

# *Developing a visual temporal modeller: applying an extensible nlp system to support learners' understanding of tense and aspect in English*

JOHN KERINS

*Department of Computer Science and Information Systems, University of Chester,  
Parkgate Road, Chester CH1 4BJ, UK (e-mail: j.kerins@chester.ac.uk)*

ALLAN RAMSAY

*School of Computer Science, University of Manchester, Manchester M13 9PL, UK  
(e-mail: Allan.Ramsay@manchester.ac.uk)*

---

## Abstract

This paper reports on the development of a prototype tool which shows how learners can be helped to reflect upon the accuracy of their writing. Analysis of samples of freely written texts by intermediate and advanced learners of English as a foreign language (EFL) showed evidence of weakness in the use of tense and aspect. Computational discourse modelling techniques were applied to the data to generate semantic models of fragments of the narratives with particular focus on their temporal structure. These models have been converted into dynamic graphical representations of the temporal relationships between discourse events as the narratives are written. The system also provides access to the ontology devised to model individual events and this offers learners insights into the events' semantic properties. These techniques provide the basis for a stimulating learning tool capable of capturing key elements of written narratives, and prompting learners' awareness of language use, particularly tense and aspect.

Keywords: Computational discourse modelling, tense and aspect in English, graphical representation of learner narratives, language awareness, learner autonomy

---

## 1 Introduction

We consider how Natural Language Processing (NLP) techniques can be applied to CALL to extend the range of software resources that facilitate language awareness, reflection and increased autonomy in learners. We place our development within the context of NLP-based CALL. We evaluate examples from the learner data that motivated our application, and explain how discourse models were constructed to capture the temporal entailments of samples of these data.

We analysed a sample of scripts selected from different groups of intermediate and advanced learners of English as a foreign language which revealed a number of

difficulties in handling tense and aspect. We show how computational discourse modelling techniques can be used to capture and portray the temporal entailments of fragments of the data. Learners find the forms that express temporal relationships challenging, and they often produce expressions which are at odds with their intentions. It is almost impossible to construct a computational tool which will ‘diagnose’ these problems, because the relationships that the texts do express are not impossible, and hence cannot easily be detected as ‘errors’. What we can do is provide learners with a visual representation of the relationships between events that are encoded by what they have said, and hence help them reflect on their use of these forms.

We exploit the machinery provided by an existing NLP system to construct graphical representations of events and the relationships between them to help learners visualise the temporal structure of what they have written. We present the underlying ontology and demonstrate how specific events can be defined in terms of the properties of a set of generic event types to expand this representational framework. We extend the formalism by developing rules governing the relationships between events in a discourse based upon their tense and aspect properties. Finally, we interpret the contents of the discourse and convert them into a dynamic graphical representation showing the relationships between events as they are generated. This depiction operates at three levels: the system displays stylised images of the events and their temporal orientation with respect to each other; individual events in this display conform to standardised colour and shade patterns to reinforce their event type; by activating an individual event, users can access a model of its internal properties. We show how different uses of tense and aspect generate different graphical configurations.

Discourse models are depicted as a series of shaded blocks denoting individual events stretching from left to right reflecting the timeline of the discourse as it evolves. Speech time is represented by a vertical line. Information to the left of the line indicates events that took place in the past. The example in Figure 1 illustrates how the following discourse, which is derived from the learner data, is interpreted:

*I visited England for the first time.  
My parents had decided they would travel to England.  
I saw London.  
I also saw more cities.*



Fig. 1. Sample discourse

Our models of samples of learner data suggest that meaningful representations of the temporal entailments of discourses can be constructed. This dynamic visual modelling tool has the potential to draw learners' attention to their use of written English and to reinforce their understanding of temporal expression in English. Significant scope remains to expand the range and flexibility of the system. This could potentially be achieved by extending the lexicon and by integrating more powerful logical axioms to handle additional discourse relations such as causality or contextual knowledge, for example.

## **2 Applying NLP techniques to CALL**

Our work draws upon principles that underpin NLP, which is concerned with how computers can best analyse, store, sort and search language (Nerbonne, 2005). Nerbonne acknowledges that NLP techniques are absent from many CALL applications but he identifies a number of successful technologies: concordancing, text alignment, speech recognition and synthesising, morphological and syntactic processing, which have been used to illustrate linguistic structure and to help with error correction and making language more comprehensible. The German on-line dictionary (LEO GmbH, 2011) is an example of how some of these techniques can provide learners with valuable support. Nerbonne also acknowledges the limitations of NLP's ability to handle complex, ambiguous texts.

Researchers have sought to apply NLP techniques to language learning in a number of contexts. Much of this activity has taken place in universities where learners can be expected to be reasonably proficient in the target language. Examples of grammatical parser-based approaches are used to detect errors in user input and to model learning processes, as we discuss here, and our work draws upon these principles. There is evidence that these systems share the wider CALL aims of developing an interactive learning environment supporting awareness, reflection and learner autonomy.

Vandeventer Faltin (2003) reports on the integration of a number of NLP technologies in a tool designed to facilitate grammatical construction, error analysis and French verb conjugation. The system aims to support communicative language learning tasks (Beatty, 2003) and to enable learners to identify and correct their own errors. Despite the sophistication of the software, Vandeventer Faltin stresses the limitations of the tool but points out its value in encouraging learners to reflect on the target language and its structure. Sanders (1991) is concerned with providing helpful error messages to student users of a parser-based writing aid. While the parser detects grammatical errors in limited domains, the problem of identifying underlying semantic content poses problems when constructing learner feedback.

More recent developments in Intelligent CALL (ICALL) aim to link grammatical formalisms to principles of Second Language Acquisition (Schulze and Penner, 2008). The authors apply Construction Grammar (Fried and Oestman, 2004) to build a framework within which they can develop measures of learner text complexity in order to model language learning processes and improve the individualisation of ICAL systems.

The aim to individualise systems and to actively engage learners by facilitating noticing and raising language awareness underpins the development of open learning models. These techniques, which give learners responsibility for learning (Xu and

Bull, 2010; Bull, Dong *et al.*, 2008), are believed to support language acquisition, (Rutherford and Sharwood-Smith, 1985; Schmidt, 1990). A similar learner modelling system sees the software as a collaborator (Dimitrova, 2003). Users can inspect the model of their beliefs about domain knowledge while the system tries to discover possible reasons for erroneous learner knowledge.

