Preventing Existence

Cleo Condoravdi, Dick Crouch, Martin van den Berg
John Everett, Valeria Paiva
Reinhard Stolle, Danny Bobrow

Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304
rdc+@parc.xerox.com

FXPAL
3400 Hillview
Palo Alto, CA 94304
vdberg@pal.xerox.com

Abstract: We discuss the treatment of prevention statements in both natural language semantics and knowledge representation, with particular regard to existence entailments. First order representations with an explicit existence predicate are shown to not adequately capture the entailments of prevention statements. A linguistic analysis is framed in a higher order intensional logic, employing a Fregean notion of existence as instantiation of a concept. We discuss how this can be mapped to a Cyc style knowledge representation.

Categories & Descriptors: I.2.4 [**Knowledge Representation Formalisms and Methods**]: Representation languages

General Terms: Theory

Keywords: Existence, knowledge representation, natural language semantics

1. Introduction

Assertions such as (1) pose problems for both knowledge representation and natural language semantics.

(1) Negotiations prevented a strike.

While (1) claims that no strike came into being, one can nevertheless make true or false claims about the non-existent strike, such as (2):

(2) The strike would have been long and bitter.

It will not do to represent the content of (1) along the lines of (3), which asserts the existence of a strike that never came to be.

(3) $\exists x\ \exists y.\ \textit{negotiation}(x) \wedge \textit{strike}(y) \wedge \textit{prevent}(x, y)$

By contrast, a similar representation is at least plausible for causal statements:

(4) a. Negotiations caused a strike.
b. $\exists x\ \exists y.\ \textit{negotiation}(x) \wedge \textit{strike}(y) \wedge \textit{cause}(x, y)$

FOIS '01, October 17-19, 2001, Ogunquit, Maine, USA.

Cause and prevent statements make different existence claims about the event caused or prevented. More generally, the contrast between cause and prevent statements raises the kinds of problems with existence discussed by Hirst [8]. This paper focuses on problems concerning existence that arise specifically with prevent statements. Going beyond the treatments of non-denoting names recently discussed in [4], we hope to sharpen the requirements for a general treatment of existence in knowledge representation and natural language semantics.

The work reported here grows out of a multi-disciplinary research effort to link language processing to knowledge representation, subject to inference, in a domain where causation and prevention feature prominently. The Xerox Linguistic Environment and a broad coverage grammar of English are used to parse English texts. Semantic and discourse interpretation then produce fully scoped logical forms representing their possible linguistic meaning(s). Domain specific reasoning is then used both to convert the linguistically motivated logical forms to Cyc style representations [12], and also to disambiguate them by ruling out linguistically possible interpretations that are implausible given the domain. Within this setting, an adequate treatment of prevention statements is necessary even for low-level information collating tasks, such as matching the objects referred to in two texts. Ontological decisions must be made about what objects, and what kinds of objects, a text refers to, and these decisions must address both linguistic and knowledge representation concerns.

The paper is organized as follows. Section 2 describes some linguistic properties and entailment patterns of cause and prevent statements with particular emphasis on the downward monotonicity entailments and negative polarity licensing of prevent statements. Section 3 discusses Hirst's first order logic with existence predicates, its failure to capture downward monotonicity entailments, and problems of individuation in an ontology that admits non-existent objects. Section 4 presents two higher-order representations that solve the downward monotonicity problem. The first preserves Hirst's existence predicates, but goes higher order in allowing propositions as arguments to predicates. The second dispenses with existence predicates, and builds on Fregean notions of intension and concept [6] that have been mainstream in linguistic semantics since Montague [15]. Section 5 argues that the Fregean treatment can be re-formulated in a way suitable for knowledge representation and reasoning using a first order description logic plus contexts. Section 6 indicates how this approach accounts for certain counterfactual inferences.

2. Properties of Prevent and Cause Statements

2.1 Linguistic properties

Consider past tense[1] statements of the form "*Arg1 caused Arg2*" and "*Arg1 prevented Arg2*". Arg1 is a Noun Phrase (NP), e.g. an NP denoting events, states or properties (e.g., *an accident, my driving too fast, the situation, the wetness of the road, the wet road*), or an NP denoting a causal agent (e.g., *John, the gene*). Arg2 is an event, state or property denoting NP or non-finite clause (e.g., *X prevented/caused an accident/*John, X caused John to see Mary, X prevented John from seeing Mary*).

2.1.1 Scope variation

Prevent and cause statements both permit scope ambiguity.

(5) Negotiations prevented/caused every strike.

