuine, full-blown extra-nuclear existence, and the existent golden mountain still doesn't have the latter. A similar story can be told about the possible round square; its possibility is merely the watered-down variety.

The watering-down operation on an extra-nuclear predicate creates a nuclear predicate true of a superset of the objects of which the original predicate was true. That is, if a given object has an extra-nuclear predicate true of it, it will have the corresponding watered-down nuclear predicate true of it as well (but not necessarily vice versa). Anything that exists full-strength also exists in a watered-down way; anything that is full-strength-possible is also watered-down-possible. But it's not clear exactly what sort of a thing these watered-down properties can be if they don't really do anything. What exactly is it that the watered-down-existent gold mountain has that the regular gold mountain doesn't? Just, it seems, an abstract attribution that has no effect on anything except in serving to distinguish the two.

Parsons develops a formal language, called $\mathcal{O}$, for talking about this universe. $\mathcal{O}$ is a second-order modal language with belief contexts; quantification is explicitly over all objects in the universe. The language distinguishes the two types of predicates, and the extra-nuclear predicate of existence, denoted $E!$, has special axiomatic properties. The watering-down operation on extra-nuclear predicates is defined. Using Montague-like techniques [Montague 1973], Parsons shows how $\mathcal{O}$ can act as a semantics for a fragment of English, treating sentences such as

(52) The King of France doesn't exist.
$\neg(\iota x)(E!(x) \wedge \mathit{King\text{-}of\text{-}France}(x))[\lambda y E!(y)]$

Roughly, this says that it is not true that there is—in the actual world—a unique x that both is the King of France and exists in the world; if there is indeed no King of France, this formula is true. Also included in the fragment is the sentence *Every good modern chemist knows more about chemical analysis than Sherlock Holmes* (*cf* sentence (v) of footnote 6).

If we are willing to accept Parsons's approach, then a number of our problems are solved. We can talk about Sherlock Holmes and dragons and other fictional objects all we like (Parsons devotes two chapters to fictional objects). We also have Alan Turing available, and, presumably, all future objects. And we have all objects that don't exist, including the strike that was averted and the lecture that was cancelled—that is, we have the objects that have exactly the properties required, with no necessity that they exist. And the existence of God is expressible in $\mathcal{O}$, but is not a theorem.

Parsons's approach is not without problems. (See [Rapaport 1985c] for a detailed critique.) The most obvious, especially for a computational implementation, is the profligate scope of the quantifiers. A free-logic insight that must be retained is that quantification scope must be restrained. Parsons's universe is much too large to quantify over (although Meinong would do so). But there is no single correct constraint on quantification. For example, it would normally be silly to quantify over all the unwritten books, unthought ideas, or unlived lives; but sometimes, one might have to do so. (An unwritten book is surely reified in the sentence *Ross is going to start writing a book.*) In KR systems, this may not be a practical problem, for the size of the universe is limited by the size of the knowledge base anyway, and even within that, searches would normally be further constrained. This is not to say that a knowledge base cannot contain (finite representations of) infinite objects—the set of integers, for example—but a practical system will normally limit itself to the entities it already knows about and won't capriciously start generating new ones just to see what turns up.

Another problem is that while the averted strike and cancelled lecture are available as objects, we can't do everything with them that we would like. $\mathcal{O}$ can say that an existent Ross stands in a *cancelled* relation to a non-existent lecture, but it is not possible, I think, to explicate the meaning of this as Ross *causing* the non-existence; Parsons did not consider such things.

7 Naïve ontology: The ontology of natural language

The real problem with the Russell–Quine position, the free logic approach, and even Parsons's approach is that they equivocate about existence; they speak as if all things that exist exist in the same way. This is clearly not so. Margaret Thatcher exists, and so does the number 27, but they do so in different ways: one is a physical object in the world, while the other has only abstract existence. But even Quine is willing to grant the existence of mathematical entities—and of concepts in general. If we admit these two kinds of

existence, then perhaps we can find even more kinds if we look. And arguments about the nature of one kind need not hold true of the others.

In fact, we can identify about nine different kinds of existence. In doing so, we will follow Meinong, Parsons, and Rapaport in not limiting existence to things in the world, but attributing it to anything that can be spoken of.[17] In this view, everything exists, but not as Quine meant that slogan. In particular, we solve the problems of sentences (44) and (45) by attributing existence (but not physical actuality) to the lecture and the strike concerned. Thus all terms will denote, and all sentences will be about existent objects, in some sense of existence, and will have the potential to be true.

