Pages 67–72 in *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, Chicago, August 1991

Generating Expressions Referring to Eventualities

Jon Oberlander[1] and **Robert Dale**[2]
Centre for Cognitive Science
University of Edinburgh
2 Buccleuch Place
Edinburgh EH8 9LW Scotland

Abstract

We note (a) the well-rehearsed linguistic observation that eventualities can be referred to by using either noun phrases or sentences, and (b) the seductive ontological parallels drawn by Bach [1986] between eventualities and individuals. We show how the mechanisms for knowledge representation and referring expression generation in an existing natural language generation system [Dale 1988, 1989] can be easily extended to combine these two insights in the generation of a wide variety of forms of reference to eventualities.

Referring to Eventualities

Most domains consist both of things and events. In order to communicate effectively, it is necessary to be able to generate appropriate references to ordinary physical objects, like chairs, cats and carrots, and to **eventualities**, like being in love, chopping carrots and having two hours of negotiations. A natural language generation system will thus have to encode information about both things and eventualities in its underlying knowledge base (KB) (cf. [Kowalski and Sergot 1986; Moens and Steedman 1987]). Now, consider the following examples:

(1) a. John and Bill discussed the paper.
    b. John and Bill had a discussion about the paper.

(2) a. Caesar died slowly.
    b. Caesar's slow death upset Mary.

The data indicate that very similar information about eventualities can be conveyed via either sentences or noun phrases. In the context of theoretical linguistics, Chomsky [1970] observed that to avoid problems arising from syntactic parallelism, verbs and their derived nominals could be made to share lexical entries.[3] In the context of NL understanding, Dahl *et al* [1987] urge the desirability of returning the same semantic representation when interpreting clauses and their nominalizations. In a generation system, we want to be able to generate both syntactic forms. It is convenient, for instance, to be able to refer anaphorically to an eventuality (introduced by a clause) using a pronoun, as in this example from Schuster [1988:602]:

(3) I want to move a block of text as a unit. How do I do <u>it</u>?

On this basis, we may impose the following condition on a generation system. Desideratum 1: we must select a sufficiently neutral KB format, and provide sufficiently strong mapping rules from KB entities to surface strings to allow the generation of both forms. We urge below that this in turn leads us towards another relevant desideratum.

Eventualities and the Mass–Count Distinction

The space of possible eventualities possesses considerable structure, and this structure has been taxonomised in various ways. Following Vendler [1967], much consideration has been given to the 'aspectual types' of utterances of English sentences (cf. [Mourelatos 1978; Hinrichs 1986; Dowty 1986; Moens and Steedman 1987]). Bach [1986] takes the space to include **states** and **non-states**; in turn, states consist of **dynamic** and **static** states, while non-states consist of **processes** and **events**. Events are then either **protracted** or **momentaneous**; momentaneous events are either **happenings** or **culminations**.

[1]Also of the Human Communication Research Centre at the University of Edinburgh. HCRC is supported by the UK Economic and Social Research Council. The work was also supported in part by the UK Science and Engineering Research Council through project grant GR/G22077. Email: `J.Oberlander@uk.ac.edinburgh`

[2]Also of the Department of Artificial Intelligence at the University of Edinburgh. Email: `R.Dale@uk.ac.edinburgh`

[3]The issues involved in nominalization are broad; some provisos are therefore in order. We here consider only NPs which explicitly refer to eventualities; thus we do not discuss nominals with 'hidden events' in their analyses (cf. [Pustejovsky and Anick 1988]). We do not here consider the relation between tense and aspectual systems and NPs referring to eventualities. Further, although we pursue a relationship between nominal reference and eventualities, there is no direct relation to Partee [1984], where a DRT treatment is given of the thesis that tense is an expression allowing anaphoric reference to *times*.