[1] In this paper we treat only past tense, episodic statements of cause and prevention. Present tense introduces an added element of genericity that interacts in complex ways with underlying core aspects of cause and prevention. By confining our attention to episodic statement we avoid these complicating interactions and are better able to observe the core phenomena involving existence and non-existence.

This can mean either that one set of negotiations prevented/caused all the strikes (the narrow scope reading of *every strike*), or that each strike was prevented/caused by a different set of negotiations (the wide scope reading of *every strike*).

2.1.2 Negative polarity

A significant difference between *prevent* and *cause* is that *prevent* licenses so-called negative polarity items, such as *any* or *ever*, in its second argument, whereas *cause* does not.

(6) a. X prevented any accidents.
b. *X caused any accidents.

(7) a. X prevented him from ever falling again.
b. *X caused him to ever fall again.

The acceptability of negative polarity items in prevent statements and their unacceptability with cause statements is a manifestation of fundamentally different entailments for cause and prevent statements. (6a) and (7a) provide evidence that Arg2 of *prevent* is within the scope of a *downward monotone expression*, since negative polarity items are acceptable only in such a context [11]. This kind of evidence has a direct bearing on the different entailment properties of cause and prevent statements, discussed in the following section.

2.2 Inference patterns

2.2.1 Monotonicity entailments

Existence or non-existence entailments are tied to the monotonicity inference patterns. *Prevent* is downward monotone in its second argument, *cause* is upward monotone. An expression $\phi[\alpha]$ is downward (or upward) monotone in its argument α if it validates the inferences shown:[2]

Downward Monotone (*prevent*)	Upward Monotone (*cause*)
$\alpha_1 \sqsubseteq \alpha_2 \quad \phi[\alpha_2]$	$\alpha_1 \sqsubseteq \alpha_2 \quad \phi[\alpha_1]$
$\phi[\alpha_1]$	$\phi[\alpha_2]$
Negotiations prevented a long strike	*Negotiations caused a long strike*
$\not\models$ *Negotiations prevented a strike*	$\models$ *Negotiations caused a strike*
Negotiations prevented a/any strike	*Negotiations caused a strike*
$\models$ *Negotiations prevented a long strike*	$\not\models$ *Negotiations caused a long strike*

Note that the reverse entailments for *cause* and *prevent* do not hold. Even if heroic negotiations managed to prevent a long strike, a short strike might still have occurred. And if negotiations caused a strike, it does not follow that they caused a long strike.

2.2.2 General and specific prevention

There are certain interpretations of prevent sentences, which we will call 'specific' interpretations, under which downward monotonicity entailments appear to be suspended. As an example of the difference between general and specific interpretations of *prevent*, consider

(8) Safety procedures at Chernobyl prevented a serious nuclear accident.

[2] $\sqsubseteq$ is a relation of semantic specificity: α_1 is at least as specific as α_2. It amounts to entailment if α_1 and α_2 are propositional denoting expressions, and to the subset relation if α_1 and α_2 are set denoting.

Under the general interpretation (8) is false, since one accident was notoriously not prevented. But under a specific interpretation, managers at Chernobyl may be able to point to at least one near miss where the safety procedures did save the day.

Intuitively, a general claim that "X prevented Y" asserts that there were no occurrences of Y within a certain spatio-temporal location and involving certain participants, and that the reason for this was X. A specific claim that "X prevented Y" asserts that there was at least one chain of events that the speaker has in mind which would have led to Y had it not been for X. The specific claim remains true even if there were other occurrences of Y in the same spatio-temporal location and involving the same participants.

2.2.3 Existential commitment

As noted in the introduction, prevent statements do not commit one to the existence of the kinds of event prevented. Specific and general prevention statements make slightly different existential commitments. In the specific prevention of a strike, other relevantly similar strikes may still occur. In the general prevention, no such strikes may occur.

2.2.4 Counterfactual entailments

As well as being able to make counterfactual claims about prevented events, as in (2), prevent statements support counterfactual entailments of the form

(9) a. Negotiations prevented a strike.
b. Therefore, if there had not been any negotiations, there would have been a strike.

It is in fact this counterfactual entailment that allows for a felicitous use of the definite NP *the strike* in (2).

3. Existence Predicates in a First Order Approach

Hirst [8] describes a first order formalism for capturing patterns of existential inference. He distinguishes several different kinds of existence via particular existence predicates. The ones of relevance here are actual existence and non-actual existence.

Hirst proposes representing *Negotiations prevented a long strike* along the lines of (10) and assumes axioms like (11) to capture the lexical entailments of *prevent* and *cause*.