The various kinds of existence are as follows:[18]

- Physical existence in the present real world, with causal interaction. Margaret Thatcher exists this way.
- Physical existence in a past world (with causal interaction therein, and some indirect causal connection to the present world). The late Alan Turing, for example, exists in a world of the past; he doesn't exist now, but nevertheless he *is*, in the present, a celebrated mathematician, and likewise he *is* dead (see section3.2.2 above).
- Abstract, necessary existence, as of mathematical objects such as 27 and the least prime greater than 27.
- A sort of doubly abstract existence, granted to objects such as $\sqrt{-1}$; mathematicians distinguish such 'imaginary' numbers from other numbers for good reason.
- Existence outside a world, but with causal interaction with that world. In most Western religions, this kind of existence is attributed to God; that is, God is not thought to exist merely the way the number 27 does.
- Abstract, contingent existence in the real world. Freedom, suavity, and fear would come into this category.
- Existence as a concept, which is abstract but contingent, such as the concept of Margaret Thatcher, which need not have existed.[19]
- Unactualized existence,[20] as of the baby that Diane wants to have after she graduates, the strike that the faculty would have held if they hadn't got a pay rise, or Margaret Thatcher's brushing her teeth the day after tomorrow. This category includes objects that could become actual in the future, objects in counterfactuals, 'past' objects that never came into being, and perhaps also impossible objects. Again: objects with this sort of existence are distinct from concepts—it is not just the concept of the faculty strike that was averted, it was the strike itself.
- Existence in fiction. This is the sense in which Sherlock Holmes and dragons exist.[21]

My point here is not to argue for exactly this list of types of existence—that is a topic in philosophy, not knowledge representation—but to demonstrate that however many distinct types of existence there are, it's rather more than two. Any KR formalism that is to be adequate to the task of NLU will need to be able to account for them all—that is, it will treat existence as a set of properties, and, given a particular object's mode of existence, draw inferences accordingly.

8 Conclusion

Before closing, a word should be said about possible worlds. It might be suggested that objects such as the averted strike, the gold mountain, or Nadia's baklava-loving dragon *do* have physical existence, just like Margaret Thatcher, but have it in a different possible world. This misses the point. What we want to talk about and represent is one particular world, usually the actual world, and the question is therefore how dragons and averted strikes exist in the particular world of interest. It is insufficient to say merely

[17] The view presented here goes beyond these authors in that it imposes a taxonomy of existence upon their basic ontology.

[18] It should be clear that these kinds of existence can't all be accounted for just by organizing the IS-A hierarchy the right way. It is true that one can, at the top, make a distinction between abstract and concrete entities. But past existence and unactualized existence are certainly orthogonal to the hierarchy of concrete entities.

[19] One may wish to combine this category with the previous one, saying that concepts are not ontologically distinct from other abstract entities like suavity. I will not take a position on this. Alternatively, one might argue that the existence of a concept may be necessary or contingent depending on its extension. That is, the concept of Margaret Thatcher is as contingent as Margaret Thatcher is, but the concept of the least prime greater than 27 is necessary because its extension is. This category of existence would then be split over the three abstract categories above.

[20] I use this horrible term for want of a better one.

[21] "Everyone knows that dragons don't exist. But while this simplistic formulation may satisfy the layman, it does not suffice for the scientific mind ... The brilliant Cerebron, attacking the problem analytically, discovered three distinct kinds of dragon: the mythical, the chimerical, and the purely hypothetical. They were all, one might say, non-existent, but each non-existed in an entirely different way." (Stanislaw Lem. "The third sally, or The dragons of probability." *The cyberiad: Fables for the cybernetic age* (Michael Kandel, translator). NY: Avon books, 1976, p. 76.)

that the status of dragons is that they exist in a different possible world, for so, after all, does Margaret Thatcher. That tells us nothing about the difference between dragons and Margaret Thatcher in the world that we *are* representing.

KR formalisms that are to be suitable for NLU must take account of the many different modes of existence that can be spoken of in natural language. Traditional Russellian approaches are inadequate, as are free logics. However, by taking existence to have a variety of modes, and treating it as a property of objects, an adequate approach can be developed.

Acknowledgements

I am grateful to Chrysanne DiMarco, Diane Horton, Andrew Malton, Chris Mellish, Bill Rapaport, Barry Richards, Graeme Ritchie, Stephen Regoczei, Stuart Shapiro, John Sowa, and Nadia Talent for discussions on these matters and comments on earlier drafts of this paper.

References

[Anselm 1078] Anselm. *Proslogion* (M. J. Charlesworth, translator). Oxford: Clarendon Press, 1965.

[Barnes 1972] Jonathan Barnes. *The ontological argument* (New studies in the philosophy of religion). London: Macmillan, 1972.

[Brachman and Levesque 1985] Ronald Jay Brachman and Hector Joseph Levesque (editors). *Readings in knowledge representation.* Los Altos, CA: Morgan Kaufmann, 1985.

