Logicism in Formalizing Common Sense and in Natural Language Semantics

Richmond H. Thomason
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
thomason@isp.pitt.edu

Abstract

I briefly rehearse a general characterization of logicism that includes projects within philosophical logic, linguistics, and artificial intelligence. After a digression in which I survey work in natural language semantics that I claim is very relevant to common sense logicism, I describe a logicist project that attempts to formalize semantic relations among words in natural languages.

Introduction

This paper has three main goals. I will (1) advocate a more general account of logicism than the one that is current in philosophy of mathematics; (2) motivate a logicist program in an area of natural language semantics that turns out to be closely related to problems in the formalization of common sense; and (3) provide a brief survey of the research tradition in "natural language metaphysics." This tradition, which seeks to provide a broadly logical theory that will account for meaning relations in human languages, is probably best regarded as a subdiscipline of common sense logicism. But since this work is mostly done by linguists who are not computer scientists, and relies to some extent on a different research tradition, it is largely complementary to the work by computer scientists. It would probably be beneficial for everyone if there were more mutual influences between these areas.

Logicism

I can be brief about this, since I have said similar things elsewhere: see, in particular, (61) and (62).

Let X be a topic of inquiry; X logicism is the view that X should be presented as an axiomatic theory from which the rest can be deduced by logic.

Science logicism is expressed as an ideal in Aristotle's *Organon*.[1] But Aristotelian logic is far too weak to serve as a means of formalizing Aristotelian science, and logicism remained impracticable until the 17th century, when a separation of theoretical science from common sense simplified the task of designing an underlying logic.[2]

There is a moral here about logicism. X logicism imposes a program: the project of actually presenting X in the required form. But for such a project to be feasible, the logic must be well matched to the demands of the topic; for instance, a topic that (like Aristotelian science) depends crucially on relational reasoning needs a logic that can deal with relational reasoning.

Thus, assuming that a logic must present explicit formal patterns of valid reasoning, the *first* problem for X logicism is then to articulate formal patterns that will be adequate for formalizing X. The *second* problem is to demonstrate the adequacy of the logic by carrying through the formalization of X.

Solutions to the problem of mathematics logicism eventually emerged from the work of Frege, Russell, Bernays, Tarski, and many others.[3] These formalizations are specific to mathematics; they are not readily extendable to non-mathematical domains. It proved remarkably difficult to develop solutions that met even elementary criteria such as consistency. And as techniques for reasoning about formal systems developed, the idea of consistency and consistency proofs proved to be surprisingly problematic. The moral seems to be that strategic logicist projects (projects that aim at formalizing a new domain) can be difficult and protracted. Success in such projects depends essentially on choosing an appropriate formalism; and the deployment of the formalism can be limited by incompleteness results.

Extensions to Empirical Domains

Philosophical attempts to extend logicist methodology to the empirical world and the common sense foundations of science were disappointing. They rapidly be-

[1] See, for instance, (2, I.2–3).

[2] Despite the simplification, of course, a workable formalism did not begin to emerge until the 19th century.

[3] In the sense of logicism that I am advocating, all those who sought to formalize mathematics were mathematics logicists, whether or not they considered mathematics to be a definitional extension of logic.

came unfashionable and were abandoned by the 1960s. Here, I want to mention just one of the contributing factors: Carnap's inability to provide full definitions of 'dispositionals'.

In (14), Carnap points out that, for instance, 'x is water soluble' cannot be defined by the condition 'At any time at which x is put in water, x dissolves', because this does not give the right results for cases in which an object is made of a non-soluble material, but happens to be never placed in water. This forces him to settle for partial definitions of dispositionals,[4] although the inability to provide full definitions weakens his account of empirical theories considerably.

Linguistic logicism

Richard Montague is largely responsible for the continuation of a logicist tradition in linguistics and (to some extent) in philosophy. In (44), he explores the philosophical problem of determinism by formalizing celestial mechanics. And in a number of places he explicitly advocates the (possibly quixotic) project of *philosophy logicism*. But he is best known for his work in *natural language semantics* logicism, which probably was intended in part as a demonstration of the adequacy of the formalism that he intended for philosophy. This work proved to be much more influential in linguistics than in philosophy. Philosophers tend to resist formalism, but many linguists felt that Montague grammar showed how to remedy the semantic gap in the linguistic theories of the late 1960s. This provided the impetus for the development of an active subfield in linguistics, with connections to logic, philosophy, and artificial intelligence.[5]

This research tradition illustrates the logicist themes that were formulated above. Some of the most interesting work in the area has involved explicit attempts to match the logical formalism to the needs of natural language semantics. Since the logic is inherited from mathematics logicism, it isn't surprising that the research emphasizes areas in which mathematical formalisms overlap with natural language—in particular, there has been a great deal of attention to the semantics of quantification.[6] Thus, one theme has been the exploration of dynamic logics of quantification. Another has been the attempt to develop a theory of "abstract data types" for natural language, based on problems in the correspondences that Montague proposed between syntactic categories and logical types.[7] This work, it seems to me, is less closely related to common sense logicism than cases where natural language semantics seems to have revealed the need for formalisms to deal with phenomena that are rare or entirely absent in mathematical domains, but pervasive in natural language. In the following survey, I will concentrate on the latter areas. This survey is, of course, partial, and reflects my own interests, and my guesses at the interests of the common sense logicists in AI.

[4] A dispositional concept is one that expresses a tendency or disposition for a thing to behave in some way.

