Prepositions as abstract relations

Allan Ramsay
School of Informatics,
University of Manchester, M60 1QD, UK

Abstract

Many natural languages use prepositions to mark relations between entities of various kinds – between physical entities and their spatial locations, between temporal entities and their temporal locations, between abstract entities of various kinds (e.g. between ideas and their 'mental locations'). In the current paper I will show that the consequences of using prepositions to relate *temporal* entities emerge naturally from the basic interpretations of the prepositions themselves together with a very weak logic of events, where I take it that prepositions denote abstract relations whose significance only emerges when properties of the related items are taken into consideration.

1 Prepositions as abstract relations

> "*in: preposition expressing inclusion or position within limits of space, time, circumstance, etc.*"
> *Pocket Oxford Dictionary of Current English*, 1970

Consider the following sets of sentences:

(1) a. He drove from London to Paris.[1]
b. He drove from dawn to dusk.
c. He copied it from his hard drive to a floppy.
d. I got the RAE results from the Guardian[2].

(2) a. I saw a man in the park.
b. I saw him in January.
c. I've got an idea for a paper about time in my mind, but I don't know how it will work out.
d. I've got 230 people in my AI class next term.

(3) a. I saw him at the station.
b. I saw him at two o'clock.
c. He seemed to be at his wits' end.

All of these, apart perhaps from the very last one, seem to be entirely natural. The prepositions that appear in them, however, clearly link very different kinds of things. In the (a) examples the ground for the relation denoted by the PP is a physical location, in the (b) examples it is a temporal location, in the (c) and (d) examples it seems to be some sort of abstract entity. Is one of these relations primary, with the others as some kind of metaphorical extension, or is there some common element which gets extended in different ways depending on the nature of the ground?

This is not, in this paper, a question about the history of prepositions, either over the evolution of the language across the centuries or within each language learner. It may be that, at least historically, the spatial interpretations of some of these prepositions come first, but there is very little evidence that the spatial readings predominate in the language as

[1](1a) to (1c) are taken from [Dowty, 1979].

[2]because they ranked UMIST higher than anyone else did!