Putting states to one side, we wish here to focus on Bach's observation of the seductive parallels between the mass–count distinction in nominal systems and the aspectual classification of verbal expressions.[4] Bach exploited Link's [1983] treatment of nominal systems in the context of verbal expressions. In particular, he expounded 'this proportion: events: processes :: things: stuff' [1986:5]. The algebra of events and processes derived has a number of appealing features, and explains why 'temporal mass' indicators, like *twice last night*, occur with some VPs but not others, as well as indicating under what circumstances we can switch from a count-expression to a mass-expression, and *vice versa*.

However, Bach does not discuss a class of expressions which share some of the features of both sides of the analogy, and these are precisely the NPs which refer to eventualities we adverted to above. Further examples include: *war*, *concert*, *discussion* and *the destruction of Carthage by the Romans*. We would argue that just as Bach's interpretation of the mass–count distinction applies to verb phrase reference to eventualities, so too it applies to nominal reference to eventualities; and that the distinction should again be understood in terms of a process–event division. For example:

(4) I had a discussion with Fred.

(5) I had two discussions with Fred.

(6) I had a lot of discussion with Fred.

(7) I had two hours of discussion with Fred.

Here, (4) and (5) represent the use of *discussion* as a count noun, appearing with an indefinite and a number expression respectively; (6) and (7) represent its use as a mass noun, appearing with the quantifier *a lot of* and a measure expression respectively. Lewis's [1983] 'Universal Grinder', which converts a count expression into a mass one, seems to apply rather easily to eventualities, and, following Bach, we may say that we have moved from the discrete event of discussion into the process stuff that composes it. Since grinding and packaging seem to be zero-morphological operations, leaving no trace of their operation in lexical realisations, it is perhaps unsurprising that English should contain many eventuality-denoting words that appear both as count nouns and mass nouns; *killing*, for example, can co-occur both with mass-measures and count-measures.

We therefore wish to place a new condition to our generation system. Desideratum 2: if our generation system is to produce reference to eventualities, it must be able to generate both mass and count eventuality expressions. Given the two desiderata, we want to be able to generate sentential and noun phrase references to eventualities, in the latter case allowing singular, mass or plural expressions. We demonstrate below that this is possible by relatively simple augmentation of the ontological structures and generation mechanisms in an existing natural language generation system [Dale 1988, 1989].

[4] In a fuller analysis, we would argue that states are to be represented via objects possessing temporally-constrained properties.

The Representation of Entities

Dale's [1988, 1989] EPICURE system contains knowledge base structures appropriate for generating reference to count, mass and plural reference to things. The system makes use of a notion of a generalized physical object or **physobj**. This permits a consistent representation of entities irrespective of whether they are viewed as individuals, masses or sets, by representing each as a **knowledge base entity** (KBE) with an appropriate **structure** attribute. To construct a referring expression corresponding to a KBE, we first build a **deep semantic structure** which specifies the semantic content of the noun phrase to be generated. We call this the *recoverable* semantic content, since it consists of just that information the hearer should be able to derive from the corresponding utterance, even if that information is not stated explicitly: in particular, elided elements and instances of *one*-anaphora are represented in the deep semantic structure by their more semantically complete counterparts. From the deep semantic structure, a **surface semantic structure** is then constructed. Unlike the deep semantic structure, this closely matches the syntactic structure of the resulting noun phrase, and is suitable for passing directly to a PATR-like unification grammar. It is at the level of surface semantic structure that processes such as elision and *one*-anaphora take place. The transitions from KBE to deep semantic structure (the issue of content determination) and from deep semantic structure to surface semantic structure (the choice of anaphoric strategies and broad linguistic choice) are performed by means of independent sets of **mapping rules** which rewrite the appropriate structures from one level into those of the other.[5]

Dale's [1988] system encoded only rudimentary information about eventualities, although still in the same general form as the encoding of physobjs. Eventualities were treated simply as operators which change the properties of objects. Hence (as appropriate to the sublanguage chosen for the system), eventualities appeared as (imperative) sentences, such as *peel and slice the onions* or *simmer the soup*. Effectively, they were treated, in accordance with situation-calculus based planning, as transitions from global state to global state. However, by representing eventualities as the (near) equals of physical objects, we can satisfy both of our desiderata. Below, we extend the parallels between the two types of entities.