[5] The leading journals in the field are *Linguistics and Philosophy* and *Natural Language Semantics*. For one relatively recent survey, see (66).

[6] See (10) for recent work and references.

[7] See, for instance, (17).

Survey of work in natural language metaphysics

The phrase 'natural language metaphysics' was coined, I believe, by Emmon Bach.[8] I like this term. It certainly has the right associations, suggesting as it does a combination of linguistic methodology with metaphysics. In its origins in the book that comes after the Physics in Aristotle's works, metaphysics is closely connected with the principles and foundations of natural language semantics. And Aristotle's work is remarkably insightful when it comes to the shape of common sense. In many cases—particularly in the treatment of verbal semantics, which is explicitly inspired by passages like (3, 1048^b18–36), (4, 1140^a30–$^b12$), (1, 417^a30–$^b2$), through Anthony Kenny's work[9]—the influence is direct and explicit.

Precursors

Perhaps the most important effect of Montague's contribution was the creation of a sustained, unified research tradition combining logical formalization with the empirical approach to language developed by linguists. Before the 1970's, however, there were a number of isolated projects that still remain important and useful as resources for the formalization of common sense as well as semantics. This work deserves to be noted and remembered, though much of it is not much cited nowadays.[10]

Otto Jespersen was far ahead of his time in his attention to foundational matters. (28) is a monumental study of English; (29) concentrates on conceptual foundations, and contains especially useful discussions of tense and aspect, indirect speech, mood, and negation. Jespersen used actually occurring examples, and his discussions are indexed to a large number of quotations, mainly from English-language authors.

There is a tradition in philosophical and linguistic anthropology that uses empirical methodology in an attempt to discover basic categories of human thought. Ernst Cassirer produced, in (16), perhaps the last

[8] See (9, p.112). This book, byt the way, is very readable introduction to the area; readers familiar with logic can skip the earlier chapters.

[9] See (30, Chapter 8).

[10] I'm not a good historian, and I have not tried to sort out influences. Most probably, these people were working independently.

sustained work in this tradition. Using comparative linguistic and anthropological evidence, Cassirer discusses topics such as time, space, and the concept of self.

Carnap's intellectual autobiography in (15) makes it clear that he was more interested in constructing artificial languages than in studying natural ones. But the techniques that he developed (especially those relating to intensionality) could be adapted for use in natural language semantics. Montague and Carnap seem to have shared the goal of formalizing philosophy; like many tasks in knowledge representation, this leads naturally to an interest in human languages.

Arthur Prior was a leading figure among the logicians who tried to extend the logical formalisms designed for mathematical domains for use in domains that are closer to common sense. Prior is largely responsible for the development of tense logic,[11] but his work is wide-ranging enough to cover most of the areas of philosophical logic.[12] Prior's work, and that of other philosophical logicians with similar goals, provides resources that were used in the theories that I will now discuss.

Eventualities and Aktionsarten

As a way to deal with adverbial modification (which also avoids appeal to mechanisms like higher-order logic), Donald Davidson proposed in (18) to introduce quantification over events into the logical form of action sentences. A sentence like

(1) *George surrendered suddenly,*

for instance, would become

(2) $\exists x[surrender(x) \wedge agent(e) = George$
$\wedge\ sudden(e)]$.

This style of formalization has many linguistic advantages, and can be incorporated in Montague-like formal treatments; see (47). (This style of formalization is also used in AI; see (25; 26).) If that is done, it is natural (for the sake of uniformity) to introduce such quantifiers into all sentences, so that

(3) *Ellen is asleep*

would become

(4) $\exists e[asleep(e) \wedge undergoer(e) = Ellen]$.

In this case, it is more natural to say that a *state* is invoked—a state of being asleep. We then need a term for the happenings that are invoked by verb phrases in this style of interpretation. 'Eventuality' is Emmon Bach's term[13] for any state, process, event, or action. Various subtypes of eventualities can be distinguished, based on inferential and syntactic tests. There is a large literature on eventualities and their subtypes;[14] see, for instance, (63; 19; 8; 57; 47). We have already mentioned states and events. Most linguists also recognize a *telic* subtype of events; these are events that inherently involve a goal. (Thus, while *walking* is atelic, *walking to the store* is atelic.) Most linguists feel that telic eventualities involve causality, in part because many of the world's languages use explicit causal language as the normal way of expressing telicity.

Argument structure, thematic roles and diathesis alternations

Readers familiar with knowledge representation will recognize the functions *agent* and *undergoer* in Examples (2) and (4), above, as roles of events. Since the late 1960's, linguists (though they are not in general familiar with the notion of a role hierarchy) have been exploring how to organize such notions into a role hierarchy and how to use this hierarchy to state universal generalizations about the ways in which languages bind verbal constructions to their arguments. As linguists would put it, the problem is how to state the relation between *thematic roles* such as *agent, theme* (or *undergoer*), *instrument*, *experiencer*, *source of path*, and *goal of path* to patterns consisting of a verb and its arguments, and involving devices such as agreement, linguistic case, and word order.

Attempts to provide a semantic characterization of thematic roles lead to very general problems concerning change, agency, and affectedness. These problems are vague and difficult to resolve, and the top level of the role hierarchy remains controversial. But there is interesting information to be obtained from the linguistic insights. The literature on this subject is rather large; I would recommend (20) as a good start for the AI community.[15]

Incorporating eventualities into the formalization of simple sentences also draws attention to the phenomenon of *diathesis alternation*: variations in the syntactic ways in which the same verb can be associated with noun phrases expressing arguments. Often, as with alternations like

(5a) *She searched the pages for a clue*
and
(5b) *She searched for a clue in the pages,*

it seems as if the same eventuality is described in different ways; in others, the syntactic alternations are accompanied by systematic effects on the meaning.