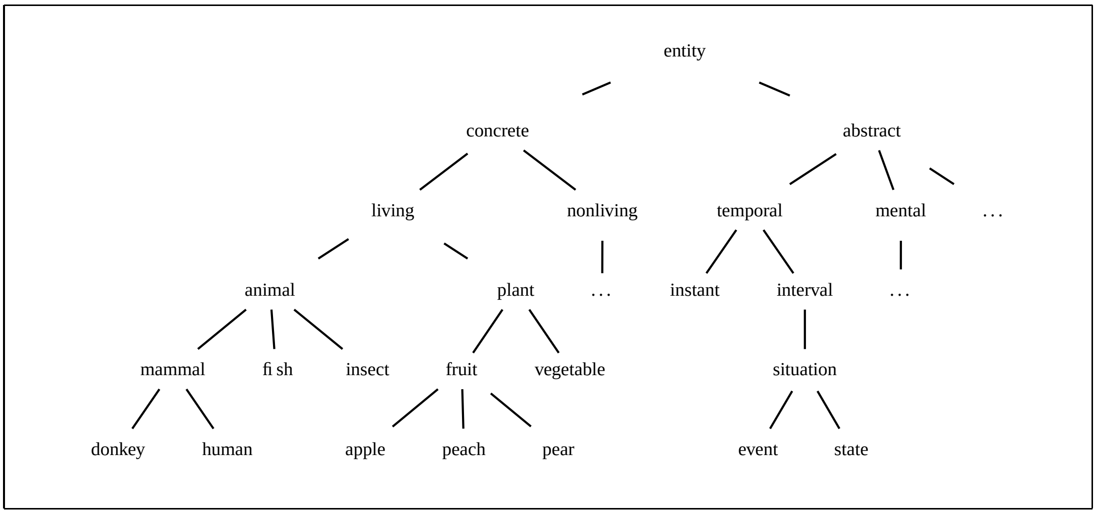

Figure 1: Extracts from the basic sort hierarchy

it is used today (see Appendix A for an extract from the first 9000 noun-*in*-noun triples in the BNC. There is very little indication here that spatial or temporal uses of *'in'* are particularly common[3]). The aim of the current paper is to see whether it is possible to provide a single uniform account of the meaning of prepositions like the ones in (1)–(3) which shows how the contribution they make to the meaning of the sentence, and in particular to the temporal structure of the reported event, arises from a simple core meaning for the preposition in conjunction with the key properties of the entities being related.

The key observation is that most objects can be seen from multiple points of view. A reading event, for instance, can be seen as a physical object involving light being reflected from a marked surface into someone's eyes, as a temporal object with a duration over which the event takes place, and as a mental object involving transfer of information into someone's mind.

Different events have different sets of views, where each view is assigned a position within a natural kind hierarchy for which we have a constant (very short!) time algorithm for checking whether two types are compatible (i.e. that they lie on the same branch of the hierarchy). We write $X \sim Y$ to say that X and Y have compatible types (so that if some object was described as being *human* at one point in our knowledge base and at another as being *animal* then the algorithm would check that these assignments are compatible with reference to the type hierarchy, extracts from which are shown in Fig. 1).

Reading has physical, temporal and mental views, eating has physical and temporal views, knowing has mental and temporal views, and so on. This is unsurprising. Different events are different, involving different kinds of things. If the object of an event is a mental entitity such as a thought then the event will necessarily have a mental view, if its object is a physical entity such as a peach then the event will have a physical view. It is, however, notable that *all* events have a temporal view. We therefore take it that the primary view of an event is as a temporal object. The idea that you can see a single object from a number of different points of view is similar to Pustejovsky [1991]'s 'qualia structure', except that the current analysis allows an arbitrary number of views, of arbitrary types, whereas Pustejovsky argues for a small fixed range of qualia. We denote the fact that X is of type E by writing $X \in E$.

We exploit this in the analysis of natural language utterances by constructing a logical form which aims to contain all and only the

[3]The process of extracting these was not sensitive to the possibility that the PP modifies a VP rather than a nominal, but in any case it is the nature of the ground that has most effect on the type of relation, and there will not be much wrong with the choice of ground in these examples.

$$\exists A : \{A \in interval \& real(A) \& ends\_before(ref(\lambda B(now(B))), A)\}$$
$$\quad \exists C : \{aspect(simple, A, C)\}$$
$$\quad\quad \theta(C, agent, ref(\lambda D(centred(D, ref(\lambda E(cdiscourse(E)))) \& m(D))))$$
$$\quad\quad \& drive(C) \& C \in event$$
$$\quad\quad \& from(C, ref(\lambda F(named(F, London) \& to(F, ref(\lambda G(named(G, Paris)))))))$$

Figure 2: A driving event that started in London and ended in Paris

information that is explicitly encoded by the sentence[4], and then backing this up with a set of meaning postulates which flesh out the content of the terms that appear in the logical form. For (1a), for instance, we obtain the logical form in Fig. 2.

The logical form in Fig. 2 is constructed by orthodox compositional techniques, using a 'codebook' to make it easy to derive different interpretations (e.g. containing different levels of detail) for different purposes, in much the same way as [Konrad et al., 1996]. The particular variant given in Fig. 2 contains information about the thematic roles of the various entities, where $\theta(X, T, Y)$ says that X is playing the role T with respect to the event Y; about the temporal viewpoint, encoded by recording the aspectual form of the sentence; and about the use of referring expressions (pronouns, definite NPs, referential tenses, etc.) which have to be 'anchored' with respect to the current discourse state. Most of the content of Fig. 2 is fairly standard: all we are doing is producing a formal paraphrase in some appropriate language (we use Turner [1987]'s 'property theory', but there are plenty of defensible alternatives) where the formal paraphrase contains, as far as possible, all and only the information explicitly encoded by the given natural language sentence.