(10) $\exists x\ \exists y.\ \textit{negotiation}(x) \wedge \textit{strike}(y) \wedge \textit{long}(y) \wedge \textit{prevent}(x, y)$

(11) a. $\forall x\ \forall y.\ \textit{prevent}(x, y) \rightarrow (\textit{actually-exists}(x) \wedge \neg \textit{actually-exists}(y))$
b. $\forall x\ \forall y.\ \textit{cause}(x, y) \rightarrow (\textit{actually-exists}(x) \wedge \textit{actually-exists}(y))$

This representation has the considerable advantage of being first order. However, it faces a number of problems.

First, it cannot capture the downward monotone inference patterns of *prevent* that we have discussed. It predicts that, like *cause*, *prevent* is upward monotone in its second argument. For example the standard entailment (12)

(12) $\exists x\ \exists y.\ \textit{negotiation}(x) \wedge \textit{strike}(y) \wedge \textit{long}(y) \wedge \textit{prevent}(x, y)$
$\models \exists x\ \exists y.\ \textit{negotiation}(x) \wedge \textit{strike}(y) \wedge \textit{prevent}(x, y)$

means that the representation in (10) validates the inference *If negotiations prevented a long strike, then they prevented a strike*. As shown in section 2.2.1 this inference is not valid for general prevention. Consequently, even if this strictly first order representation were adequate for specific preventions, it does not handle general preventions.

A second closely related problem is that a Hirst-style analysis would assign the representation (13c) to both (13a) and (13b).

(13) a. Negotiations prevented a long strike.

b. Negotiations prevented a strike. The strike would have been long.

c. $\exists x\ \exists y.\ \textit{negotiation}(x) \wedge \textit{strike}(y) \wedge \textit{prevent}(x, y) \wedge \textit{long}(y)$

But the statements differ in their entailments: (13b) says that there was no strike, whereas (13a) and (13c) say that there was no long strike, though there may still have been a short one. Hirst's suggested representation cannot capture this distinction because negation over actual existence is too local; it has scope over only one predication, the *actually-exists* predication.

Third, what are the identity criteria for non-existing entities? Is a prevented strike that failed to start at 6am the same as an otherwise identical strike that failed to start at 7am? For actually existing strikes the two would be different, since the properties of 'starting at 6am' and 'starting at 7am' cannot consistently be predicated of the same individual. Suppose this co-predication is also inconsistent for non-existing strikes, so that the non-existent strikes starting at 6am and 7am are distinct. Then (10) has the wrong truth conditions. It allows negotiations to prevent a strike just so long as there is at least one non-existent strike, say the one starting at 6am, that was prevented. This leaves open the possibility that there was another strike starting at 7am, or even at 6:01, that did occur. That is, delaying a strike by an hour or a minute would count as preventing the strike. Intuitively this is incorrect. We must therefore retract our supposition, and conclude that the non-existent strikes starting at 6am and at 7am are identical. But this makes the identity criteria for non-existent entities completely different to those for existent entities. We are not aware of any adequate account of such identity criteria[3].

A more intuitive account of talk of preventing a specific strike is that specificity is invariably taken relative to some description. This description may leave the precise starting time open, merely stating that it is on a certain day rather than at a certain hour. The prevent statement claims that there are no entities meeting the description. This is in contrast to Hirst's account, where there is some non-existent entity that does meet the description.

In summary, Hirst's first order approach captures some of the inference properties of specific interpretations of prevent statements. But the local scope of negation precludes it from accounting for the downward monotone entailments of general prevent statements. Negation cannot be given wider scope because the second argument to *prevent* is entity denoting. Moreover, the lack of clear identity criteria for non-existent entities means that Hirst does not even capture all the entailments of specific prevention.

4. Higher-Order Representations

Hirst's approach fails to capture downward monotonicity entailments as a consequence of being strictly first order, so that the second argument to *prevent* can only be entity denoting. This section discusses two distinct non-first order representations that account for downward monotonicity inferences. The first preserves Hirst's use of existence predicates, but goes considerably beyond it in allowing propositions as arguments to predicates. However, we do not pursue this approach since (i) it inherits Hirst's problem of individuating non-existent entities, and (ii) it does not lend itself to tractable inference. The second approach dispenses with existence predicates, and follows the traditions of Fregean and Montagovian semantics by employing terms referring to higher-order intensional entities such as concepts. We will argue that this linguistically more traditional approach not only captures the correct entailments, but is also well suited to the purposes of knowledge representation, and can be given a more or less first order reformulation using description logics.

[3] Though see [5] for a related attempt concerning arbitrary objects.