[11] His most important work in this area is (49); but also see the papers in (50).

[12] For instance, see (51).

[13] See (8).

[14] Often these subtypes are called *Aktionsarten,* adopting a grammatical term from German.

[15] The current interest of this topic is largely due to (24) and (21). The literature is by now rather extensive. The bulk of it, however, is concerned with aspects pertaining to syntax and universal grammar that are probably of less interest to those interested in knowledge representation and reasoning. See also (27, Chapter 10), (47) and (65).

The patterns of diathesis alternations that are associated with verbs provide a means of classifying verbs. Systematic studies of this phenomenon, especially (38), make a very good case for the claim that the resulting taxonomy is very closely connected to general semantic distinctions. For example, Beth Levin argues[16] that differences in the diathesis alternations of 'hit' and 'cut' are due to the fact that the latter expresses a change of state verb, while the former does not.

I suspect that results of this methodology provides a useful complement to using direct intuitions about the meaning and reasoning patterns associated with words in formalizing common sense information. If the history of the claims philosophers have made on behalf of common sense are any indicator, common sense can be elusive once you start to reflect on it and systematize it. The linguistic methods provide an independent check on theories of common sense, if we assume, as many have, that natural language and common sense are closely related. Also, the linguistic methods can be supplemented in this area with results obtained from online dictionaries and corpora; see, for instance, (7).

Tense and aspect

Aktionsarten interact strongly with the aspectual systems of language—the apparatus that deals with the way in which eventualities come about, rather than with strictly temporal information. Tense logic is a natural extension of the logics designed for the formalization of mathematics, but theorists of natural language gradually realized that aspect was more problematic than tense.

Aspect is a fairly open-ended category; for purposes of this discussion, let's count *progressive, perfective*, and *habitual* as aspects. Progressive aspect is marked in English by a combination of auxiliary 'be' and suffix '-ing', and can occur in present and past tense, as in (6a) and(6b).[17] Perfective aspect is marked in English by a combination of auxiliary 'have' and the perfective participle, and also can occur in present and past tense, as in (7a) and(7b).

(6a) *The boat is sinking.*
(6b) *The boat was sinking.*
(7a) *The boat has sunk.*
(7b) *The boat had sunk.*

In an early attempt to deal logically with progressive aspect, Dana Scott (as related in (45)) proposed using topological concepts; the idea was that a progressive formula $[ing](A)$ is true at a moment t if and only if A is uniformly true over an open interval containing t. The use of interval semantics to account for aspect has been explored since then, but does appear to have been very successful. The direct appeal to properties of eventualities seems to have worked out better; for a systematic account, see (47).

[16] (38, p. 8).
[17] Future is not a grammatical tense in English; it is usually expressed with the modal auxiliary 'will'.

Linguistic semanticists soon noticed that Montague's account of progressive was inadequate in fundamental respects. In particular, it failed to address the "imperfective paradox"—which is simply the failure of past progressives that involve telic verbs to entail the corresponding past or present perfective forms. Thus, for example, (8a) does not entail (8b).

(8a) *The horse was walking to the stable.*
(8b) *The horse walked to the stable.*
(8c) *The horse has walked to the stable.*

In (19), the idea on an "inertial world" is developed to deal with the imperfective paradox. Inertial worlds are histories in which ongoing telic processes (which can be thought of as processes involving momentum towards the completion of a certain goal) develop normally; progressive aspect is interpreted as a modality that is restricted to these inertial histories. Dowty's logic is monotonic, and his semantics for progressive antedates the development of nonmonotonic logics. But there are clear affinities to nonmonotonic temporal logics. And in fact, nonmonotonic logics have been applied to the semantics of aspect in recent work; see (5) and (35).

This use of nonmonotonic logic is not accepted by everyone; see (34), for example. Many of the objections to this application are really general objections to nonmonotonic logic. However, one puzzle that many authors have mentioned seems worth mentioning; if nonmonotonicity is part of the progressive, the defaults are local, and some conflicting defaults, at any rate, can apparently be true at once. This is illustrated by examples like

(9) *The bear was pulling the hiker down as he was climbing the tree.*

On the other hand, other examples, such as (10), definitely do sound contradictory.

(10) *The hiker was climbing a tree, and at the same time a bear was preventing him from climbing it.*

I am not sure what to make of this. Some of these examples, including perhaps (9), can be reconciled by taking into account the relative granularities of the events, or vagueness in the sort of completion that is associated with a telic event. But we might be driven by such examples to say that inertial defaults hold relative to a context of some sort, and are true if there is an appropriate context in which they hold.

Present habituality is usually expressed in English by simple present, as in the most natural interpretation of

(11) *I walk to work.*

Jespersen's term in (29) for habituality is "generic tense," which (rightly, I think) assimilates this aspect to other sorts of generic constructions.

Generic constructions

(12) *Birds fly,*

the stock example of a default in the nonmonotonic logic literature, is a *generic plural*. There are many other sorts of generic constructions; the indefinite article is also used generically, as well as mass nouns.

Since there is a recent collection of research articles on genericity in natural language, with an excellent introduction,[18] I will not say much about genericity here. In fact, I will only say two things.