2 Multiple views

A logical form of this kind, however, does nothing for you unless you know what driving is like, what agents do, and so on. We therefore back up our logical forms with sets of meaning postulates, as in Fig. 3.

$$\forall A : \{drive(A)\}$$
$$\quad extended(A)$$
$$\quad \& bounded(A)$$
$$\quad \& closed(A)$$
$$\forall A : \{drive(A)\}$$
$$\quad \exists B view(A, B)$$
$$\quad\quad \& path(B) \& (B \in concrete)$$

Figure 3: What is driving like?

These MPs start to describe what driving is like. The first MP in Fig. 3 says things about the temporal view of a driving event. The second one says that you can view a driving event as a concrete (i.e. physical) path. There is plenty more to be said about what driving events involve (e.g. that they typically involve some kind of vehicle), but the observations in Fig. 3 are all we need here.

We thus have two views of our driving event, one coming directly from the logical form saying that it can be seen as a temporal entity and one arising from the MP saying that it can be seen as physical path. What can we do with them?

Clearly we have to know what extended closed bounded intervals are like, since the MPs in Fig. 3 say that driving events can be seen as closed bounded intervals, and what paths are like, and we also have to know what *'to'* means, since the logical form says that this driving event is *'to'* Paris. Fig. 4 shows the basic axioms we use for reasoning about time and about paths in general, and Fig. 5 explains what *'to'* means.

[4]The notations $\forall X : \{P\}Q$ and $\exists X : \{P\}Q$ are convenient shorthands for $\forall X(P \rightarrow Q$ and $\exists X(P \& Q)$. We use them in the construction of logical forms because they help manage the context of the restrictor for universal quantifi ers when these are subjected to rescoping.

$$\forall A : \{A \in instant\}$$
$$\quad \forall B : \{B \in instant\}$$
$$\quad\quad \forall C : \{C \in instant\}$$
$$\quad\quad\quad A > B \& B > C$$
$$\quad\quad\quad \rightarrow A > C$$

$$\forall A \forall B \neg(A > B \& B > A)$$

$$\forall A : \{lbounded(A)\}$$
$$\quad \exists B \forall C : \{member(C, A)\}$$
$$\quad\quad C > B$$

$$\forall A bounded(A)$$
$$\quad \leftrightarrow lbounded(A) \& ubounded(A)$$

$$\forall A : \{lclosed(A)\}$$
$$\quad \exists! B member(B, A) \& start(B, A)$$

$$\forall A : closed(A)$$
$$\quad \leftrightarrow lclosed(A) \& uclosed(A)$$

$$\forall A \forall B : \{member(B, A)\}$$
$$\quad \forall C : \{start(C, A)\} B \geq C$$

$$\forall A : \{extended(A) \& A \in interval\}$$
$$\quad \forall B : \{start(B, A)\}$$
$$\quad\quad \forall C : \{end(C, A)\} C > B$$

$$\forall A : \{path(A)\} closed(A)$$

Figure 4: A weak logic of time (and paths)

The set of axioms in Fig. 4 is very weak. It says that the set of instants is partially ordered; that some intervals are bounded above and/or below; that some bounded intervals have greatest lower bounds and/or least upper bounds, which we call their starts and ends; and that some intervals contain their starts and ends[5]. Note that we do *not* insist that every bounded interval has a start and an end: the temporal logic that underlies natural language does not require intervals to have this property, and indeed in some cases it seems useful to consider bounded intervals that lack greatest lower bounds or least upper bounds. We are therefore assuming that the idea of time that underpins natural language has a model which is either smaller than the set of real numbers (e.g. that it maps onto the rationals) or larger than the set of reals (e.g. that it includes infinitesimals [Keisler, 1976]). The axioms in Fig. 4 also say nothing about whether the time line is dense, though most reasonable models will have this property.