4.1 Propositional arguments

One way of accounting for the downward monotone entailments of *prevent* is to maintain Hirst's use of existence predicates, while dispensing with his use of first order logic. Instead we move to an intensional logic that allows propositions as arguments to predicates. This, of course, is a substantial departure from Hirst's original proposal, and commits one to the kind of intensional machinery that he was at pains to avoid.

A representation for (14a) is

(14) a. Negotiations prevented a strike.
b. $\exists y.\ \textit{negotiation}(y) \wedge \textit{prevent}(y, \exists x.\ \textit{strike}(x) \wedge \textit{exists}(x))$
c. $\exists y.\ \textit{negotiation}(y) \wedge \textit{cause}(y, \neg\exists x.\ \textit{strike}(x) \wedge \textit{exists}(x))$

where analyzing *prevent* as *cause not to hold* makes (14b) equivalent (14c). This equivalence can alternatively be captured by means of the following axioms that capture lexical entailments of *cause* and *prevent* about the truth of their propositional arguments:

(15) a. $\forall x\ \forall p.\ \textit{cause}(x, p) \rightarrow p$
b. $\forall x\ \forall p.\ \textit{prevent}(x, p) \rightarrow \neg p$

The monotonicity inference patterns that were problematic under Hirst's original proposal now follow straightforwardly: the negation implicit in *prevent* has scope over more than just an atomic existence predication. For example, we account for the fact that preventing a long strike does not necessarily prevent a strike, whereas causing a long strike necessarily causes a strike through the following standard entailment relations

(16) a. $\neg(\exists x.\ \textit{strike}(x) \wedge \textit{long}(x) \wedge \textit{exists}(x)) \not\models \neg(\exists x.\ \textit{strike}(x) \wedge \textit{exists}(x))$
b. $\exists x.\ \textit{strike}(x) \wedge \textit{exists}(x) \wedge \textit{long}(x) \models \exists x.\ \textit{strike}(x) \wedge \textit{exists}(x)$

Quantifier scope variation is as easily expressed in this representation as in a first order one:

(17) a. Negotiations prevented every strike.
b. $\forall x.\ \textit{strike}(x) \rightarrow (\exists y.\ \textit{negotiation}(y) \wedge \textit{prevent}(y, \textit{exists}(x)))$
c. $\exists y.\ \textit{negotiation}(y) \wedge \textit{prevent}(y, \forall x.\ \textit{strike}(x) \rightarrow \textit{exists}(x))$

However, unlike a first order account scope variation also potentially accounts for the distinction between general and specific prevention:

(18) a. Negotiations prevented a strike.
b. Specific: $\exists x.\ \textit{strike}(x) \wedge \exists y.\ \textit{negotiation}(y) \wedge \textit{prevent}(y, \textit{exists}(x))$
c. General: $\exists y.\ \textit{negotiation}(y) \wedge \textit{prevent}(y, \exists x.\ \textit{strike}(x) \wedge \textit{exists}(x))$

The specific reading (18b) says that some given strike had its existence prevented. The (implicitly negated) *exists* predication is necessary to (i) provide the propositional argument to *prevent*, and (ii) to give an explicit statement of the existential import of the wide scope NP. The general reading (18c) says that the existence of any strike was prevented. However, scope variation is not the only source of the general/specific distinction. Another source is the degree of contextual restriction. By placing sufficiently many contextual restrictions on the domain of potential strikes, either scoping may appear specific.

Contextual restriction relates to an apparent problem raised by Hirst [8]: (18c) would seem to entail that negotiations prevented all possible occurrences of strikes ever. However, any adequate linguistic analysis must include the background context of an utterance, thereby relativizing the description of the strike to some intended set of participants, grievance(s), and spatio-temporal location cf. [16]. A more fleshed-out version of the prevented strike might be

(19) $\exists x.\ \textit{strike}(x) \wedge \textit{strikers}(x, \textit{ford-Employees}) \wedge \textit{s-location}(x, \textit{dagenham-UK}) \wedge$
$\textit{t-location}(x, \textit{june-2000}) \wedge \textit{grievance}(x, \textit{layoffs})$

We return below (section 5.2.1) to the question of how this linguistic contextualization can best be represented.

Combining explicit existence predicates with the use of propositional arguments to predicates like *prevent* addresses downward monotonicity and allows a distinction between specific and general prevention. Neither is possible under Hirst's extensional, first order analysis. But there are a number of drawbacks to this style of intensional analysis. First, it inherits Hirst's problem of individuating non-existent entities due to the presence of an explicit existence predicate. Second, the original motivation for an existence predicate was to remain first order. But if one is going to adopt an intensional logic anyway, there are means for representing existence and non-existence without the need for an explicit predicate, as we will see in section 4.2. Third, from the point of view of practical reasoning, it is unclear that a representation taking propositions as its basic intensional objects will lead to tractable inference. We argue below that taking concepts as the basic intensional objects both avoids the individuation problem, and leads to more tractable representations.