Natural language semantics draws attention to sentences with multiple generics. An example like

(13) *A doctor prescribes medicine for his patients to take.*

actually contains four generics: each of the three noun phrases is generic, and the aspect is habitual. The formalization and interpretation of such sentences is problematic.

Surprisingly enough (since generic plural, after all, is the stock example of a default, the use of nonmonotonic logic to interpret natural language generics is controversial. The difficulty is that natural language generics seem to blend smoothly into cases where contextuality seems to be more important than normalcy;

(14) *Deer are eating my shrubs,*

for instance, appears to be generic.[19] A unified theory of contextuality and normalcy might help to resolve this difficulty, if it ever becomes available.

Plurals

There is much disagreement about how to formalize nongeneric plurals, especially group plurals like

(15) *The crowd of citizens gathered in the square*

that are not readily represented by quantification. But the linguistic literature on the topic does give a good picture of the range of phenomema that need to be covered by an adequate theory, though there is no consensus about how to formalize it. Nor is there a clear consensus on the interpretation of the evidence; for instance, (23) and (64) differ on whether there is a difference in the "distributive" and "collective" readings of a sentence like

(16) *Ann and Bob decorated a Christmas tree.*[20]

Here, I will only try to describe very briefly some of the ideas deriving from Godehard Link's treatment of plural in (40); see also (41) and (42).

Logical systems incorporating a mechanism for dealing with the part-whole relation (or *mereological systems*) were developed in the first half of this century by the Polish logicians and the American nominalists.[21] In this spirit, and motivated in part by a desire to find a uniform treatment of plurals and mass terms, Link adds lattice operations to the basic domains needed for natural language interpretation, the lattice ordering being *part-of*. for instance, a whole made up out of Ann and Bob is an individual of the extended domain, and is the least upper bound of the individuals Ann and Bob. The "plural objects" in the extended domains are the nonatomic elements of the lattice. In keeping with the mereological tradition, these pluralities are distinguished from sets. For instance, there is no lattice theoretic counterpart of a set like $\{a, \{a\}\}$. Link shows how to provide a systematic semantic theory using these ingredients which handles many uses of natural language plurals. It was soon noticed, however, that some phenomena, like the group predication in (15), seem to require an extension of the theory. This topic is discussed in (41), (42), and (33). See also (54).

A recent extended study, (52), contains an extended defense of the idea that events can be used to avoid an ontology containing higher-order objects. This work also contains a good bibliography to the literature on plurals.

Stuff

Syntactically, count terms and mass terms differ in that the former can take the indefinite article, while the latter cannot; for instance, 'a lump of snow' is grammatical, 'a snow' is not. The quantifier expressions that go with mass and count nouns differ, though there are systematic similarities: 'no pebbles', 'few pebbles', 'many pebbles', 'all pebbles', as compared with 'no snow', 'little snow', 'much snow', 'all snow'. Count terms can be derived from mass terms: 'a lump of snow', 'a bushel of snow'. There are in fact many similarities between mass terms and plurals; 'splitting a log', for instance, is telic, while 'splitting wood' and 'splitting logs' are atelic.

It is very natural to postulate a distinction between countable and continuous things to account for the grammatical distinction. I will not go into the details of the problems that arise in a systematic attempt to do this, since an excellent survey is available in (48). One promising line is the exploitation of the lattice-theoretic (or more generally the algebraic) approach mentioned in the previous section, in conjunction with a general theory of measurement.[22] This approach has the advantage of explaining some of the systematic correspondences between plurals and mass terms. This work has been pursued by Godehard Link and Manfred Krifka; see, for instance, (40), (42), and (31).

[18] The book is (13), the introduction is (32). Also, see (53).

[19] See (13) for discussion and references.

[20] On the distributive reading, they each have the property of decorating a Christmas tree; on the collective reading, they decorated a Christmas tree together.

[21] (37), (36). See (56) for a useful recent study in this tradition.

[22] Such a theory can be taken from sources such as (58) and (59).

This survey leaves many important areas of natural language semantics untouched. But I have tried to note the areas that strike me as most relevant to the formalization of common sense. At this point, I return to the main train of thought of the paper.

Word meaning logicism

The working hypothesis

Among the available theories of defeasible reasoning that could be applied in lexical semantics, I find circumscription the most congenial to use in attempting to apply these theories to problems of natural language semantics, for the following reasons.

(i) Circumscription is based on second-order logic, and this second-order foundation can easily be generalized to Montague's Intensional Logic.[23] Montague's apparatus does conflict with preferences that McCarthy has expressed from time to time about how to formalize intensional constructions, but this philosophical disagreement does not seem to be an obstacle to the absorption of circumscription into the framework of Intensional Logic.

(ii) The more sophisticated versions of circumscription provide an explicit formalism for dealing with abnormalities.[24] I believe that such a formalism is needed in the linguistic applications.

The thesis, then, is that when the meaning of a word is to be derived from the meaning of another word, the definition can be framed in a relatively natural manner in a circumscriptive version of Intensional Logic. For reasons that are similar to those that have come up in common sense reasoning, it may be necessary to extend the logic to include a theory of context.[25]

The following sections illustrate how the thesis can be developed and tested. This work is still preliminary, and is still changing rapidly.

Telicity

I am abstracting here away from all problems having to do with time and the progressive, and concentrating on the relation between a telic event and its culminating state.[26]

It is convenient to think of events as classified by a system of event types, from which abnormalities and other features are inherited. Thus, in axioms expressing normalities deriving from a verb, the abnormalities will be indexed to the corresponding event type.