The key MP for *'to'* is given in Fig. 5 (the one for *'from'* is too similar to include).

$$\forall A \forall B : \{to(A, B)\}$$
$$\quad \forall C : \{view(A, C)\}$$
$$\quad\quad \forall D : \{view(B, D)\}$$
$$\quad\quad\quad \forall E : \{end(E, C) \& D \sim E\}$$
$$\quad\quad\quad\quad D = E$$

Figure 5: What *'to'* means

This MP says that if *to(A, B)* holds then if C and D are views of A and B and the end E of C is of the same general type as D then D and E are the same thing. This MP looks vacuous, in the same way as the definition of *'in'* quoted at the start of the paper, since it seems merely to reduce *to(X, Y)* to *end(X, Y)* without saying what ends are like. In fact it does most of the important work of explaining what *'to'* means. Its job is to pick out the relevant facets of the figure and the ground – to pick out that it is the concrete view of the driving events that goes to Paris, but the temporal view that goes to dusk, saying that if the ground is the kind of thing that *could* be the end of the figure then it *is* the end of the figure.

The role of the preposition, then, is to *connect* appropriate views. Anything further that emerges will follow from the nature of the relevant items: Fig. 6 (overleaf) shows the temporal model that emerges from assigning dawn and dusk to be the start and end points of the temporal view of John's drive (this temporal model contains those instants and intervals that are linked to the temporal view of the event itself. Interesting facts that the system has about those instants and intervals are listed below the bar chart[6]).

[5] Fig. 4 concentrates on lower bounds: the axioms dealing with upper bounds are exactly analogous.

[6] the model shown in Fig. 6, like all the other models shown in this paper, was obtained automatically from the given axioms and logical forms.

[7] We are using Reichenbach [1956]'s tri-partitite distinction between event time, speech time, called *now* for short, and 'reference time', which is related to speech time by the tense of the utterance and to event time by its aspect.

#34 (the driving event itself, viewed as an interval)
#33 (reference time[7])
#41 (speech time)
82 (dawn)
#39 (dusk)

INTERVALS
$\#34 = [drive(\#34), extended(\#34), telic(\#34),$
$caused(\#86, \#34), intended(\#86, \#34),$
$\mathbf{to}(\mathbf{\#34}, \mathbf{\#39}), \mathbf{from}(\mathbf{\#34}, \mathbf{\#82}),$
$view(\#34, \#64(\#34)), aspect(simple, \#33, \#34),$
$\theta(\#34, agent, \#86), \theta(\#34, object, \#86)]$
$\#33 = [real(\#33), ends\_before(\#41, \#33), aspect(simple, \#33, \#34)]$
$\#41 = [now(\#41)]$
INSTANTS
$\#82 = [dawn(\#82), from(\#34, \#82)]$
$\#39 = [dusk(\#39), to(\#34, \#39)]$

Figure 6: the drive started at dawn, ended at dusk, and is now over

We can axiomatise *'in'* similarly, as in Fig. 7. To paraphrase the dictionary definition of *'in'* from the start of the paper, X is in Y if each part of X is a part of of the 'envelope' of Y, where the envelope of Y is a set of 'limits of space, time, circumstance'.

$$\forall A \forall B : \{in(A, B)\}$$
$$\quad \forall C : \{view(A, C)\}$$
$$\quad\quad \forall D : \{view(B, D)\}$$
$$\quad\quad\quad \forall E : \{envelope(D, E)$$
$$\quad\quad\quad\quad \& C \sim D\}$$
$$\quad\quad\quad\quad \forall A : \{member(F, C)\}$$
$$\quad\quad\quad\quad\quad member(F, E)$$

Figure 7: If A is in B then every part of A is a part of B

As with Fig. 5, the lexical item *'in'* invokes the more abstract notion of the *envelope* of some view of the ground. Again there is no restriction on what kind of thing this view is. We just have to find compatible views of the figure and the ground, and then the membership relation will do the rest.