4.2 Quantification over concepts

A Fregean approach holds that existence is not a property of individuals, but of concepts: to say something exists is to say of a concept that it has a non-empty extension. Hirst peremptorily dismisses this as a merely terminological change: from existence of entities to instantiation of concepts. But it is not mere terminology. If nothing else, identity and specificity criteria for concepts are markedly different from those for individuals. In section 3 we pointed out some of the difficulties inherent in individuating non-existent individuals, which are avoided if existence is a property of concepts. More generally, the observation that meaning cannot be reduced to reference has ensured that Fregean notions of concept and intension have been mainstream in linguistic semantics since its formal inception in Montague semantics.

Montague [15] provides an analysis of referentially opaque verbs such as *want*, where the argument to *want* is a (set of) concepts rather than an individual[4], and where differing existential commitments are accounted for scopally:

(20) a. John wants an affordable house in Silicon Valley
b. $\exists x.\ \mathit{house}(x) \wedge \mathit{want}(j, [\lambda P.P(x)])$): There is a specific house he wants
c. $\mathit{want}(j, [\lambda P.\exists x\ \mathit{house}(x) \wedge P(x)])$: He wants there to be one

Here, $\lambda P.P(x)$ is the set of concepts P true of some house x, and $\lambda P.\exists x(\mathit{house}(x) \wedge P(x)$ is the set of concepts true of at least one house, which will be empty if there are none.

We cannot directly apply Montague's analysis of verbs like *want* to *prevent*. This is because the specific reading in (20b) commits one to the existence of a house: specific readings for *prevent* should not carry this commitment. The root of the problem is that in Montague's original analysis NPs quantify over individuals. In more recent elaborations [17; 3] NPs can quantify over concepts. This permits the following more satisfactory style of linguistic analysis[5]

(21) a. Negotiations prevented a strike.
b. Specific: $\exists X.X \sqsubseteq \mathit{Strike} \wedge \exists Y.Y \sqsubseteq \mathit{Negotiation} \wedge \mathit{prevent}(Y, X)$
c. General: $\exists X.X = \mathit{Strike} \wedge \exists Y.Y \sqsubseteq \mathit{Negotiation} \wedge \mathit{prevent}(Y, X)$

Here *Strike* and *Negotiation* refer to the concepts of strike and negotiation, and $\sqsubseteq$ means 'subconcept'. Lexical entailments for $\mathit{prevent}(X, Y)$ ensure that for any concepts X and Y for which the predicate holds, concept X has an instance and concept Y does not.

[4] Strictly speaking the argument to *want* is the *intension* of a set of properties/concepts. We turn to the intensional aspect of concepts in section 5.2.2.

[5] The representations in (21) are simplifications of what a compositionally driven linguistic semantics would provide, ignoring such things as tense, plural/singular and collective/distributive distinctions.

The difference between general and specific prevention is, under this analysis, not scopal. Rather it results from a systematic ambiguity as to whether NPs quantify over subconcepts of the noun or the concept itself. The general reading entails that there is no instance of the concept *Strike*, whereas the specific reading says that there is no instance of some sub-concept of *Strike*. The specific reading still allows for instances of different sub-concepts of *Strike*.

Despite the detour through higher-order intensional logic and Montague semantics, the core of the semantic representation is strikingly close to a knowledge representation that we believe supports tractable inference. It also validates the correct monotonicity entailments and existence commitments, as we will shortly see.

5. A Knowledge Representation with Reified Concepts

In implemented reasoning systems like Cyc [12], concepts are first order entities distinct from their instantiations. Basic concepts are reified as part of a static ontology. While non-basic concepts need to be constructed dynamically, the existence of a static ontology covering much of the domain substantially facilitates the reasoning process. Description logics [1] provide a reasonable set of dynamic concept (property) constructors, and they enable a useful and tractable class of reasoning capabilities.

The use of concept-based knowledge representation to issues of existence and referential opacity is not new (see, e.g., [13]). In this section we extend this kind of treatment first to discuss how knowledge representations can be systematically linked to linguistically motivated semantic representations. We then relativize concept instantiation to contexts [14], and discuss how KR notions of context can be linked to linguistic contexts restricting domains of discourse.