Among the information that attaches to event types there may be constraints on the parts that we can expect events to have. Here, to a large extent, I rely on the ideas of (57).

In particular, I assume that events of *telic* type have three associated parts: an inception, a body, and a culmination. The inception is an initiating event. The culmination is the state that normally results. The body is the process that normally leads to the culmination; often (as in closing a door or filling a glass), the body will consist of stages in which the goal is progressively achieved.[27] We can lay it down as a general default on telic events that the culmination of such an event will occur if its inception and body occurs. In many cases the body will also normally occur if the inception occurs. Thus, I am likening unfulfilled telic events to Manx cats—they are objects that belong to a type that normally has a certain part, but that for some *ad hoc* reason happen to lack this part.

The axioms are then as follows.

(17a) For all events e of type *telic-type,* if e occurs then so does $inception(e)$.

(17b) For all events e of type *telic-type* that are not *telic-type-abnormal*1, if e occurs then so does $body(e)$.

(17c) For all events e of type *telic-type* that are not *telic-type-abnormal*2, if e occurs as well as $body(e)$, then so does $culmination(e)$.

Case study: Carnap's dispositionals

The natural way to define 'x is water-soluble' is with a counterfactual conditional, as follows.

(18) *If x were put in some water, then x would dissolve in the water.*

So at first glance, it may seem that the resources for carrying out the definition that Carnap found problematic will be available in a counterfactual logic. But suppose it so happens that if one were to put this salt in some water, it would be this water, and this water is saturated with salt. The fact that the salt would not then dissolve is no reason why the salt should count as not water-soluble. This and other such thought experiments indicate that what is wanted is not the bare subjunctive conditional, but a "normality conditional" of the sort that is used in some nonmonotonic formalisms.[28]

In a circumscriptive framework, normality is obtained by imposing conditions on a number of abnormality predicates, which are then circumscribed, or minimized relative to certain background assumptions, in obtaining models of the nonmonotonic theory.

In treating the dissolving example, I will make the following assumptions.

[23] See (60) for a brief description and application of the combined theory (to discourse, rather than to lexical semantics).

[24] See (39).

[25] See (43).

[26] It should be clear, though, that I have in mind an account that would relate the truth of a progressive sentence to the occurrence of the body of an event.

[27] For ideas that are in some respects similar to these, see (57).

[28] See (11), (6).

(19a) There is an event type *putin-type* of *put-in events.*[29] Associated with any event e falling under this type, there is a container $container(e)$, and a thing moved, $movee(e)$.

(19b) The event type *putin-type* has a subtype *putsaltinwater-type*. For any event e of this type, $container(e)$ is a quantity of water and $movee(e)$ is a quantity of salt. There is an abnormality predicate associated with *putsaltinwater-type*.

(19c) There is an event type *dissolve-type* of *dissolving events.* Associated with any event e of this type there are eventualities $inception(e)$, $body(e)$, and $culmination(e)$ (where the first two are events and the last is a state); also, there is an associated medium $medium(e)$ and a thing dissolved $dissolvee(e)$. There is an abnormality predicate associated with this type.

It will follow from general considerations about *putsaltinwater-type* that if a *putsaltinwater-type*-normal event of this type occurs, its associated culmination state will also occur. (See the remarks above on telicity.)

Given this information about event types, the sort of analysis that I currently favor for dissolving amounts to this.[30]

(20) Every *putsaltinwater-type*-normal *putsaltinwater-type* event e_1 is also the inception of a *dissolvesaltinwater-type* event e_2 such that

$container(e_1) = medium(e_2)$
and
$movee(e_1) = disolvee(e_2)$.

Agency

Some formalisms of agency in AI involve a separation of events into those that are in an agents' immediate repertoire and those that are not.[31] If such a division is adopted for linguistic purposes, we can capture agency—at least, for telic events that normally follow from their inceptions—as follows.

(21) $Do(x, e)$ holds iff the inception e' of e is identical to an immediate action e'' that is performed by x, provided that the body of e'' occurs.[32]

[29] This event type itself has a decompositional analysis, but we can ignore that for purposes of the example.

[30] SOme of the complexities in the formalization are due to a decision to use other mechanisms, such as appeal to the parts of events, to do away with any need for a notion of "direct causality." I am suspicious of direct causality, and would prefer to avoid any appeal to it.

[31] See especially (46).

[32] I want to say that the body of a telic event occurs even if it is partial or incomplete.

For example, it follows from this account that in case someone puts a piece of salt in water and it then dissolves in the ordinary way, then this person has also performed the action of dissolving it in water, assuming that putting the salt in water is immediate an immediate action. Moreover, the action of putting the salt in water will be the inception of the dissolving event.

On this treatment, we dispense with an explicit use of any causal notions in the analysis of agency—though causal notions are certainly implicit if we believe that there is a connection between sequences of events conforming to patterns of normality and causal sequences. Since an explicit theory of common sense causality is not likely to be easy, I prefer such eliminative accounts. However, I'm not sure if explicit causality can be eliminated in general from the theory of agency.

Causality

It is very common for causative verbs to be lexically derived from predicates denoting states; in English, for instance, the verb 'darken' is derived in this way from the adjective 'dark'. And constructions like 'wipe clean' also seem to involve a causal component. This means that a notion of causality needs to be available for axiomatizing lexical derivations.