Just what the membership relation will do depends on the kinds of sets we are talking about. If every member of one interval I_1 is a member of another, bounded, interval I_2 then, in particular, every member of I_1 is after or equal to the start of I_2 and before or equal to its end. This follows from the axioms in Fig. 4.

For other kinds of sets other conclusions will follow. This is entirely natural – being in a park is quite different from being in January or being in a class, even if in each case *'in'* just means set membership. The consequences of being a member of an ordered set such as a temporal interval are stronger than the consequences of being a member of an unordered set: that's what it means to say that a set is ordered. But the consequences of saying that one bounded interval is a subset of another are simply the consequences that arise from saying that intervals are ordered sets of instants. The temporal relations that arise from this for (2b) are shown in Fig. 8 overleaf.

The general idea, then, is that prepositions pick out compatible views of the figure and ground and relate elements of those views.

Things get a bit more complicated when we consider the examples in (3):

(3) a. I saw him at the station.
b. I saw him at two o'clock.
c. He departed at two o'clock.
d. He arrived at the station at two o'clock.

At first sight, it is tempting to deal with these examples allowing *'at'* to pick out the (physical/temporal/...) *location* of an entity, in the same way that *'to'* picks out its start, and *'in'* picks out its envelope, allowing the properties of different kinds of location to

INTERVALS

$\#1345 = [now(\#1345), ends\_after(\#1345, \#1345)]$

$\#1391 = [month(\#1391), envelope(\#1391, \#1391), in(\#1428, \#1391), named(\#1391, January)]$

$\#1428 = [extended(\#1428), see1(\#1428), \mathbf{in}(\mathbf{\#1428}, \mathbf{\#1391}), intended(\#1385, \#1428), aspect(simple, \#1427, \#1428), \theta(\#1428, agent, \#1385), \theta(\#1428, object, \#1395)]$

$\#1427 = [aspect(simple, \#1427, \#1428)]$

Figure 8: The seeing event is contained in January, and is now over

support different sets of consequences. The envelope and the location, however, seem to be generally very closely linked, and the idea of characterising *'at'* by exploiting the envelope again, rather than introducing yet another slippery term is very attractive.

If we say, as in Fig. 9, that $at(F, G)$ means that F must be a member of the envelope of G then we can capture most of the entailments of (3).

$$\begin{array}{l} \forall A \forall B : \{at(A, B)\} \\ \quad \forall C : \{view(A, C)\} \\ \qquad \forall D : \{view(B, D)\} \\ \qquad\quad \forall E : \{envelope(D, E) \\ \qquad\qquad \& C \sim D\} \\ \qquad\qquad member(D, C) \end{array}$$

Figure 9: If A is at B then B is part of A

The MP in Fig. 9 will derive the appropriate consequences for (3a) and (3b) fairly straightforwardly. The envelope of the physical view of the seeing event include the station (3a), and the envelope of the temporal view of this event (which is extended, and hence includes a number of instants) includes two o'clock (3b).

For (3c) and (3d), we need to think about what the envelopes of arriving and departing events are like. Arriving is instantaneous, in the sense that you know when someone has arrived somewhere (they are at that place and are no longer moving) even if it is difficult to articulate whether there is a first moment at which they arrived, or a last moment before they arrived. We will take it, then, that the temporal envelope of an arriving event contains exactly one moment i such that there is no earlier time when you are there. So if he arrived at two o'clock, then i must be two o'clock, since it is the only member of the set. If we let the physical envelope of an arriving event contain just the destination, as seems perfectly reasonable, then the PP *'at the station'* in (3c) also makes perfectly good sense.