The idea of reflection helps awareness and the promotion of autonomous learning (Kohonen, 1992). Bull (1997) demonstrates that CALL software can be deployed to help raise learners' awareness of their learning style and to interactively prompt alternative styles to be adopted. The aim is to develop a greater awareness of the range of strategies available, and to help students discover strategies which will benefit them. The critical feature is to facilitate autonomy in the learner.

While the systems discussed so far, whether NLP-based or not, have featured purpose-built CALL software, CALL embraces a much wider range of information and communications resources (Warschauer, 1996). Nevertheless, common themes which contribute to language development such as noticing, raising awareness and remembering, for example, feature in many approaches. This is apparent in the environment developed by Dekhinet (2008) who claims that non-native speakers can become active, self-regulated learners through on-line written interactions with native-speaking tutors. Jaen and Basanta (2009) enumerate the benefits of using DVD-based multimodal texts to raise learners' awareness of contextual factors in conveying meaning in developing their understanding of spontaneous speech behaviour. This authentic input is used to direct attention to noticing and storing forms and meanings in context (Long, 1991), which is aimed, in this case, at enhancing conversational proficiency.

The work reported in this paper utilises NLP techniques to construct models of learners' written discourses. As with Schulze and Penner (2008) and Xu and Bull (2010), the models are designed to be processed in real-time to provide feedback to learners. Our experimental approach is at an early stage of development but its aims are consistent with the attempts to build interactive environments which prompt learners to question their own language use and to become aware of discrepancies between their own and native speaker constructions. This is consistent with the idea of learners noticing the distance between the target language and their own (Schmidt and Frota, 1986).

The use of graphical representation seems an intuitive means of capturing the temporal entailments of a discourse. There has been much work into the effectiveness of graphical representation to capture and convey meaning (Shimojima, 1999; Sowa, 1984). Dimitrova, Brna *et al.* (2002) report on learners' ability to communicate using conceptual graphs. Significantly, their research indicates that learners are generally adept at interpreting graphical representations. Anderson-Hsieh (1994) comments on the use of visual feedback for teaching phonetics and concludes that the role of the teacher is important in selecting and presenting appropriate examples. While CALL environments are largely designed for independent use, it is likely that initial training is necessary to familiarise learners with the selected medium. Clearly, there is a need for empirical analysis of users' responses to our system, which will inform decisions about the effectiveness of the interface, and suggestions about how it might be improved.

### 3 Analysing learner data

Data were gathered from native German speaking learners of English. In each case, students were asked to write freely about their first impressions of Britain. One group of learners was studying at Cambridge Proficiency Level 6; a second group was studying English at a private language school; two learners were advanced students who were on an undergraduate exchange programme at a college of Higher Education.

A common theme emerging from the data is a misrepresentation of temporal relations within a number of the narratives. These difficulties are related to the use of tense and aspect in English, and show evidence that the sequencing of events is inconsistent. The discrepancies are more pronounced in the less experienced learners, although there is evidence of subtle misrepresentation in the more accomplished learners.

Interestingly, empirical work on a sample of German learners of English focuses specifically on the acquisition of tense and aspect and concludes that evidence suggests avoidance of a number of compound tenses denoting temporal relations (Duerich, 2005). The findings are evaluated with respect to the view that the non-native speakers have particular difficulty with expression of aspectual relations (Dorfmueller-Karpusa, 1985). It is, clearly, very difficult to detect avoidance strategies in user input: by definition, such strategies lead to the absence of errorful text. If we want to help students acquire the ability to use temporal expressions, we have to work with what they do write, rather than looking for what they do not. The goal of the work here, then, is to reveal the temporal relationships that are encoded by the learner's input text in an unthreatening environment which will let them reflect on their writing.

#### 3.1 Sample data

Samples of the kinds of error appearing in the texts illustrate the motivation for the focus of the learner modelling software:

*I was so impressed and fascinated by the city that I decided to follow my mother and to go to London as an au pair as well after I've passed my exams. This was 8 years ago but I have held the promise I had made to myself.*

The author identifies a period in the past at which a decision to go to London is made. The proposed time of the visit is at some time in the future when the author has passed her exams. The next sentence places us in the present, at which point we infer from the reference to fulfilling her promise to herself, that she has already passed her exams. The final past perfective clause possibly refers to the fact that the promise was made before, and was contingent upon, passing her exams. This interpretation would have been easier to recover if the past perfect had been used in reference to passing her exams: ... *after I had passed my exams.*

The following example raises subtle issues about aspectual use of English:

*I must say my rather positive impressions of Britain and its people have changed. I have never thought there are so many differences between two cultures that both are settled in Europe. But it is true.*

The first sentence reports a change in attitude but the use of the present perfect in the next sentence seems to imply that the earlier attitude persists.

The following example also features aspectual infelicity:

*I don't know why I've chosen London for my stay as an au pair because I hated the city immediately when I came here in 1994.*

The use of the present perfect seems to imply that the stay in London has not yet begun, or possibly, has only just begun, yet within the context of the discourse this is not the case. The use of the simple past here would give a clearer description of events.

A temporal reference is at issue in the following:

*It took only a week or so and I realised that I don't really like London, it is a big city, multicultural but I more like my Berlin.*

Ordinarily, having begun events in the past, it would probably be expected that the description would remain consistent. The author's realisation that she disliked London occurred in the past; it might have been preferable to introduce an additional clause confirming that she continues to dislike it. Keeping the whole phrase in the past would not be contentious.

An analysis of the following text shows evidence of accurate use of English tense and aspect:

*I think I loved England from the very first time we arrived on the island by ferry. I had known England before as a land of fish and chips, and of course I had to try it. Very greasy, but delicious. Also this was when I started liking cheese. We tried lots of different types of cheese and even though the bread was horribly soft and full of sugar I ate lots.*

The author successfully integrates a number of temporal features into this particular fragment. The notion of the beginning of a state of loving is located in the past and is co-ordinated with a past arrival. From this point, the author successfully refers to an earlier state. In fact, the author selects the point of arrival as pivotal in this part of the narrative. The state of knowing about the concept of fish and chips is about to be modified by experiencing them for real. The point of arrival is implicitly widened to incorporate a general time around which the author first came to England. It is associated with the start of a state of liking cheese, which persists. The imperfective is used accurately to convey this.