5.1 From linguistic to knowledge representation

For the sentence *negotiations prevented a strike* our semantic interpreter produces two logical forms, repeated below:

(22) a. $\exists X.X \sqsubseteq \mathit{Strike} \ \wedge \ \exists Y.Y \sqsubseteq \mathit{Negotiation} \ \wedge \ \mathit{prevent}(Y, X)$
b. $\exists X.X = \mathit{Strike} \ \wedge \ \exists Y.Y \sqsubseteq \mathit{Negotiation} \ \wedge \ \mathit{prevent}(Y, X)$

Logical forms are converted, through a process of evidential reasoning, to more familiar Cyc style representations. In this case the conversion is straightforward,[6] relying on (i) substitution of equalities, and (ii) replacement of concepts by entities according to lexical entailments stating which concepts are instantiated, such as

(23) $\forall X \ \forall Y.\ \mathit{prevent}(X, Y) \rightarrow (\neg \mathit{instantiated}(Y) \wedge \mathit{instantiated}(X))$

This leads to final representations

(24) a. $\exists y.\ \mathit{negotiation}(y) \wedge \mathit{prevent}(y, \mathit{Strike})$
b. $\exists X.\ X \sqsubseteq \mathit{Strike} \wedge \exists y.\ \mathit{negotiation}(y) \wedge \mathit{prevent}(y, X)$

(where upper cased constants are taken to refer to concepts, and predicates are lower case).

Lexical entailments like (25) also provide rules for the knowledge base. The knowledge base has rules concluding instantiation or non-instantiation of concepts depending on their roles as arguments to particular predicates.

(25) a. $\forall x \ \forall T.\ \mathit{prevent}(x, T) \rightarrow \neg \mathit{instantiated}(T)$
b. $\forall x \ \forall T.\ \mathit{cause}(x, T) \rightarrow \mathit{instantiated}(T)$

[6] Had the semantic interpreter produced representations in the style of section 4.1, it would be problematic to retrieve the correct concept terms from a formula like (17b).

Coupled with these rules, (24a) and (24b) capture within the knowledge base the downward monotonicity and existential entailments required, while also permitting scope variations. Suppose that *prevent*$(n, \textit{Strike})$ is part of the knowledge base. Then the concept *Strike* is concluded to not be instantiated and, therefore, no more specific concept than *Strike* can be instantiated either. Suppose, by contrast, that *cause*$(n, \textit{Strike})$ is part of the knowledge base. Then the concept *Strike* is concluded to be instantiated and, therefore, any more general concept is instantiated as well.

5.2 Contexts, instantiation and intensions

We extend our concept-based representation to one in which we have reified contexts, as introduced by McCarthy [14], making it possible to talk explicitly about the context in which a sentence is interpreted. However, unlike McCarthy we will not take contexts to be undefined primitives, since we link his KR contexts to the distinct kinds of linguistic context used to resolve the interpretation of pronouns and other context-dependent expressions in natural language.

We now contextually relativize instantiations of concepts, so that (25a) instead becomes

(26) $\forall c\ \forall x\ \forall T.\ \textit{ist}(c, \textit{prevent}(x, T)) \rightarrow \textit{ist}(c, \neg \textit{instantiated}(T))$

where *ist* means 'is true in context' and c ranges over contexts. Also, instead of (24b) we write

(27) $\textit{ist}(c_0, \exists X.\ X \sqsubseteq \textit{Strike} \wedge \exists y.\ \textit{negotiation}(y) \wedge \textit{prevent}(y, X))$

where c_0 is the actual context. Although this asserts that there is no instantiation of the concept *Strike* in the actual context, there may be related (non-actual) contexts in which *Strike* is instantiated.

5.2.1 NL and KR contexts

It is important not to leave contexts as undefined primitives, for otherwise it is hard to determine whether and how to contexts are related to one another, e.g., one is a minimal revision of another. We, therefore, define contexts by specific axioms that are true of them.

One example of axioms defining contexts arises from consideration of (19), where it was observed that additional (linguistic) contextual restrictions should be used to limit the domain of strikes over which a quantifier could range. There are two ways of incorporating these restrictions given by linguistic context. One is to explicitly build them into concept expressions whose instantiations we reason about. A more interesting way is to build them into the (KR) context against which instantiations are relativized. If, for example, the linguistic context tells us that we are talking about strikes by Ford employees, one axiom defining the KR context c_0 would be

(28) $\forall x.\ \textit{ist}(c_0, \textit{instantiates}(x, \textit{Strike}) \rightarrow \textit{strikers}(x, \textit{ford-Employees}))$

In practical terms, this has the advantage of keeping the concept terms much simpler. It also means that in cases where linguistic context does not specify which particular group of workers is involved, we can state

(29) $\exists w.\ \textit{workers}(w) \wedge \forall x.\ (\textit{ist}(c_0, \textit{instantiates}(x, \textit{Strike}) \rightarrow \textit{strikers}(x, w)))$

In context c_0 all strikes involve some fixed, but unknown group of workers.