[5] There is insufficient space in the present paper to provide examples of these two levels of representation: see Dale [1988, 1989] for examples.

Common Features of Physobjs and Eventualities

The domain in which the system operates consists of a finite set of entities. Each entity is represented by a distinct symbolic constant called its **index**.[6] An entity is either a **physobj** or an **eventuality**, where these are defined as follows: a **physobj** is any (not necessarily contiguous) collection of contiguous regions of space occupied by matter; a **eventuality** is any (not necessarily contiguous) collection of contiguous regions of time occupied by process stuff. Physobjs have spatial parts; if a physobj can be decomposed into parts, those parts will be physobjs. Eventualities have temporal parts; if an eventuality can be decomposed into parts, those parts will be eventualities. Every entity has, in addition to its index and type, a **specification**. The specification of an entity provides all the information known to the system about that entity.

An entity may have as part of its specification a **location**. In a physobj, this will typically be its spatial location; in an eventuality, this will typically be its temporal location, given via begin and end points (which may coincide). In both cases, the system need not know any value for the location. However, every entity must have, as part of its specification, a **substance**. In the case of physobjs, the substance is the kind of matter from which the object is made; in the case of eventualities, the substance is the kind of process stuff that makes up the eventuality. There is a finite but extensible set of substances and process stuffs represented within the system by means of symbolic constants, which are organised into a taxonomic graph structure.

As a further part of its specification, every entity has a **structure**. This corresponds to the way in which the entity is perceived. Whether the entity is a physobj or an eventuality, its structure is either **individual**, **set** or **mass**. An eventuality is treated on a par with a physobj. Whether it is referred to as a process or an event is determined by whether its structure is mass or individual. Any entity may have any number of additional properties specified as part of its specification, where those additional properties are drawn from a finite but extensible set of properties. These properties are binary valued features with + and − being the possible values. For example, for either *Caesar died slowly* or *Caesar's slow death*, we would expect the KB specification to include [slow = +] as part.

If an entity has structure **individual**, then it also has a **packaging** as part of its specification. A packaging is a tuple consisting of a **shape** and a **size**. The possible values of **shape** and **size** are each drawn from two finite but extensible sets: one in each case relevant to physobjs, the other to eventualities. If an entity is a **mass**, it may or may not have a **quantity**. If the entity is a **set**, it may have either a **cardinality**, a **quantity**, an explicit list of **constituents**, or none of these (in which case it can only be described by a bare plural).

Cardinality is specified as a numerical value, or a **range** of numerical values where the range consists of a **lowerlimit** and an **upperlimit**. The quantity of an entity is specified as a tuple consisting of some **unit** of of measurement, and a number; or a range consisting of two such tuples, one for the lower limit and one for the upper. The possible values for the **unit** feature are drawn from finite but extensible sets appropriate to physobjs and eventualities. If an entity whose structure is **set** has the feature **constituents**, the value of this feature is a list of symbolic constants which are the indices of other entities. If an entity whose structure is **set** does not have an explicit list of constituents, then it will have an **elements** feature, which provides a **specification** which is true of all the elements of the entity.

$$\left[\begin{array}{l} \mathit{index} = x_0 \\ \mathit{stage} = s_0 \\ \mathit{spec} = \left[\begin{array}{l} \mathit{structure} = \mathit{set} \\ \mathit{quantity} = \left[\begin{array}{l} \mathit{unit} = \mathit{kilogram} \\ \mathit{number} = 2 \end{array}\right] \\ \mathit{elt} = \left[\begin{array}{l} \mathit{structure} = \mathit{individual} \\ \mathit{substance} = \mathit{onion} \\ \mathit{pack} = \left[\begin{array}{l} \mathit{shape} = \mathit{onion} \\ \mathit{size} = \mathit{regular} \end{array}\right] \end{array}\right] \end{array}\right] \end{array}\right]$$