Departing is more complex. The key to a departure is that you are no longer at the starting position, so the temporal envelope of a departing event should contain instants when you are no longer at the start point. The question is: which such instants should it include? It clearly does not include *all* subsequent instants – I am not at this moment still departing from my house to get to work – but is very hard to put a natural bound on when you are no longer departing. There are two reasonable choices: the last moment at which you were at the start point, or the first moment when you were not there. Since the crucial property of departure is that you are no longer at your start point, we assume that the temporal envelope of a departing event is the first moment at which you are not at your setting off place.

This discussion of arriving and departing revolves around notions such as 'the first moment at which you are at your destination' or 'the first moment when you are not at your setting off place'. This is, of course, anathema to any view of time that maps the time line onto the real numbers. *There is, however, no reason to assume that this is the right thing to do*. There are numerous treatments of time that allow models incorporating multiple levels of granularity, where each each level is discrete (thus providing for a coherent notion of 'the next instant') but can be expanded as required to a finer view. Models of this kind make it possible to think of the envelope of an arrival or a departure as a singleton set containing an instant i such that there is no earlier instant where I either am, or more awkwardly am not, at the relevant place without having to answer questions about where I was in between [Liu and Orgun, 1996, Merlo et al., 1999].

3 Conclusions

The work reported here aims to show that we can obtain the consequences of temporal uses of prepositions on the basis of very abstract descriptions of what a range of prepositions mean, together with a very weak logic of instants and intervals (basically that instants are partially ordered, and that intervals are sets of instants which may be bounded and may or may not have least-upper and greatest-lower bounds). Note that we *not* choosing between a number of possible different interpretations. Rather we are using conditional meaning postulates to allow different consequences to emerge depending on nature of the connected items. This fits in with Blackburn et al. [1997]'s observation that the best reading of an utterance is the most informative one, except that we allow the consequences of seeing an object from different points of view to emerge during the construction of a single model. In other words, rather than saying that *'read-in-January'* selects a temporal reading of *'in'*, we allow the fact that *'read'* and *'January'* both have temporal views to add temporal information to the model. If the figure and ground have more than one set of compatible views, then the consequences of all these links are added. The advantages of proceeding this way are that there is no backtracking, since we are not choosing between alternative readings, and that we can allow a range of different relations to emerge within a single model.

This work gives a concrete implementation of ideas discussed by Dowty [1979], who in turn refers to '. . . van Fraassen's notion of *logical space*. There will be as many axes of logical space as there are kinds of measurement. [. . .] *Each axis might have a different mathematical structure according to the discriminations that can be appropriately made in each case*' (p. 126, my italics). Successful implementation depends on two things: (i) you have to have an inference engine which can be used for generating models on the basis of sets of meaning postulates, and (ii) you have to have a set of meaning postulates. There are a number of candidate model generation algorithms – the approach we use is derived from Manthey and Bry [1988]'s 'model generation' approach to theorem proving, extended as described in [Ramsay, 2001] to cover a fine-grained intensional logic, but other authors (e.g. [Baumgartner and Kühn, 2000, Gardent and Konrad, 2000]) have used resolution-based engines for the same task. The development of appropriate sets of MPs is more problematic, since constructing large sets of MPs is a very labour-intensive activity, and it is hard to extract the necessary rules automatically. The models above were produced on the basis of a set of MPs describing the temporal structure ascribed to a number of verb classes, but clearly the more we can say about individual words and concepts the richer the models that can be generated.

The treatment of aspect described deliberately ignores issues relating to tense and aspect, as discussed by [Smith, 1991, ter Meulen, 1995, Moëns and Steedman, 1988], since including these issues would merely have obscured the main point. The aim of the current paper was to bring out the abstract nature of prepositional relations, and to show how the detailed consequences of temporal uses of such terms emerge from combining the view of intervals as sets of instants, the general properties of ordered sets, and the very abstract characterisations of the prepositions themselves. We have therefore *removed* our normal treatment of these issues to keep the presentation simple, though we do normally include them in our logical forms.