[Brachman *et al* 1983] Ronald Jay Brachman, Richard E. Fikes, and Hector Joseph Levesque. "KRYPTON: A functional approach to knowledge representation." Technical report 16, Fairchild Laboratory for Artificial Intelligence Research, May 1983. Reprinted in [Brachman and Levesque 1985], 411–429. An earlier version appeared in *IEEE Computer*, 16(10), October 1983, 67–73.

[Charniak *et al* 1983] Eugene Charniak, Michael Gavin, and James Hendler. "The Frail/NASL reference manual." Technical report CS-83-06, Department of Computer Science, Brown University, February 1983.

[Davidson 1980] Donald Davidson. *Essays on actions and events.* Oxford: Clarendon Press, 1980.

[Descartes 1641] René Descartes. *Meditations on first philosophy.* In: *The philosophical works of Descartes* (Elizabeth S. Haldane and G. R. T. Ross, translators), volume I. Cambridge University Press, 1911.

[Feigl and Sellars 1949] Herbert Feigl and Wilfrid Sellars (editors). *Readings in philosophical analysis.* New York: Appleton-Century-Croft, 1949.

[Frege 1892] Gottlob Frege. "Über Sinn und Bedeutung." *Zeitschrift für Philosophie und philosophische Kritik*, 100, 25–50. Reprinted as: "On sense and nominatum" (Herbert Feigl, translator), in [Feigl and Sellars 1949], 85–102. Reprinted as: "On sense and reference" (Max Black, translator), in Peter Thomas Geach and Max Black (editors), *Translations from the philosophical writings of Gottlob Frege*, Oxford: Blackwell, 1952, 56–78.

[Hirst 1987] Graeme Hirst. *Semantic interpretation and the resolution of ambiguity* (Studies in natural language processing). Cambridge University Press.

[Hirst 1988a] Graeme Hirst. "Semantic interpretation and ambiguity." *Artificial intelligence*, 34(2), March 1988, 131–177.

[Hirst 1988b] Graeme Hirst. "Knowledge representation problems for natural language understanding." In: Harald Trost (editor), *4. Österreichische Artificial-Intelligence-Tagung: Wiener Workshop Wissensbasierte Sprachverarbeitung: Proceedings* (Informatik-Fachberichte 176), Berlin: Springer-Verlag, August 1988.

[Hobbs 1985] Jerry Robert Hobbs. "Ontological promiscuity." *Proceedings, 23rd annual meeting of the Association for Computational Linguistics*, Chicago, June 1985, 61–69.

[Horton 1987] Diane Lynn Horton. *Incorporating agents' beliefs in a model of presupposition.* MSc thesis, Department of Computer Science, University of Toronto; published as technical report CSRI-201, Computer Systems Research Institute, University of Toronto, August 1987.

[Horton and Hirst 1988] Diane Lynn Horton and Graeme Hirst. "Presuppositions as beliefs." *Proceedings, International conference on computational linguistics (COLING-88)*, Budapest, August 1988, 255–260.

[Hume 1779] David Hume. *Dialogues concerning natural religion* (Norman Kemp Smith, editor). Oxford: Clarendon Press, 1935.

[Kant 1787] Immanuel Kant. *Critique of pure reason* (second edition) (Norman Kemp Smith, translator). Revised edition, London: Macmillan, 1933.

[Lambert 1983] Karel Lambert. *Meinong and the principle of independence* (Modern European philosophy). Cambridge University Press, 1983.

[Levinson 1983] Stephen C. Levinson. *Pragmatics* (Cambridge textbooks in linguistics). Cambridge University Press, 1983.

[Maida and Shapiro 1982] Anthony S. Maida and Stuart Charles Shapiro. "Intensional concepts in propositional semantic networks." *Cognitive science*, 6(4), October–December 1982, 291–330.

[McCarthy 1977] John McCarthy. "Epistemological problems of artificial intelligence." *Proceedings, 5th International Joint Conference on Artificial Intelligence*, Cambridge, Massachusetts, August 1977, 1038–1044.

[McCarthy 1979] John McCarthy. "First order theories of individual concepts and propositions." In: Jean Elizabeth Hayes, Donald Michie, and L. I. Mikulich (editors). *Machine intelligence 9.* Chichester: Ellis Horwood, 1979, 129–147. Reprinted in [Brachman and Levesque 1985], 523–533.

[Meinong 1904] Alexius Meinong. "Über Gegenstandstheorie." In: Alexius Meinong (editor). *Untersuchungen zur Gegenstandstheorie und Psychologie*, Leipzig: Barth, 1904. Reprinted in: Alexius Meinong. *Gesamtausgabe* (Rudolf Haller, and R. Kindinger, editors), Graz: Akademische Druck- und Verlaganstalt, 1969–1978, Volume II, 481–535. In English as: "The theory of objects" (Isaac Levi; D. B. Terrell; and Roderick Milton Chisholm, translators). In: Roderick Milton Chisholm (editor). *Realism and the background of phenomenology* (The library of philosophical movements), Glencoe, IL: The Free Press, 1960, 76–117.