Figure 1: The knowledge base entity corresponding to *two kilograms of onions*

$$\left[\begin{array}{l} \mathit{index} = x_0 \\ \mathit{stage} = s_1 \\ \mathit{spec} = \left[\begin{array}{l} \mathit{structure} = \mathit{mass} \\ \mathit{quantity} = \left[\begin{array}{l} \mathit{unit} = \mathit{kilogram} \\ \mathit{number} = 2 \end{array}\right] \\ \mathit{substance} = \mathit{onion} \\ \mathit{chopped} = + \end{array}\right] \end{array}\right]$$

Figure 2: The knowledge base entity corresponding to *two kilograms of chopped onion*

Special Features of Physobjs

Since a physobj can change its properties through time, we let properties hold true of physobjs at **stages**.[7]

[6] We will not discuss here the obvious benefits of this generalisation for anaphoric reference to eventualities.

[7] This was termed 'state' in Dale [1988, 1989].

The invariant properties (**substance** and **quantity**) hold true of an object at all stages; all other properties of objects may change. In general, any change in a physobj's properties is mediated by the occurrence of an eventuality. Figures 1 and 2 show two different KBEs, which might be realized by the noun phrases *two kilograms of onions* and *two kilograms of chopped onion* respectively; these correspond to the same physobj at different stages.

Special Features of Eventualities

As stated above, eventualities have temporal parts, while physobjs have spatial parts. Just as we don't always know the spatial location of a physobj, we don't always know the temporal location of an eventuality. But it is an option: for some eventualities the **location** will take as value a **begin** time and/or an **end** time. Then, an eventuality will have, as part of its specification, a finite but extensible set of **participants**. These are collected into two substructures, corresponding to the **in** participants and the **out** participants: the **in** participants are those participants which exist in the **begin** location of the eventuality, and the **out** participants are those participants which exist in the **end** location of the eventuality. This allows us to model eventualities which result in the destruction of existing entities or the creation of new ones.[8] The participants in each set will include those entities known by the system to fill various participant roles within the eventuality, such as **agent**, **patient** or **theme**. Each role has as its value a symbolic constant that corresponds to an entity in the domain. Note that we can capture the notion of causation in this way: an eventuality may contain other eventualities amongst its participants.

Generating Reference to Eventualities

Sentences vs NPs

Consider the KBE in Figure 3, where c is the index of the KBE corresponding to the individual named *Caesar*; this represents the event of Caesar dying. The choice of whether this is to be realized as a sentence or as a noun phrase is decided by the mapping rules that determine the content of the deep semantic structures underlying a text. If a sentence realisation is chosen, then, given an appropriate speech time, the sentence will be something like

(8) Caesar died slowly.

Now consider the KBE in Figure 4, where m is the index of the KBE corresponding to the individual named *Mary*, and the agent of the event is the eventuality shown in Figure 3; this represents the event of Mary being upset by Caesar's death. This allows us various possible realizations:

(9) Caesar died slowly. It upset Mary.

[8] Noun phrase reference to eventualities will not in general require this; but *Caesar's death* may well do.

$$\begin{bmatrix} \mathit{index} = e_0 \\ \mathit{location} = \begin{bmatrix} \mathit{begin} = t_1 \\ \mathit{end} = t_2 \end{bmatrix} \\ \mathit{spec} = \begin{bmatrix} \mathit{structure} = \mathit{individual} \\ \mathit{substance} = \mathit{dying} \\ \mathit{particips} = \begin{bmatrix} \mathit{in} = \begin{bmatrix} \mathit{agent} = c \end{bmatrix} \end{bmatrix} \\ \mathit{slow} = + \end{bmatrix} \end{bmatrix}$$

Figure 3: The KBE corresponding to *Caesar dies slowly*

$$\begin{bmatrix} \mathit{index} = e_1 \\ \mathit{location} = \begin{bmatrix} \mathit{begin} = t_6 \\ \mathit{end} = t_7 \end{bmatrix} \\ \mathit{spec} = \begin{bmatrix} \mathit{structure} = \mathit{individual} \\ \mathit{substance} = \mathit{upsetting} \\ \mathit{particips} = \begin{bmatrix} \mathit{in} = \begin{bmatrix} \mathit{agent} = e_0 \\ \mathit{affected} = m \end{bmatrix} \end{bmatrix} \end{bmatrix} \end{bmatrix}$$

Figure 4: Mary's being upset by Caesar's death

(10) The fact that Caesar died slowly upset Mary.