The difficulty for learners seems to lie in the way in which English captures the relationships between events that interact over time. These relationships can be interpreted in terms of the properties that different classes of event entail and the ways in which these properties interact with aspectual representation in English. Ordinarily, a state cannot be referred to using the imperfective, for example: *he is knowing English*. The learner data demonstrate inconsistency and inaccuracy in the use of temporal representation. When seen in context, a number of temporal constructions, such as the examples discussed above, do not seem to convey the temporal relationships that the reader is expecting.

The principal aim of our development so far is to demonstrate that fragments of freely written texts can be modelled and presented to users with key semantic features highlighted, although we are far from being able to model the full complexity

and subtlety of the discourses analysed. Nevertheless, this pared-down structural representation has the capacity to raise language awareness and to prompt questions about understanding where the output differs from learners' expectations.

We need to extend the grammatical and lexical range of this early prototype. More particularly, we need to test the system thoroughly with target users to assess the benefits and limitations of the approach. We anticipate that there may be significant scope for improvement but we believe that the treatment has the potential to enhance learners' understanding of what they have written and to contribute to the range of tools and techniques that helps them reflect on their knowledge and to actively engage in further learning.

#### 4 Modelling natural language semantics

In this section we introduce key approaches to generating abstract models of linguistic utterances. We focus particularly on temporal representation. Given the difficulties that learners have with this aspect of language, providing them with an externalised representation of the temporal structure of their narratives is likely to help them see the difference between what they said and what they meant, and hence to reflect on their use of temporal expressions.

In building computational discourse models we construct a representation of the entailments of what has been said. Implicit in this process is the assumption that when individuals communicate they are each attuned to various contextual phenomena that the language denotes. This background knowledge supports a much wider range of inferences than are explicit in a given utterance. A sentence such as: *The soldiers withdrew from their advanced position* assumes knowledge about soldiers, their behaviour, and the motivation for their actions.

In discourse modelling, we aim to implement a representational framework that reflects our semantic interpretation of the world. This must support simple entailments such as the fact that reference to a single entity cannot be followed by reference to a plural entity with any sense that they are the same thing. Equally, we might wish to infer that reference to a woman signifies that we are referring to a human, and we incorporate relevant axioms capturing background knowledge to support these kinds of inference.

While our models are abstract structures which pick out salient semantic features of a discourse that signify entities and relations that have a parallel in the real world, there is some debate as to whether we construct similar mental models from a psychological perspective when we communicate (Johnson-Laird, 1983; 2006). These issues inform some of the theoretical work associated with the approaches to NLP which underpin the system upon which the material presented in this paper is based; the distinctions between model-theoretic semantics, and approaches that also address ways that belief and knowledge might be modelled, are presented by Kamp and Reyle (1993) and Barwise and Perry (1983) respectively.

The system exploits the principles of unification-based grammar (Pollard and Sag, 1987; 1994), and lexical inheritance (Pustejovsky, 1995), which are augmented by dynamic predicate logic to capture discourse structure (Groenendijk and Stokhof, 1991). The mechanism is also informed by work in Centering Theory

(Grosz, Joshi *et al.*, 1995); Rhetorical Structure Theory (Mann and Thompson, 1988; Mann, 2010) and Rhetorical Parsing (Marcu, 1997) in maintaining dynamic models of evolving discourses.

#### 4.1 Generating a discourse model

The NLP system maintains coherent discourse state models that capture the temporal properties of a series of events and of the entities that participate in them. The entailments of what is encoded in the words that are input into the system are captured in the logical form of the discourse. This representation is combined with the background knowledge, which is encoded within the system, to produce a model. If we examine a brief series of discourse events we can see how a dynamic model of the discourse is maintained as new events are added. Although we are able to generate a richly detailed representation of what we can infer from the model, we can also appreciate that the encoding is not particularly easy to read. We do include the details of the models that are generated by the system in Figures 2, 3 and 4, but these are included here in order to give the reader a clear idea of the information that the system has access to. Exposing these textual representations to a learner would be counter-productive. In Section 9 we show the more palatable graphical representations that we present to learners.

The system generates the model in Figure 2 for the simple input sentence: *The student left London*. The single event, *leave*, generates an initial discourse state. Salient entities in

Discourse state 1	
-----	
<b>#38 is interval</b>	#45 is stateOfAffairs
	#45 < #44
	distanceAt(#45, #41, #40, #48)
	member(#45, #39)
<b>#39 is event</b>	
<b>leave(#39)</b>	#47 < #46
theta(#39, agent, #41)	
theta(#39, object, #40)	#48 < #47
#39 is setOfStatesOfAffairs	
moveFrom(#39)	of(#49, #41)
denotesMovement(#39)	
dense(#39)	<b>speechTime(#39) = now</b>
validProgressive(#39)	<b>#38 &lt; now</b>
<b>named(#40, London)</b>	<b>eventTime(#39) &lt; now</b>
<b>#40 is city</b>	<b>referenceTime(#39) &lt; now</b>
<b>origin(#40, #39)</b>	0 < #48
	0 < #46
#41 is student	<b>aspect(simple, #38, #39)</b>
#42 is stateOfAffairs	accommodated(named(#40, London))
#42 < #43	
#42 < #45	accommodated(ttype(#41, student))
distanceAt(#42, #41, #40, 0)	
sTart(#42, #39)	
member(#42, #39)	
#43 is stateOfAffairs	DISCOURSE STRUCTURE
distanceAt(#43, #41, #40, #46)	*****
member(#43, #39)	0
	1
#44 is stateOfAffairs	1
#44 < #43	
distanceAt(#44, #41, #40, #47)	(1) the student left London .
member(#44, #39)	

Fig. 2. Discourse State 1



Discourse state 1	Discourse state 2
-----	-----
#53 is interval	#57 < #58
	#57 < #60
#54 is event	#65 is event
<b>leave(#54)</b>	theta(#65, agent, #56)
theta(#54, agent, #56)	theta(#54, object, #55)
theta(#54, object, #55)	#54 is setOfStatesOfAffairs
#54 is setOfStatesOfAffairs	moveFrom(#54)
moveFrom(#54)	denotesMovement(#54)
denotesMovement(#54)	dense(#54)
dense(#54)	validProgressive(#54)
validProgressive(#54)	
<b>named(#55, London)</b>	<b>named(#66, Paris)</b>
#55 is city	<b>destination(#66, #65)</b>
<b>origin(#55, #54)</b>	<b>target(#66, #65)</b>
#56 is student	#67 is stateOfAffairs
...	distanceAt(#67, #56, #66, 0)
...	member(#67, #65)
speechTime(#54) = now	end(#67, #65)
	...
#53 < now	...
eventTime(#54) < now	
referenceTime(#54) < now	
0 < #63	DISCOURSE STRUCTURE
0 < #61	*****
	0
aspect(simple, #53, #54)	
	1
accommodated(named(#55, London))	
	2
	(1) the student left London .
	(2) she went to Paris .