References

P Baumgartner and M Kühn. Abducing coreference by model construction. *Journal of Language and Computation*, 1(2), 2000.

P Blackburn, J Bos, M Kohlhase, and H de Nivelle. Inference and computational semantics. In H Bunt and E Thijsse, editors, *3rd International Workshop on Computational Semantics*, pages 5–19, University of Tilburg, 1997.

D R Dowty. *Word Meaning and Montague Grammar*. D. Reidel, Dordrecht, 1979.

C Gardent and K Konrad. Interpreting definites using model generation. *Journal of Language and Computation*, 1(2):215–230, 2000.

H J Keisler. *Foundations of Infinitesimal Calculus*. Prindle, Weber and Schmidt, Boston, Mass., 1976.

K Konrad, H Maier, D Milward, and M Pinkal. An education and research tool for computational semantics. In *Proceedings of the 16th International Conference on Computational Linguistics (COLING-96)*, pages 1098–1102, Copenhagen, 1996.

C Liu and M A Orgun. Dealing with multiple granularity of time in temporal logic programming. *Journal of Symbolic Computation*, 22(5/6):699–720, 1996.

R Manthey and F Bry. Satchmo: a theorem prover in Prolog. In R Lusk and R Overbeek, editors, *Proceedings of the 9th International Conference on Automated Deduction (CADE-9)*, volume 310 of *Lecture Notes in Artificial Intelligence*, pages 415–434, Berlin, 1988. Springer-Verlag.

I Merlo, E Bertino, E Ferrari, and G Guerrini. A temporal object-oriented data model with multiple granularities. In C Dixon and M Fisher, editors, *Proc. 6th International Workshop on Temporal Representation and Reasoning*. Orlando, Fl., 1999.

M Moëns and M Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14(2):15–28, 1988.

J Pustejovsky. The generative lexicon. *Computational Linguistics*, 17(4):409–441, 1991.

A M Ramsay. Theorem proving for untyped constructive λ-calculus: implementation and application. *Logic Journal of the Interest Group in Pure and Applied Logics*, 9(1):89–106, 2001. ISSN ISSN: 1367-0751.

H Reichenbach. *The Direction of Time*. University of California Press, Berkeley, 1956.

C S Smith. *The Parameter of Aspect*. Kluwer Academic Publishers, Dordrecht, 1991.

A G B ter Meulen. *Representing Time in Natural Language: the dynamic interpretation of tense and aspect*. Bradford Books, Cambridge, Mass., 1995.

R Turner. A theory of properties. *Journal of Symbolic Logic*, 52(2):455–472, 1987.

Appendix A: noun-*in*-noun triples from the BNC

```
abnormalities in limb
abnormalities in monkeys
abnormality in structure
abortion in principle
absolutist in tone
absorption in life
abuse in childhood
abuse in schools
abuse in siblings
abuses in debt
acceleration in bedrock
accident in history
accidents in life
accidents in summer
accommodation in advance
accommodation in areas
accommodation in rooms
accomplices in crime
account in protest
account in transport
account in volume
accountability in broadcasting
accountants in turn
accounts in newspapers
accounts in order
accuracy in illustration
achievement in arts
```

achievement in education
achievement in games
achievement in house
achievement in life
achievement in verse
achievements in animals
achievements in chess
achievements in cricket
acid in activity
acid in raindrops
acid in rivers
acid in tea
acidity in rainfall
acids in proteins
acres in coppices
acres in extent
acres in size
acres in south
act in theatres
action in cases
action in cells
action in confidence
action in defence
action in detention
action in privacy
action in sight
action in time
actions in court
actions in fiction
actions in office
activities in day
activities in drama
activities in parts
activities in words
activity in body
activity in flavonoids
activity in man
activity in painting
activity in rats
activity in school
activity in summer
activity in towns