Much of (19) is devoted to the analysis of causative constructions. Dowty's work makes it clear that the crucial form of causality that is needed for purposes of word meaning is a relation between an individual and a proposition—this is the sense in which an agent causes something to be true. Dowty attempts to define this notion in terms of a causality that relates a pair of propositions—this is the sense the truth of something causes the truth of another. But (something that as far as I know is not noted in the literature) Dowty's definition of agent causality[33] is hopelessly flawed. It implies that everyone causes anything that anyone causes.

As far as I know, there is no way to repair this flaw by refining the definition. But perhaps an action-driven approach to change might help; I am currently exploring the possibilities of taking the fundamental notion of causality to be a relation between an action and a proposition. This approach seems quite promising.

The '-able' suffix in general

Carnap's problem is really a special case of accounting for the meaning of words derived using the '-able' suffix. As with many lexical processes, there are a number of exceptions (like 'amenable') that have to be treated in an *ad hoc* manner. And there are a number of roughly similar patterns of derived meaning. Carnap's case fits into a general pattern of meanings having to do with the idea that if a test action of some sort is performed, a certain state will normally or easily result. The word 'excitable', as in 'This horse is

[33] The idea is that agent a causes p in case for some property P, the proposition that a has P causes p.

excitable' is another example of this type. In other cases, the word means not that the test action normally results in the state, but will normally result in a version of the state that is desirable: 'drinkable', as in 'This water is drinkable', is a case of that sort. In other cases, the word means that a thing is worthy of being in a certain condition: 'admirable', as in 'That person is admirable', is such a case.

To some extent, these patterns are predictable; for instance, verbs of evaluation (like 'admire') seem to invite the last pattern with '-able'.

The, second pattern, illustrated by 'drinkable', is interesting because it leads to a need for an account of practical ability, and connects with issues in the theory of planning. The account that I'll sketch here may seem circular or trivial. The reason (I hope) is that the background hasn't been filled in.

Let's suppose that there is a propositional constant *practical_abnormality* that is used in practical reasoning to reject alternatives because of utility considerations. That is, if a contemplated practical alternative is shown to lead (perhaps with the aid of defaults) to this constant, the alternative has thereby been shown to one that can be ruled out of consideration. A qualitative account of practical reasoning would have to relate this constant to desires and contingent circumstances.

The definition of practical ability would then be the following, where $\Box$ represents temporal necessity.

(22) $can(\phi) \leftrightarrow \neg\Box[\phi \rightarrow practical\_abnormality]$

Artifacts

Many artifacts are defined in terms of their normal uses. This suggests decompositional analyses such as the following example.

(23) A *fastener* is an object x such that, where *use-type* is the event type of using x, every *use-type*-normal occurrence of an event e of type *use-type* is such that $purpose(e)$ is to fasten an object to another object.

The problem of reasoning

Many researchers in the area of knowledge representation have advocated a close association between the development of systems of representation and reasoning systems. For many practitioners, this may mean only that attention must be paid to the complexity of some of the algorithms associated with the representations; for others, it may mean that the representations should be deployed in a working system whose performance has been tested in practical terms.[34] On the other hand, John McCarthy, one of the founders of AI and of the field of knowledge representation, seems to have steadfastly maintained a platonic research agenda, according to which we first clarify in logical terms the knowledge that we are trying to build into a reasoning system, without distracting ourselves with heuristics and implementation details.

Since it seems to me that we may not yet have adequate theories in which to describe our goals, and that attention to the design of successful systems is a likely source of good theories, and a good way of ensuring that they will ultimately have some reasonable relation to technology, I am a moderate on this issue.

However, I would not have been able to write this paper without pretending to be a platonist. Natural language semantics is such a large domain that it is very hard to think about it in the very general way that linguistics seems to demand while maintaining a good relation to some feasible reasoning task. The task of axiomatizing the theories that I have been contemplating in this paper is forbidding, and, if the axiomatization could be carried out, it is hardly likely that it would form the basis of a workable implementation.

While it is very difficult to see how a single project can serve the needs of logic, linguistics, and implementation, there is no reason why a larger cooperative enterprise could not do justice to all of these factors. Linguistic adequacy and logical clarity can help to make a large knowledge base intelligible, natural, and coherent. As such knowledge bases are developed, we can incrementally associate reasoning procedures of many kinds with them. This is perhaps another argument for the sort of cooperation that I called for at the beginning of this paper.

[34] See (12).

References

[1] Aristotle, "De anima."

[2] Aristotle, "Posterior analytics."

[3] Aristotle, "Metaphysics."

[4] Aristotle, "Nicomachean ethics."

[5] N. Asher, "A default, truth conditional semantics for the progressive," *Linguistics and Philosophy*, vol. 15, pp. 469–508, 1992.

[6] N. Asher and M. Morreau, "Commonsense entailment: a modal theory of nonmonotonic reasoning," in *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, (J. Mylopoulos and R. Reiter, eds.), (Los Altos, CA), pp. 387–392, Morgan Kaufmann, 1991.

[7] B. T. Atkins, J. Kegl, and B. Levin, "Anatomy of a verb entry: from linguistic theory to lexicographic practice," in *Current Issues in Computational Linguistics: Essays in Honour of Don Walker*, (A. Zampolli, N. Calzolari, and M. Palmer, eds.), pp. 237–266, Pisa and Dordrecht: Giardini Editori e Stampatori and Kluwer Academic Publishers, 1994.