Fig. 3. Discourse State 2

the discourse are assigned unique identifiers and it is evident that *leave*(#39) is an event with which an agent (the student) and an object (London) are associated. The system exploits thematic roles here in classifying participants in the event (Fillmore, 1968). Further information about the relationship between the event and its associated entities (which follows from the discussion in Section 5) also features in the model.

While *leave*(#39) is identified as an event, another entity, #38, is identified as an interval, which provides a perspective on the event. The relationship between the two is evident towards the end of the model where the event is oriented with respect to tense and aspect. The influence of interval semantics, which is central to our temporal representation, is apparent here (Dowty, 1979; Allen, 1984; Allen and Ferguson, 1994). Intervals underpin the ‘extended now’ theory, as proposed by Dowty: the perfect locates an event within a period of time that began in the past and extends up to the present moment. The simple past, by contrast, specifies that an event occurred at a past time that is separated from the present by some interval.

Three predicates determine *leave*’s temporal properties: *speechTime*(#39) = now asserts that the discourse is uttered in the present; *eventTime*(#39) < now indicates that the leaving event took place in the past; *referenceTime*(#39) < now asserts that the point of reference is also in the past. This predicate identifies the time at which the event is being referred to, with respect to speech time. The aspect predicate

```

Discourse state 1
-----
...
...
Discourse state 2
-----
...
#89 is event
...
go(#89)
#89 is setOfStatesOfAffairs
dense(#89)
validProgressive(#89)
moveTo(#89)
moveTowards(#89)
...
...
speechTime(#89) = now

#92 < now
eventTime(#89) < now
referenceTime(#89) < now

aspect(simple, #92, #89)
...
...

Discourse state 3
-----
...
#99 is event
theta(#99, agent, #80)
theta(#99, object, #79)
related(#99, #100)
love(#99)
#99 is setOfStatesOfAffairs
stative(#99)

#101 is stateOfAffairs
member(#101, #99)
lovesAt(#101, #80, #79)
...
speechTime(#99) = now
eventTime(#99) < now
referenceTime(#99) < now

aspect(perfect, #92, #99)
...
DISCOURSE STRUCTURE
*****
0
|
1
|
2
|
3
(1) the student left London .
(2) she went to Paris .
(3) she had loved London .
    
```

Fig. 4. Discourse State 3

reflects the fact that the aspectual interval has simple aspect with respect to the event: `aspect(simple, #38, #39)`. Our perspective on tense and aspect is consistent with Dowty’s (1979) analysis, which is founded upon Vendler’s ideas (Vendler, 1957). This informs Smith’s work (Smith, 1991) who also adopts Reichenbach’s notion of speech time, event time and reference time (Reichenbach, 1947).

The discourse structure appears in graphical form at the end of the model. It shows the relationships between the discrete discourse events that make up the current discourse. Discourse State 0, which is not presented in Figure 2, represents the background knowledge available to the system. For example, logical axioms encode the fact that London is a city and we see that the appearance of London in the logical form triggers the relevant inference in the model: `#40 is city`. Discourse State 1 comprises the integration of background knowledge with the specific entailments derived from the logical form.

*4.1.1 Extending the discourse.* The discourse is extended by entering another sentence and instructing the system to generate a new model incorporating the entailments of each sentence input so far. A new set of identifiers is generated; the first sentence retains the inferences from the previous discourse while the new event is modelled as an

additional discourse state. Adding the sentence: *She went to Paris* produces the abbreviated model in Figure 3.

*4.1.2 Introducing the perfect tense.* By introducing the perfect tense into the discourse with the sentence *She had loved London*, the resulting model indicates that the identifier #92 is associated with the interval that features in the simple aspect predicate in Discourse State 2. The same interval appears in the aspect predicate for Discourse State 3, as illustrated in Figure 4: aspect(perfect, #92, #99). When the perfect is introduced, it is interpreted with respect to the nearest past temporal entity. The temporal profile of the discourse places the time of the ‘loving’ event in the past with respect to speech time, which is now. We are also referring to a ‘loving’ event having taken place in the past; however, because this event shares an aspect interval with the ‘going’ event, and is labelled as perfect, then we are referring to an event that took place in the past with respect to the ‘going’ event. Therefore, we interpret the ‘loving’ event as having temporally preceded going to Paris.

While this sample discourse illustrates the fact that we can build significant amounts of background knowledge into our representations, the textual output is pretty inaccessible. Consequently, we want to interpret our models graphically. By extracting a stylised visual representation of the semantic entailments of a discourse we can explore ways of providing stimulating feedback to learners which gives them insight into the structure and meaning of what they have written.

## 5 Establishing the ontology of events

The timeline moves in a single direction from past through the present to the future (Reichenbach, 1971). Determining how to capture the properties of this phenomenon and to identify its implications for modelling language use are contentious. Allen’s work on temporal, interval-based logic to develop models that reason with beliefs and intentions, which addresses the limitations of the situation calculus (McCarthy and Hayes, 1969), demonstrates how temporal relations can be handled using interval semantics (Allen, 1984; Allen and Ferguson, 1994). Yet the linear order in which narrative events are related does not reflect the temporal order governing the way things actually happened (ter Meulen, 1997). Context, aspect, causal connection and knowledge impact upon the relationships between discourse components.