(11) Caesar's slow death upset Mary.

(12) The slow death of Caesar upset Mary.

(13) Mary was upset by Caesar's slow death.

Here, the alternation between (9) and (10) is determined by the specific mapping rules chosen to construct deep semantic structures from the underlying KBEs (a question of 'chunking' and content determination);[9] and the alternation between (11), (12) and (13) is determined by the specific mapping rules chosen to construct surface semantic structures from the corresponding deep semantic structures (a question of the choice of gross syntactic structure and clause-internal information structure).

This characterisation misses a crucial element, of course: *how* do we choose between these different possible realisations? In particular, it is not yet clear what will lead a speaker to choose an NP rather than an S as the resource for realising a particular eventuality specification. Two kinds of considerations might be brought to bear. First, where the available information is less than complete, syntactic constraints may rule out certain realisations. Facts about obligatory arguments roles are particularly relevant here. For example, Rappaport [1983:137] observes that 'Bach's generalization

[9] An alternative approach to this question, which partitions the task of language generation slightly differently, is described in Meteer [1990].

$$
\left[\begin{array}{l}
\mathit{index} = e_2 \\
\mathit{spec} = \left[\begin{array}{l}
\mathit{structure} = \mathit{individual} \\
\mathit{substance} = \mathit{discussing} \\
\mathit{particips} = \left[\mathit{in} = [\mathit{agent} = \langle s, f\rangle]\right]
\end{array}\right]
\end{array}\right]
$$

Figure 5: I had a discussion with Fred

$$
\left[\begin{array}{l}
\mathit{index} = e_3 \\
\mathit{spec} = \left[\begin{array}{l}
\mathit{structure} = \mathit{set} \\
\mathit{cardinality} = 2 \\
\mathit{elt} = \left[\begin{array}{l}
\mathit{structure} = \mathit{individual} \\
\mathit{substance} = \mathit{discussing} \\
\mathit{particips} = \left[\mathit{in} = [\mathit{agent} = \langle s, f\rangle]\right]
\end{array}\right]
\end{array}\right]
\end{array}\right]
$$

Figure 6: I had two discussions with Fred

$$
\left[\begin{array}{l}
\mathit{index} = e_4 \\
\mathit{spec} = \left[\begin{array}{l}
\mathit{structure} = \mathit{mass} \\
\mathit{substance} = \mathit{discussing} \\
\mathit{quantity} = \mathit{extended\_amount} \\
\mathit{particips} = \left[\mathit{in} = [\mathit{agent} = \langle s, f\rangle]\right]
\end{array}\right]
\end{array}\right]
$$

Figure 7: I had a lot of discussion with Fred

$$
\left[\begin{array}{l}
\mathit{index} = e_5 \\
\mathit{spec} = \left[\begin{array}{l}
\mathit{structure} = \mathit{mass} \\
\mathit{quantity} = \left[\begin{array}{l}
\mathit{unit} = \mathit{hour} \\
\mathit{number} = 2
\end{array}\right] \\
\mathit{substance} = \mathit{discussing} \\
\mathit{particips} = \left[\mathit{in} = [\mathit{agent} = \langle s, f\rangle]\right]
\end{array}\right]
\end{array}\right]
$$

Figure 8: I had two hours of discussion with Fred

[Bresnan 1982:418], that verbs of obligatory object control may never be detransitivized, does not hold for the corresponding nominals'. In the case of interpretation, the PUNDIT system [Dahl *et al* 1987:137] analyses both clauses and nominalizations; analysis fails in the case of a clause with an unfilled obligatory role (such as *John discussed with Bill*), but succeeds with the nominalization (*John and Bill's discussion*). The moral for generation should be clear: where the system lacks critical information about participants in an eventuality, an NP should be generated in preference to a clause.