[8] E. Bach, "The algebra of events," *Linguistics and Philosophy*, vol. 9, pp. 5–16, 1986.

[9] E. Bach, *Informal Lectures on Formal Semantics.* Albany, NY: State University of New York Press, 1989.

[10] E. Bach, E. Jelinek, A. Kratzer, and B. Partee, eds., *Quantification in Natural Languages, Vols. 1 and 2.* Vol. 54 and 55 of *Studies in Linguistics and Philosophy*, Dordrecht: Kluwer Academic Publishers, 1995.

[11] C. Boutilier, "Conditional logics for default reasoning and belief revision," Tech. Rep. KRR–TR–92–1, Computer Science Department, University of Toronto, Toronto, Ontario, 1992.

[12] R. J. Brachman, "The future of knowledge representation," in *Proceedings of the Eighth National Conference on Artificial Intelligence*, (Menlo Park, CA), pp. 1082–1092, AAAI Press, 1990.

[13] G. Carlson and F. J. Pelletier, eds., *The Generic Book.* Chicago, IL: Chicago University Press, 1995.

[14] R. Carnap, "Testability and meaning," *Philosophy of Science*, vol. 3 and 4, pp. 419–471 and 1–40, 1936–1937.

[15] R. Carnap, "My philosophical development," in *The Philosophy of Rudolph Carnap*, (P. Schilpp, ed.), pp. 3–84, LaSalle, IL: Open Court Publishing Company, 1963.

[16] E. Cassirer, *The Philosophy of Symbolic Forms.* Vol. 1–3, New Haven: Yale University Press, 1955. These works were first published (in German) in 1923–29.

[17] G. Chierchia, B. Partee, and R. Turner, eds., *Properties, Types and Meaning, Vols. 1 and 2.* Vol. 38 and 39 of *Studies in Linguistics and Philosophy*, Dordrecht: Kluwer Academic Publishers, 1988.

[18] D. Davidson, "Actions, reasons and causes," *Journal of Philosophy*, vol. 60, pp. 685–700, 1963.

[19] D. Dowty, *Word Meaning and Montague Grammar.* Dordrecht, Holland: D. Reidel, 1979.

[20] D. Dowty, "On the semantic content of the notion of 'thematic role'," in *Properties, Types and Meaning, Vol. 2*, (G. Chierchia, B. Partee, and R. Turner, eds.), pp. 69–129, Dordrecht: Kluwer Academic Publishers, 1988.

[21] C. J. Fillmore, "The case for case," in *Universals in Linguistic Theory*, (E. Bach and R. Harms, eds.), pp. 1–88, New York: Holt, Rinehart and Winston, 1968.

[22] H. Geffner, *Default Reasoning: Causal and Conditional Theories.* Cambridge, MA: The MIT Press, 1992.

[23] B. Gillon, "The readings of plural noun phrases in English," *Linguistics and Philosophy*, vol. 10, no. 2, pp. 199–219, 1987.

[24] J. S. Gruber, *Studies in Lexical Relations.* Amsterdam: North-Holland, 1976. Originally published in 1965 as an MIT dissertation.

[25] J. Hobbs, M. Stickel, D. Appelt, and P. Martin, "Interpretation as abduction," *Artificial Intelligence*, vol. 63, no. 1–2, pp. 69–142, 1993.

[26] C. H. Hwang and L. Schubert, "Episodic logic: a situational logic for natural language processing," in *Situation Theory and its Applications*, (S. Peters, D. Israel, P. Aczel, and Y. Katagiri, eds.), Center for the Study of Language and Information, 1993.

[27] R. Jackendoff, *Semantics and Cognition.* Cambridge, MA: The MIT Press, 1983.

[28] O. Jespersen, *A Modern English Grammar on Historical Principles.* Northampton, England: John Dickens and Company, 1961. Seven Volumes. This work was originally published in 1909.

[29] O. Jespersen, *The Philosophy of Grammar.* New York: W.W. Norton, 1965. This work was originally published in 1924.

[30] A. Kenny, *Action, Emotion and The Will.* London: Routledge and Kegan Paul, 1963.

[31] M. Krifka, "Four thousand ships passed through the lock: object-induced measure functions on events," *Linguistics and Philosophy*, vol. 13, no. 5, pp. 487–520, 1990.

[32] M. Krifka, F. J. Pelletier, G. Carlson, A. ter Meulen, G. Chierchia, and G. Link, "Genericity: an introduction," in *The Generic Book*, (G. Carlson and F. J. Pelletier, eds.), pp. 1–124, Chicago, IL: Chicago University Press, 1995.

[33] F. Landman, "Groups. I," *Linguistics and Philosophy*, vol. 12, no. 5, pp. 559–605, 1989.

[34] F. Landman, "The progressive," *Natural Language Semantics*, vol. 1, pp. 1–32, 1992.

[35] A. Lascarides, "The progressive and the imperfective paradox,". Synthese, forthcoming.

[36] H. Leonard and N. Goodman, "The calculus of individuals and its uses," *Journal of Symbolic Logic*, vol. 5, 1940.

[37] S. Leśniewski, "Podstawy ogólnej teorii mnogości I.," [Foundations of a General Set Theory I.] Moscow, 1916.

[38] B. Levin, *English Verb Classes and Alternations: a Preliminary Investigation.* Chicago, IL: University of Chicago Press, 1993.