The granularity at which people interpret the allocation of time to events is variable and dependent upon context. Some events are perceived as instantaneous while others are seemingly endless processes. Our approach is to identify a stylised representation for different types of event, with each encapsulating its generalised individual properties. These sparse stylisations are indicative of general temporal signatures: theoretically the underlying temporal structures are infinitely divisible into smaller time slices. We wish to identify those properties that we think it useful to highlight in depicting the temporal features of narratives. We represent events as intervals that are manifested as sets of states of affairs. A state of affairs (SOA) is an implicitly temporal entity with which other properties can be associated to build richer models of what events denote.

Three classes of sets of SOAs are defined within the ontology: a state consists of a single SOA whose properties persist throughout the state; an achievement captures

the fact that something occurs pretty much at once: this is represented by two SOAs, one where the property does not hold followed by one where it does; an activity consists of four SOAs in which a set of dynamic properties is represented.

The semantic entailments of lexical items are encoded as logical axioms, or meaning postulates. Axioms define the properties of sets of SOAs; they capture the sense that a SOA is a temporal object and that the members of sets of such objects are ordered. Each set of SOAs is oriented with respect to event time, reference time and speech time. This is consistent with the earlier reference to work by Smith (1991), and by Reichenbach (1947).

A number of verbs to be modelled reflect movement and we capture the generic properties of these events. This entails a stylised interpretation of the fact that particular properties are in a state of flux throughout a compacted, indeterminate series of time slices: between any two instants in time a further instant exists (the set of instants that spans the event is 'dense'). This is clearly impossible to represent graphically, so we adopt a canonical representation which captures these properties at a very coarse level of detail. This is to indicate that certain properties change at each instant during the process. Dense sets of SOAs comprise three canonical members which are ordered, and none of them represents the start or the end of the set. This configuration provides a foundation for modelling activities which can be used to handle movement.

Movement is modelled to or towards some target or from some origin in activities such as *go* or *drive*. The entailments of such events are that some agent gets progressively nearer to or further from a target or origin as time increases. The axioms dealing with movement from specify that they represent a setOfStatesOfAffairs and they are defined as dense so they inherit the stylised profile incorporating three SOAs.

Additional rules specify that some events include a point which is their greatest lower bound—the 'start' of the event. This member will be a SOA which temporally precedes all other members of the set. Identifying properties such as starts and ends of sets of SOAs enables us to capture significant features of particular types of event. Some events have starts, some have ends, some have both, and some have neither. By linking events to intervals and integrating information about their bounds we can build detailed structures reflecting physical and temporal constraints implicit in the events and the entities participating in them (Dowty, 1979). Our aim is to encapsulate sufficient detail to portray a standardised set of features for our event types. While we can focus on the fact that a leaving event has a start, for example, we seem to have no way of determining when such an event ends. Equally, the starts and ends of states such as knowing something can also be indistinct. Our axioms capture these event signatures. We reinforce these underlying characteristics when we generate event images, as described in Section 6.

Spatial properties are added to the temporal framework to determine the relationship between the agent and the origin or destination of the event as the activity proceeds. An order is established between the distances associated with each SOA, which reflects the temporal relationship between them. Just as time increases as the agent moves from the origin, so the distance between the agent and the origin increases.

A parallel set of axioms handles movement in the opposite direction. A distinction is made between moving towards and moving to a destination. Again, these sets of SOAs are defined as dense but the framework is extended by adding a member

representing the end of the event. Mirror axioms establish temporal and spatial relationships between the agent and the target. Time continues to increase but this time the distance between the two entities decreases as the event proceeds. The distinction between approaching and reaching a destination is that the distance between the agent and the destination is 0 when it is reached, and additional rules capture this property.

Achievements are represented as sets of two temporally ordered SOAs. Implicit in this view is the fact that something occurs, after which, some property holds. This is a canonical representation which focuses on a particular point of change.

Events classed as stative are defined as a setOfStatesOfAffairs comprising a single member. This device adopts a canonical format to indicate that properties hold throughout the duration of a state. The verb *know* represents a state. The *knowsAt* predicate asserts a relationship between the agent and the object that obtains throughout the duration of the event.

## 6 Generating visual images of individual events

Individual events, depicted by labelled, shaded blocks, encapsulate the underlying ontology and we make them interactive so that this information is accessible.

Activities involving movement capture the sense of getting closer to or further from a point of arrival or departure by increasing or decreasing the shading to reflect the effects of time and distance. The following example illustrates how the system handles the sentence: *The student left London*. The user enters the sentence into the text box and invokes the *show time lines* action, as shown in Figure 5.

This generates the Temporal Discourse window which displays the event, as illustrated in Figure 6.

Images convey the entailments of an event from left to right: the leftmost bar indicates a point of departure. The shading decreases in intensity to signify movement away from the point of departure. This reflects the fact that there is no determinate end point to a leaving event.



Fig. 5. Invoking the event graphic



Fig. 6. Activity graphic: *The student left London*




Fig. 7. Achievement graphic: *My parents decided they would travel to England*

The template for achievement-type verbs displays two SOAs depicting the transition from a SOA before the event and a SOA during which the entailments of the achievement hold. A standard colour pattern distinguishes this particular type of event. Figure 7 shows the system's interpretation of an example based upon the learner data where the system uses a stylised format to depict the future tense in the internal event implicit in the use of *decide*: *My parents decided they would travel to England*.

The relationship between speech time and the future travelling event is uncertain in this single sentence. If we had more contextual information we might want to place the travelling event before speech time, but without further information we do not even know whether this event took place at all, let alone when.

By clicking on the event image, users are able to access the verb's underlying properties, which are encoded in the meaning postulates. These graphical features provide some contextual structure for the purely textual models presented in Section 4.1. Users can generate both an interactive image of the stylised representation of the event and a model of the event's properties.

Users can view the properties of specific SOAs by clicking on the individual SOA images in the centre of the window. We can see that the first of the two SOAs has been selected in Figure 8. The second option from the *actions* menu generates the semantic model depicting the complete set of SOAs. This presents the identifiers associated with the event and its main constituents; it also reveals the relationships between individual SOAs and between entities participating in the event. Generating the semantic model for *decide* shows the temporal ordering between the two SOAs. The second SOA indicates a commitment by the agent to some entity, which in this case, represents another event, *travel*, as we can see from Figure 8.