It is also part of the linguistic context of prevent statements that there should be some combination of factors that, were it not for the preventing factor, would interact to cause the occurrence of the prevented event. Often linguistic context will make it clear what these 'causal origins' are. That these causal origins hold true should also be axioms defining the context c_0. In some cases linguistic context may not make it clear what the causal origins are; when this happens, we define the KR context by existentially quantifying over some set of origins, in the same way that we existentially quantified over some set of workers.

5.2.2 Contexts and possible worlds

So far, we have viewed contexts as a way of restricting the extensions of concepts to capture the effects of linguistic restrictions on the domain of discourse. But contexts can do more than this. Contexts in our proposed knowledge representation play the same role as partial possible worlds in the higher order, intensional semantic representation. The objects instantiating a concept can vary from context to context, so that objects may exist in one context and not in another.

Quantifying over individuals across modal contexts raises well known philosophical problems (see [2] for a review).[7] Different possible worlds can contain different individuals. How can we be sure that an individual identified in one world is the same as a counterpart individual identified in another world, or if it even exists in any particular world? This problem is clearly shown by the example of the prevented strike, which exists as a strike in an alternative world but does not correspond to an individual in the actual world.

Our proposed analysis of existence and non-existence avoids this problem by quantifying over concepts and not over individuals. While the domain of individuals may vary from world/context to world/context, the domain of concepts is fixed — all that varies are the extensions of these concepts. For example, the concept associated with the prevented strike is the same in all worlds, though it may be instantiated differently. Moreover, we need not ask whether any entity instantiating the prevented strike in one world is the same as some entity instantiating it in another: all that matters is that the entities instantiate the same concept. In this respect, our analysis embodies a form of essentialism, where what counts is whether two objects instantiate the same concept (essence). But rather than having to identify a fixed set of essential properties in advance, our 'essences' are relativized to the particular linguistic context at hand.

5.2.3 Levels of intensional commitment

In section 4 we discussed two different intensional logics, one taking propositions as its basic intensional units, and the other taking concepts. Here we briefly argue that taking concepts as basic involves a lesser degree of intensional commitment, and that this facilitates more tractable knowledge representation.

Consider the use of concepts. Our ontology identifies a finite set of base concepts, from which complex concepts can be built using a limited set of operations (e.g. as given by description logics). Moreover, we use these mechanisms for building complex concepts in a constrained way. This is because we encode linguistic restrictions on possible instantiations of concepts as axioms defining contexts. This allows us to keep our concept terms much simpler than if we had encoded these linguistic restrictions in the concepts themselves. By starting from a finite base of intensional objects (concepts) and limiting how we construct complex intensional objects, it is feasible to attempt a first order reconstruction that in effect compiles out a lot of the intensionality.

Taking propositions as basic is rather different. There is no realistic way of identifying a base set of propositions from which more complex ones can be constructed. Nor is there any plausible way to limit the construction of complex propositions. This means that we are less able to compile out intensional commitments in a first order reconstruction.

6. Minimal Context Changes & Counterfactuals

We have so far accounted for scope variation, monotonicity and existential entailments for prevent statements, but said nothing about counterfactual entailments. We will focus on just two issues. First, part of the truth conditions of prevent statements includes

(30) a. If A prevented B, then
b. If A had not occurred, then B would have occurred.

[7] Hirst cites these problems as an argument against an intensional treatment.

Second, we can make counterfactual claims about the prevented event, e.g., the strike would have been long. The intuition guiding our treatment is the standard one, e.g. [7]: that although negotiations prevented a strike in the actual context/world, there is some minimally different context/world where the negotiations didn't occur (or failed), and the strike went ahead.

Recall that our KR contexts are defined by axioms that at least (a) restrict the possible instantiations of concepts, and (b) state what causal factors/origin hold true in the context. In a prevention context, it should also be definitional of the context that the preventing factor holds true. Given this, there is an appealing way of defining what should be a minimal revision to c_0 to get a counterfactual context c_1. We simply remove the definitional axiom stating that the preventing factor holds but leave all the other axioms intact. If it follows from the remaining definitional axioms for c_1, plus other general world knowledge axioms that the prevented event occurs in c_1, then the counterfactual is true.