The second kind of consideration concerns issues that might be regarded as rather more stylistic in nature, although we would prefer to follow Scott and Souza [1990:47] in viewing these constraints as being cognitive rather than aesthetic. In particular, several of Scott and Souza's heuristics for generating effective texts in the framework of Rhetorical Structure Theory (RST) may help here. For example, they suggest that a generator (a) should attempt to make a single sentence out of every rhetorical relation [1990:54], and (b) should prefer syntactically simple embeddings over more complex ones [1990:60]. Taken together, in appropriate rhetorical circumstances these heuristics support the use of an NP reference to an eventuality embedded within a simple S. This would be preferred to an S reference to that same eventuality, either embedded within a complex S or within a multi-sentential structure; on these grounds, (11) is preferable to (9).

One might wish to extend such a cognitively-oriented account to claim that, in general, an NP is to be preferred to an S on the grounds of efficiency, since NPs offer rather convenient reduced forms, such as pronouns, for referring to previously mentioned entities. But such a claim cannot be supported: Ss are equally flexible, since they allow VP ellipsis. It should also be borne in mind that the use of more concise forms carries a price. Just as *an antique lover* is ambiguous as to whether we mean a lover of antiques or a very old lover, using an NP in preference to an S may introduce an ambiguity in interpretation: *Sue Ellen's betrayal* could refer to Sue Ellen betraying J. R. or to J. R. betraying Sue Ellen.

Considerations of efficiency are more likely to choose between NPs than to choose between an NP and an S; it is to the choice between NPs that we now turn.

Mass vs Count Descriptions

Consider examples (4)–(7), repeated below as (14)–(17):

(14) I had a discussion with Fred.

(15) I had two discussions with Fred.

(16) I had a lot of discussion with Fred.

(17) I had two hours of discussion with Fred.

The KBEs underlying these eventualities are shown in Figures 5–8 (with s corresponding to the speaker and f corresponding to *Fred*).

In each case, the structure of the underlying KBE determines, as in the case of descriptions of objects, whether a mass or count noun phase is to be used in the realization. Note, however, that we have four distinct KBEs here; perhaps at least some of these ought to be represented by the same underlying object. Arguably, one and the same eventuality underlies (14) and (17) above; but we have represented these by distinct KBEs. Recall that our KBEs correspond to the way in which entities are *perceived*; this same issue arises, although less obviously so, in the case of physobjs, since we can often describe one and the same object, quite apart from any physical transformations that might be applied to it, by means of count or mass descriptions (again, cf. Lewis's [1983] 'Universal Grinder' and 'Universal Packager'). Ultimately, we require some means of tying different KBEs together as different **perspectives** on the same entity. We are then faced with a question analogous to that raised in the previous section: how we choose between the different perspectives? There are clear differences amongst the candidates in terms of informational content and specificity, and so the criteria would appear to be more a matter of what to say rather than how to say it. Such issues are beyond the scope of the present paper, but must ultimately be addressed, since strategic decisions have an impact on subsequent tactical decisions; indeed, in some models of the generation process, the two mutually constrain each other.

Conclusions and Future Directions

Generation of appropriate reference to eventualities in discourse will involve a number of interacting factors. Whether we refer to an eventuality with a mass or count noun (in NP reference), and the aspectual class of a reference (in sentential reference), is currently a matter decided within the KB, where the entity's structure is encoded as **individual**, **set** or **mass**. We have suggested here that it is a relatively simple matter to *permit* the generation of either sentences or NPs referring to eventualities. This leaves open the question of 'chunking' information: issues that have been touched upon in the use of RST relations in generation are relevant here (cf. [Hovy 1990; Scott and Souza 1990]).