This tabular format provides access to bundles of information associated with each SOA. This facility can potentially be developed further as a learning resource. It illustrates how modelling and interacting with the ontology of events offers the potential to enrich learners' understanding of what linguistic components entail and how they are constrained in language use, although it also needs to be refined in response to learner interaction.

Stylistically, a larger image is used to convey the sense that the properties of states persist over a longer duration than either activities or achievements. States are also drawn at a lower level than the other two classes of event so that they can be represented as overlapping with other events to indicate that their properties can hold simultaneously. States are identified visually with their own shading template

which fades into and out of colour at the beginning and end of the event respectively. This denotes the generalisation that no specific start and end points are associated with such events, as the image in Figure 9 shows:

### 7 The orientation of speech time

The examples we have seen so far have been based on discourses describing past events and so they have preceded speech time. As most of our sample data concern narratives of past events, these constructions feature prominently. We can reverse the process to represent future events which occur after speech time, as in Figure 10.

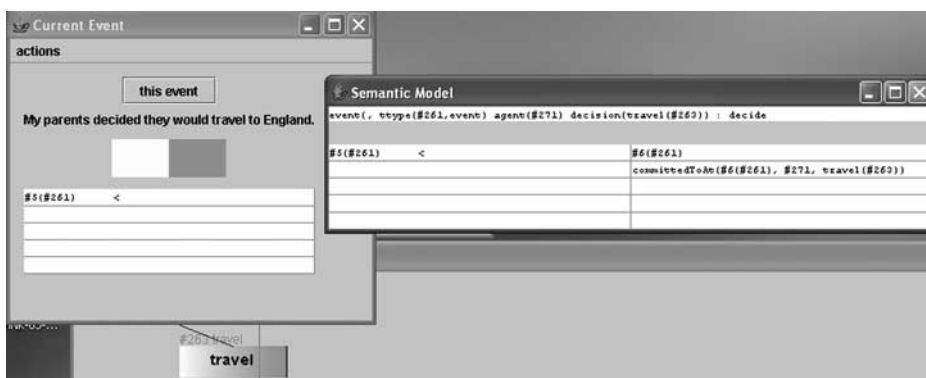


Fig. 8. Properties of decide



Fig. 9. Stative graphic: *The student loved London*



Fig. 10. Future: *The student will go to London*



Fig. 11. Simple present: *The student leaves Paris*



Fig. 12. Present perfect: *The student has arrived in London*

We also distinguish between the simple present tense and the present perfect. We depict the former as incorporating and extending beyond speech time. This conveys the sense that an event is taking place now with the assumption that it will continue occurring in the immediate future, as Figure 11 shows.

We draw events in the present perfect abutting speech time: information about the event extends from the past up to now, as Figure 12 illustrates.

## 8 Discourse relations

Verb type and aspect are the principal features in determining discourse relations. The system assigns these properties to each sentence as it is generated. The system builds a discourse tree to which new discourse states are added as either sisters or daughters, depending upon the tense used, as the discourse progresses. The aim is to utilise this structure by exploring the relationships between the current Discourse State and the previous nodes in the tree.

It is possible to use a number of rules to establish relations between sentences. The default rule for discourse structure is that events in a narrative are related in the order in which they occur. Often this is not the case, and apart from the temporal ordering of events, there can be other factors affecting the relationship between utterances, such as causality, for example.

Reference time is introduced by sentences with simple aspect and new simple sentences introduce new reference times. The most straightforward narratives using simple aspect are interpreted as a sequence of events. We exploit the system's structural framework by searching for and extracting relevant relationships that exist between the current discourse state and its antecedents. The principles are similar to Centering Theory (Grosz, Joshi *et al.*, 1995). Relations are dependent upon the tense and aspect properties of pairs of events within specific discourse configurations.



Relations can either be sequential or overlapping. Within the class of overlapping relations, the idea of a simultaneous relation is determined by the fact that related events appear in the progressive. Similar notions of overlapping temporal events feature in ter Meulen (1997). The rules are sufficiently general and flexible to capture a significant range of relations between events based upon their tense and aspect properties.

### 9 A graphical interpretation of discourses

To be of value to learners, relationships between events are translated into a visual display that approximates to the temporal entailments of their narratives. The visual interpretation of the discourse is regenerated dynamically as it evolves. In this section we model a series of test discourses to illustrate the system's behaviour.

The simplest kind of narrative relates a series of events as a sequence:

*The student left Berlin.  
She approached London.  
She arrived in London.*

Each sentence uses the simple past tense to denote an event; this means that each sentence is assigned a new aspect interval by the system. Implicit in the interpretation is that speech time is now and that all events precede speech time, as shown in Figure 13.

Statives appear at lower levels, as they are introduced, to give the impression that they overlap with surrounding events. The following test discourse illustrates:

*The student left Berlin.  
She knew London.  
She loved London.  
She arrived in the city.*

Successive statives are shown to be overlapping with each other while non-statives continue to appear at the same level, as we see in Figure 14(a).

Use of the progressive tense can imply that events are occurring simultaneously. While there are circumstances in which the progressive is used to indicate a planned or anticipated event, the following test discourse is interpreted as denoting simultaneous events:

*The student approached London.  
She was leaving Paris.  
Her friend was arriving in Paris.  
She loved the city.*

In this discourse the aspect interval which the system assigns to *approach* is also assigned to *leave*. Both events are placed temporally in the past and the use