To test the truth of (31a) in context c_0

(31) a. Negotiations prevented a strike.
b. If there had been no negotiations, there would have been a strike

we remove the negotiations from c_0, leaving all other causal origins present in c_1, and see if we can infer the existence of a strike from the axioms defining c_1.

There are other kinds of minimal counterfactual revision. In

(32) Negotiations prevented a strike. It would/might have been long.

we can envisage counterfactual contexts where negotiations took place, but for some reason failed. To account for this, other minimal revisions are obtained by removing some of the causal origins defining c_0.[8] If the revised context entails the existence of a long strike, then the strike *would* have been long. If the revised context entails the existence of a strike and is consistent with it being long, then the strike *might* have been long.

7. Conclusions

This paper has highlighted the downward monotonicity entailments of prevent statements, and the failure of a first order logic with existence predicates to account for them. A higher order intensional logic, of the kind widely used in linguistic semantics, is required for the task. We argue for one incorporating a Fregean notion of existence as instantiation of concepts. For the purposes of inferentially tractable knowledge representation, this semantic analysis can be reformulated in a Cyc style description logic with contexts.

ACKNOWLEDGMENTS

We are grateful for discussions with David Ahn, Ash Asudeh, Mary Dalrymple, Ron Kaplan, Livia Polanyi and Annie Zaenen, and for the comments of three anonymous reviewers.

References

[1] Brachman, R. J., D. L. McGuinness, and P. Patel-Schneider. Reducing classic to practice: Knowledge representation theory meets reality. *Artificial Intelligence* 114, 203–237, 1999.

[2] Cocchiarella, N. B. Philosophical perspectives on quantification in tense and modal logics. In D. Gabbay & F. Guenthner (eds.) *Handbook of Philosophical Logic, Vol II*, 309–353, Reidel, Dordrecht, 1984.

[8] Identifying causal origins and dependencies, and what constitutes a minimal revision, is beyond the scope of this paper. Various proposals exist in the literature, e.g. [10]. See also [9] for a recent discussion of causal dependencies in counterfactuals.

[3] van der Does, J. *Applied Quantifier Logics: Collectives & Naked Infinitives*. PhD thesis, University of Amsterdam, 1992.

[4] Everett, A. and T. Hofweber (eds.). *Empty Names, Fiction and the Puzzle of Non-Existence*. CSLI Publications, Stanford, 2000.

[5] Fine, K. *Reasoning with Arbitrary Objects*. Blackwell, Oxford, 1985.

[6] Frege, G. On sense and reference, 1892. Reprinted in P. T. Geach & M. Black (eds), *Translations from the Philosophical Writings of Gottlob Frege*, Blackwell, Oxford, 1952.

[7] Gärdenfors, P. and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. *Proceedings 2nd Conference on Theoretical Aspects of Reasoning About Knowledge*, 83–95, Morgan Kaufmann, San Francisco, 1988.

[8] Hirst, G. Existence assumptions in knowledge representation. *Artificial Intelligence* 49, 199–242, May 1991.

[9] Kaufmann, S. *Aspects of the Meaning and Use of Conditionals*. PhD thesis, Stanford University, 2001.

[10] Kratzer, A. An investigation of the lumps of thought. *Linguistics and Philosophy* 12, 607–653, 1989.

[11] Ladusaw, W. *Polarity Sensitivity as Inherent Scope Relations*. Garland, New York, 1980.

[12] Lenat, D., R. V. Guha, K. Pittman, D. Pratt, and M. Shepherd. Cyc: Towards programs with common sense. *Communications of the ACM* 33(8), August 1990.

[13] McCarthy, J. First order theories of individual concepts and propositions. In J. E. Hayes, D. Michie and L. I. Mikulich, *Machine Intelligence 9*, Ellis Horwood, Chichester, 1979.

[14] J. McCarthy and S. Buvač. Formalizing context (expanded notes). In A. Aliseda, R. v. Glabbeek, and D. Westerståhl, editors, *Computing Natural Language*. CSLI Publications, Stanford, 1997.

[15] Montague, R. *Formal Philosophy. Selected Papers of Richard Montague*. Yale University Press, 1974. Edited by R. Thomason.

[16] Westerståhl, D. Quantifiers in formal and natural languages. In D. Gabbay & F. Guenthner (eds) *Handbook of Philosophical Logic, Vol IV*, 1–131, Reidel, Dordrecht, 1989.

[17] Zimmermann, E. On the proper treatment of opacity in certain verbs. *Natural Language Semantics* 1, 149–179, 1993.