The major area for future study is the generation of appropriate reference to eventualities in connected discourse. The existing generation system is geared toward the production of reference to objects in multi-sentence texts; hence we would argue that the extensions to it described here should allow a natural path into the area. A current limitation of our approach is that the nominals are de-coupled from the generation of tense and aspectual class; nonetheless, the KB entities contain the appropriate attributes, and we hope to remedy this limitation.

Acknowledgments

We thank Judy Delin, Ewan Klein, Alex Lascarides, and members of the Centre for Cognitive Science's Discourse Workshop for helpful discussions.

References

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

Bresnan, J [1982] Control and complementation. *Linguistic Inquiry*, **13**, 343–434.

Chomsky, N [1970] Remarks on nominalization. In R Jacobs and P Rosenbaum (eds) *Readings in English transformational grammar*, pp184–221. Waltham, Mass.: Ginn and Co.

Dahl, D, Palmer, M and Passonneau, R [1987] Nominalizations in PUNDIT. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, Stanford, Ca., July, 1987, pp131–139.

Dale, R [1988] Generating Referring Expressions in a Domain of Objects and Processes. PhD Thesis, Centre for Cognitive Science, University of Edinburgh.

Dale, R [1989] Cooking up Referring Expressions. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, Vancouver BC, pp68–75.

Dowty, D [1986] The effects of aspectual class on the temporal structure of discourse: semantics or pragmatics? *Linguistics and Philosophy*, **9**, 37–61.

Hinrichs, E [1986] Temporal anaphora in discourses of English. *Linguistics and Philosophy*, **9**, 63–82.

Hovy, E H [1990] Unresolved Issues in Paragraph Planning. In R Dale, C Mellish and M Zock (eds) *Current Research in Natural Language Generation*, pp17–45. London: Academic Press.

Kowalski, R and Sergot, M [1986] A logic-based calculus of events. *New Generation Computing*, **4**, 67–95.

Lewis, D [1983] Scorekeeping in a Language Game. Chapter 13 in *Philosophical Papers*, Volume 1. Oxford: Oxford University Press.

Link, G [1983] The logical analysis of plurals and mass terms: a lattice-theoretical approach. In R Bauerle, C Schwarze and A von Stechow (eds) *Meaning, Use and Interpretation of Language*. Berlin: de Gruyter.

Meteer, M [1990] The "Generation Gap": The Problem of Expressibility in Text Planning. COINS Technical Report 90–04, Computer and Information Science Department, University of Massachusetts, February 1990.

Moens, M and Steedman, M [1987] Temporal Ontology in Natural Language. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, Stanford, Ca., 6–9 July, 1987, pp1–7.

Mourelatos, A P D [1978] Events, Processes and States. *Linguistics and Philosophy*, **2**, 415–434.

Partee, B H [1984] Nominal and temporal anaphora. *Linguistics and Philosophy*, **7**, 243–286.

Pustejovsky, J and Anick, P G [1988] On the semantic interpretation of nominals. In *Proceedings of COLING–88*, Budapest, Hungary, 22–27 August, 1988, pp518–523.

Rappaport, M [1983] On the Nature of Derived Nominals. In L Levin, M Rappaport and A Zaenen (eds) *Papers in Lexical Functional Grammar*. Bloomington, Indiana: Indiana University Linguistics Club.

Schuster, E [1988] Anaphoric reference to events and actions: a representation and its advantages. In *Proceedings of COLING–88*, Budapest, Hungary, 22–27 August, 1988, pp602–607.

Scott, D R and Souza, C S [1990] Getting the Message Across in RST-based Text Generation. In R Dale, C Mellish and M Zock (eds) *Current Research in Natural Language Generation*. London: Academic Press.

Vendler, Z [1967] Verbs and times. Chapter 4 in *Linguistics in Philosophy*. Ithaca, NY: Cornell University Press.

Vendler, Z [1968] *Adjectives and Nominalization*. The Hague: Mouton.