Fig. 13. Sequential relations

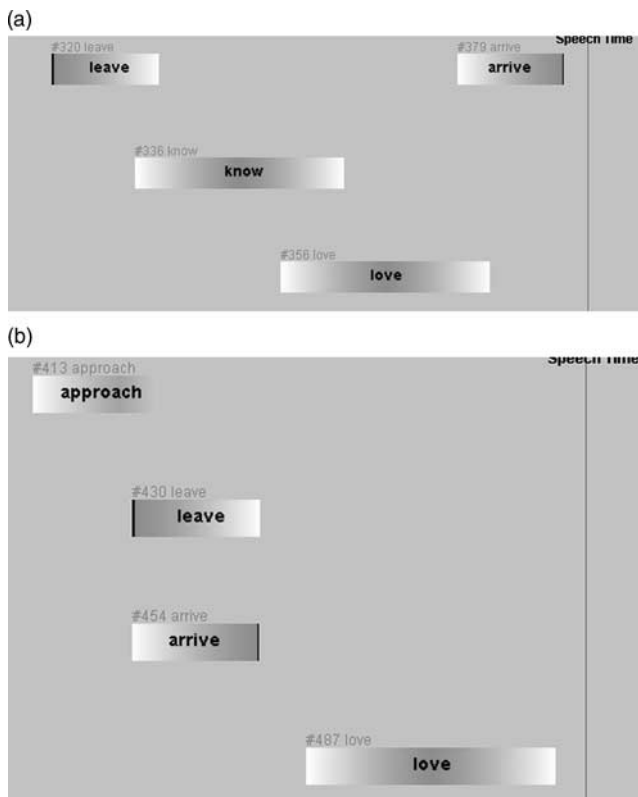


Fig. 14. Multiple overlapping and simultaneous relations

of the progressive signals that we are referring to events that share the same past aspect interval with respect to speech time. This common temporal reference point extends to the next event in the discourse, *arrive*, which is also designated the same aspect interval. The final sentence is assigned a new aspect interval indicating that a new reference point has been introduced. We utilise our convention of representing a state as a generally overlapping event. The depiction in Figure 14(b) shows how the discourse contrasts with non-simultaneous overlapping representation.

Introducing perfect aspect to discourses in the past tense alters their temporal configuration. In the following example event time and reference time for each event are located in the past. However, they share a common aspect interval. This denotes the fact that the event marked by the perfect, *leave* in this case, viewed from the past referential perspective of *approach* here, was already complete:

*The student approached London.*  
*She had left Berlin.*

Temporally, the leaving event precedes the approach. The system rules interpret this as a ‘precedes’ relation, and the graphical component reverses the location of the events in the sequence, as in Figure 15(a).




Fig. 15. Past perfect and multiple past perfect relations



Fig. 16. Changing temporal viewpoint

Additional past perfect sentences are also assigned the same aspect interval. By extending the discourse in the past perfect all three of the following events share a common aspect interval in the regenerated extended discourse:

*The student approached London.*  
*She had left Berlin.*  
*She had passed her exams.*

The system interprets a series of past perfect sentences such as this as a sequence. The order in which the information is conveyed in the narrative is preserved in the graphical representation, as we can see from Figure 15(b).

### 9.1 Signalling changes of temporal viewpoint

The system indicates where a change in the temporal viewpoint of the speaker occurs in a discourse, since this often introduces a degree of disfluency. This is signalled by altering the event labelling to alert the learner to the change. We mark the shift in viewpoint by using a larger italicised font, as the example in Figure 16 illustrates.

*The student left London.*  
*She has arrived in Paris.*

When the system generates models for this kind of discourse there are separate designations for event time, reference time and aspect interval. Detecting these discrepancies, the system signals the change.

## 10 Modelling learner data

Data were collected from German learners of English, as discussed in Section 3. The samples are not dealt with verbatim as some of the constructions are quite

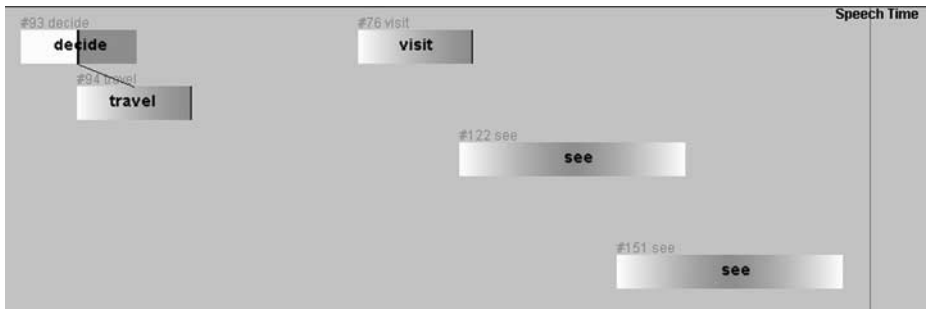


Fig. 17. Model one

complex. The models aim to capture the underlying temporal relationships within the discourses. The initial examples model sample data in which tense and aspect are used correctly.

### 10.1 Model one

The first model is simplified to eliminate a number of syntactic difficulties:

*My first visit to England was in 1990 when I was about 12 years old. My parents had decided to travel and not to stay in one place. So I was able to see not only London but also many other cities of this small island.*

The visual model in Figure 17 represents the following modified input:

*I visited England for the first time.  
My parents had decided they would travel to England.  
I saw London  
I also saw more cities.*

The depiction of events is consistent with the rules presented earlier. The system infers that the decision to travel was taken before arriving in the country.

### 10.2 Model two

The next example captures a straightforward sequential narrative in which the learner uses tense and aspect correctly:

*We took the ferry from Hamburg to Harwich and then went to Brighton by car.  
We travelled along the south coast and then went back to London to get an impression of this multicultural city.*

The discourse, shown in Figure 18, is modelled from the following paraphrase of the original:

*We travelled from Hamburg on the ferry.  
We went to Harwich.  
We went to Brighton by car.  
We travelled along the south coast back to London.*



Fig. 18. Model two

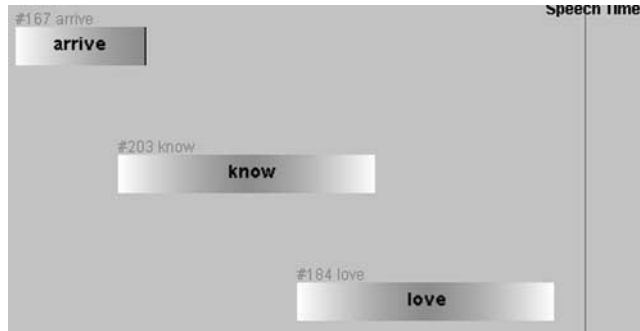


Fig. 19. Model three

### 10.3 Model three

The next example demonstrates that the writer has used tense and aspect correctly:

*I think I loved England from the first time we arrived on the island by ferry.  
I had known England before as a land of fish and chips, and of course I had to try it.*

Although the narrative has been simplified, it reflects the temporal structure of the original. The output is shown in Figure 19.