[Montague 1973] Richard Montague. "The proper treatment of quantification in ordinary English." In: Kaarlo Jaakko Juhani Hintikka, Julius Matthew Emil Moravcsik, and Patrick Colonel Suppes (editors). *Approaches to natural language: Proceedings of the 1970 Stanford workshop on grammar and semantics.* Dordrecht: D. Reidel, 1973, 221–242. Reprinted in: Richard Montague. *Formal philosophy: Selected papers of Richard Montague* (Richmond Hunt Thomason, editor). New Haven: Yale University Press, 1974, 247–270.

[Parsons 1980] Terence Parsons. *Nonexistent objects.* New Haven: Yale University Press, 1980.

[Patel-Schneider 1986] Peter Frederick Patel-Schneider. "A four-valued semantics for frame-based description languages." *Proceedings, Fifth national conference on artificial intelligence (AAAI-86)*, Philadelphia, August 1986, 344–348.

[Plantinga 1967] Alvin Plantinga. *God and other minds: A study of the rational justification of belief in God.* Ithaca: Cornell University Press, 1967.

[Quine 1948] Willard Van Orman Quine. "On what there is." *Review of metaphysics*, 1, 21–38. Reprinted in: *Proceedings of the Aristotelian Society*, supplementary volume 25: *Freedom, language, and reality*, 1951, 216–233. Reprinted in: Willard Van Orman Quine. *From a logical point of view: Logico-philosophical essays.* Harvard University Press, first edition, 1953; second edition, 1961; second edition, revised, 1980, 1–19.

[Rapaport 1981] William J. Rapaport. "How to make the world fit our language: An essay in Meinongian semantics." *Grazer philosophische Studien*, 14, 1981, 1–21.

[Rapaport 1985a] William J. Rapaport. "Nonexistent objects and epistemological ontology." *Grazer philosophische Studien*, 25/26, 1985/86, 61–95.

[Rapaport 1985b] William J. Rapaport. "Meinongian semantics for propositional semantic networks." *Proceedings, 23rd annual meeting of the Association for Computational Linguistics*, Chicago, June 1985, 43–48.

[Rapaport 1985c] William J. Rapaport. "To be and not to be" [Review of Parsons 1980]. *Noûs*, 19(2), June 1985, 255–271.

[Routley 1980] Richard Routley. *Exploring Meinong's jungle and beyond: An investigation of noneism and the theory of items* (interim edition). Departmental monograph 3, Philosophy Department, Research School of Social Sciences, Australian National University, 1980.

[Russell 1905] Bertrand Russell. "On denoting." *Mind*, n.s. 14, October 1905, 479–493. Reprinted in: Bertrand Russell. *Essays in analysis* (Douglas Lackey, editor), London: George Allen & Unwin, 1973, 103–119. Reprinted in [Feigl and Sellars 1949], 103–115. Reprinted in [Russell 1956], 39–56.

[Russell 1918] Bertrand Russell. "The philosophy of logical atomism." In [Russell 1956], 175–281.

[Russell 1956] Bertrand Russell. *Logic and analysis: Essays 1901–1950* (Robert Charles Marsh, editor), London: George Allen & Unwin, 1956.

[Schock 1968] Rolf Schock. *Logics without existence assumptions.* Stockholm: Almqvist & Wiksell, 1968.

[Shapiro and Rapaport 1987] Stuart Charles Shapiro and William J. Rapaport. "SNePS considered as a fully intensional propositional semantic network." In: Nick Cercone and Gordon McCalla (editors). *The knowledge frontier: Essays in the representation of knowledge* (Symbolic computation series), New York: Springer-Verlag, 1987, 262–315.

[Smullyan 1978] Raymond M. Smullyan. *What is the name of this book?—The riddle of Dracula and other logical puzzles.* Prentice-Hall, 1978.

[Sowa 1984] John F. Sowa. *Conceptual structures: Information processing in mind and machine* (The systems programming series). Reading, MA: Addison-Wesley, 1984.

[Strawson 1950] Peter Frederick Strawson. "On referring." *Mind*, n.s. 59, 1950, 320–344.

[Williams 1981] Christopher John Fards Williams. *What is existence?* (Clarendon library of logic and philosophy). Oxford: Clarendon Press, 1981.

[Woodruff 1970] Peter W. Woodruff. "Logic and truth value gaps." In: Karel Lambert (editor), *Philosophical problems in logic: Some recent developments* (Synthese library), Dordrecht: D. Reidel, 1970, 121–142.

[Zalta 1983] Edward N. Zalta. *Abstract objects: An introduction to axiomatic metaphysics* (Synthese library 160). Dordrecht: D. Reidel, 1983.