[39] V. Lifschitz, "Circumscriptive theories," in *Philosophical Logic and Artificial Intelligence*, (R. Thomason, ed.), pp. 109–159, Dordrecht: D. Reidel, 1989.

[40] G. Link, "The logical analysis of plurals and mass terms: a lattice-theoretical approach," in *Meaning, Use, and Interpretation of Language*, (R. Bäuerle, C. Schwarze, and A. von Stechow, eds.), pp. 302–323, Berlin: Walter de Gruyter, 1983.

[41] G. Link, "Hydras: on the logic of relative clause constructions with multiple heads," in *Varieties of Formal Semantics*, (F. Landman and F. Veltman, eds.), pp. 245–257, Dordrecht: Foris, 1984.

[42] G. Link, "Plural," in *Semantik/Semantics: an International Handbook of Contemporary Research*, (D. Wunderlich and A. von Stechow, eds.), Walter de Gruyter, 1991.

[43] J. McCarthy, "Notes on formalizing context," in *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, (Los Altos, CA), pp. 81–98, Morgan Kaufmann, 1993.

[44] R. Montague, "Deterministic theories," in *Decisions, Values, and Groups*, pp. 325–370, Oxford: Pergamon Press, 1962. Reprinted in *Formal Philosophy*, by Richard Montague, Yale University Press, New Haven, CT, 1974, pp. 303–359.

[45] R. Montague, "Pragmatics and intensional logic," *Synthese*, vol. 22, pp. 68–94, 1970. Reprinted in *Formal Philosophy*, by Richard Montague, Yale University Press, New Haven, CT, 1974, pp. 119–147.

[46] R. C. Moore, "A formal theory of knowledge and action," in *Formal Theories of the Commonsense World*, (J. R. Hobbs and R. C. Moore, eds.), pp. 480–519, Norwood, NJ: Ablex Publishing Corporation, 1988.

[47] T. Parsons, *Events in the Semantics of English: a Study in Subatomic Semantics.* Cambridge, Massachusetts: MIT Press, 1990.

[48] F. J. Pelletier and L. Schubert, "Mass expressions," in *Handbook of Philosophical Logic, Volume 4*, (D. Gabbay and F. Guenthner, eds.), pp. 327–407, Dordrecht: D. Reidel, 1989.

[49] A. Prior, *Past, Present and Future.* Oxford: Oxford University Press, 1967.

[50] A. Prior, *Papers on Time and Tense.* Oxford: Oxford University Press, 1968.

[51] A. Prior, *Objects of Thought.* Oxford: Oxford University Press, 1971. Edited by Peter Geach and Anthony Kenny.

[52] B. Schein, *Plurals and Events.* Cambridge, MA: MIT Press, 1993.

[53] L. Schubert and F. J. Pelletier, "Problems in the representation of generics, plurals, and mass nouns," in *New Directions in Semantics, Volume 2*, (E. LePore, ed.), pp. 385–451, London: Academic Press, 1987.

[54] B. Schwarzschild, *On the Meaning of Definite Plural Noun Phrases.* PhD thesis, University of Massachusetts at Amherst, Amherst, Massachusetts, 1991.

[55] Y. Shoham, *Reasoning About Change: Time and Causation From the Standpoint of Artificial Intelligence.* Cambridge, Massachusetts: The MIT Press, 1988.

[56] P. Simons, *Parts: A Study in Ontology.* Oxford: Oxford University Press, 1987.

[57] M. Steedman and M. Moens, "Temporal ontology in natural language," in *Proceedings of the 25th Annual Conference of the Association for Computational Linguistics*, pp. 1–7, Association for Computational Linguistics, 1987.

[58] S. S. Stevens, "Measurement, psychophysics and utility," in *Measurement: Definitions and Theories*, (C. W. Churchman and P. Ratoosh, eds.), pp. 18–63, New York: John Wiley and Sons, 1959.

[59] P. Suppes and J. Zinnes, "Basic measurement theory," in *Handbook of Mathematical Psychology, Volume 1*, (R. D. Luce, R. Bush, and E. Galanter, eds.), John Wiley and Sons, 1963.

[60] R. Thomason, "Accommodation, meaning, and implicature: interdisciplinary foundations for pragmatics," in *Intentions in Communication*, (P. Cohen, J. Morgan, and M. Pollack, eds.), pp. 326–363, Cambridge, Massachusetts: MIT Press, 1990.

[61] R. Thomason,, "Logicism, artificial intelligence, and common sense: John McCarthy's program in philosophical perspective," in *Artificial Intelligence and Mathematical Theory of Computation*, (V. Lifschitz, ed.), pp. 449–466, San Diego: Academic Press, 1991.

[62] R. Thomason, "Logicism: exact philosophy, linguistics, and artificial intelligence," 1995. Available via http://www.pitt.edu/~thomason/thomason.html.

[63] Z. Vendler, "Verbs and times," *Philosophical Review*, vol. 46, pp. 143–160, 1957.

[64] H. Verkuyl and J. van der Does, "The semantics of plural noun phrases," Tech. Rep. LP-91-07, Institute for Language, Logic and Information, University of Amsterdam, Faculty of Mathematics and Computer Science, Roeterssraat 15, 1018WB Amsterdam, Holland, 1991.

[65] W. Wilkins, ed., *Syntax and Semantics 21: Thematic Relations.* New York: Academic Press, 1988.

[66] D. Wunderlich and A. von Stechow, eds., *Semantik/Semantics: an International Handbook of Contemporary Research.* Berlin: Walter de Gruyter, 1991.