*We arrived on the island by ferry.  
I loved England.  
I had known England as a country of chips.*

### 10.4 Model four

In the following example, the use of the simple past in the first clause is problematic. The adjunct, *over the last nine months*, implies that the author is referring to a period that includes the present:

*However, I got to know many people over the last nine months and I feel at home here.*

The following interpretation retains the temporal relationships of the original and the rules flag up the change in tense (through italics), as illustrated in Figure 20.

*I met several friends over recent months.  
I feel at home.*




Fig. 20. Model four

The representation in Figure 21, using the present perfect for the initial sentence, places the perspective in the present, and shows that the system signals a consistent use of tense and aspect:

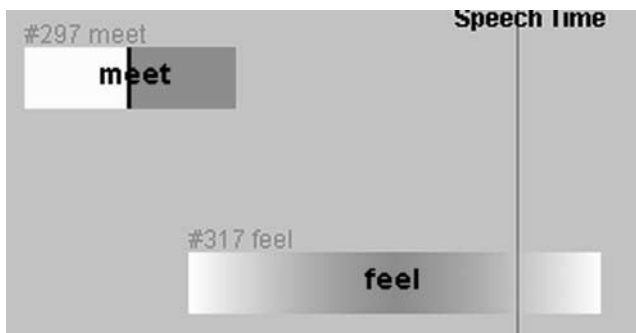


Fig. 21. Model four updated

### 10.5 Model five

The final sample shows similar temporal characteristics to the previous fragment:

*As I'm studying History and Geography here in Chester, I think it helped me to understand this country a lot better in terms of historical background, environmental, social and economic conditions.*

Much of the surrounding text is eliminated from the analysis in order to focus on the temporal relationship between *study* and the secondary verb *help*, which seems to be problematic in this example, as illustrated in Figure 22.

*I am studying history in Chester.  
I think it helped me.*

If the initial sentence is retained and the past is replaced by the present perfect in the next sentence then consistent tense and aspect is maintained, as Figure 23 shows.

*I am studying history in Chester.  
I think it has helped me.*




Fig. 22. Model five




Fig. 23. Model five updated

## 11 Conclusions

This paper has presented the application of an experimental and extensible NLP system to the support of language learners by generating real-time discourse models of their narratives and converting them into dynamic graphical displays. These encapsulate the semantic entailments of events and their temporal relations within a discourse. By utilising the tense and aspect properties of discourse relations, the system is able to provide visual representations of alternative temporal constructions in English. As sample learner data indicate misuse of tense and aspect in English, the application of discourse modelling techniques suggests a potentially beneficial contribution to the development of software capable of providing support to learners composing freely written scripts.

Many of the theoretical discussions of time in natural language use visualisations to convey temporal properties and relations: (Dowty, 1979; Smith, 1991; ter Meulen, 1997; Kazanina and Phillips, 2003). This seems to be an intuitive way of representing these phenomena, and as we observe in our concluding remarks in Section 2, there is evidence that graphical techniques can enhance learning. While further work is required to evaluate and refine our particular approach in response to learner feedback, we anticipate that our visual representation will enable learners to reflect on the match (or otherwise) between what they meant and the picture that their use of temporal expressions produced. Externalising and reflecting on your knowledge is widely regarded as an important step in learning (Chen, Cannon *et al.*, 2005; Lehtonen, Appelqvist *et al.*, 2002; Flanagan, Eckert *et al.*, 2007), and we would expect the same to be true in this challenging area of language learning.

## References

- Allen, J. F. (1984) Towards a General Theory of Action and Time. *Artificial Intelligence*, **23**: 123–154.
- Allen, J. F. and Ferguson, G. (1994) Actions and Events in Interval Temporal Logic. *Journal of Logic and Computation*, **4**(5): 531–579.
- Anderson-Hsieh, J. (1994) Interpreting Visual Feedback on Suprasegmentals in Computer Assisted Pronunciation Instruction. *CALICO*, **11**(4): 5–22.
- Barwise, J. and Perry, J. (1983) *Situations and Attitudes*. Cambridge, MA: Bradford Books.
- Beatty, K. (2003) *Teaching and Researching Computer-assisted Language Learning*. London: Pearson Education.
- Bull, S. (1997) Promoting Effective Learning Strategy in CALL. *Computer Assisted Language Learning*, **10**(1): 3–39.
- Bull, S., Dong, X., Britland, M. and Guo, Y. (2008) Can Students Edit their Learner Model Appropriately? In: Woolf, B. (ed.), *ITS 2008*. Heidelberg: Springer-Verlag, 674–676.
- Chen, H. L., Cannon, D., Gabrio, J., Leifer, L., Toye, G., and Bailey, T. (2005) *Using Wikis and Weblogs to Support Reflective Learning in an Introductory Engineering Design Course*. Proceedings of the American Society for Engineering Education Annual Conference and Exposition. [http://riee.stevens.edu/fileadmin/riee/pdf/ASEE2005\\_Paper\\_Wikis\\_and\\_Weblogs.pdf](http://riee.stevens.edu/fileadmin/riee/pdf/ASEE2005_Paper_Wikis_and_Weblogs.pdf)
- Dekhinet, R. (2008) Online enhanced corrective feedback for ESL learners in higher education. *Computer Assisted Language Learning*, **21**(5): 409–425.
- Dimitrova, V. (2003) STyLE-OLM: Interactive Open Learner Modelling. *International Journal of Artificial Intelligence in Education*, **13**: 35–78.



- Dimitrova, V., Brna, P. and Self, J. (2002) The Design and Implementation of a Graphical Communication Medium for Interactive Open Learner Modelling. In: Cerri S., Gouardères, G. and Paraguaçu, F. (eds.), *Intelligent Tutoring Systems*, Vol. 2363. Berlin/Heidelberg: Springer, 432–441.
- Dorfmueller-Karpusa, K. (1985) The Expression of Time in Texts. In: Huellen, W. and Schulze, R. (eds.), *Tempus, Zeit und Text*. Heidelberg: Winter, 39–51.
- Dowty, D. R. (1979) *Word Meaning and Montague Grammar*. Dordrecht: D. Reidel Publishing Company.
- Duerich, K. (2005) The Acquisition of the English Tense and Aspect System by German Adult Learners. *Philosophische Fakultät*. Chemnitz: Technische Universität Chemnitz.
- Fillmore, C. (1968) The Case for Case. In: Bach, E. and Harms, R. T. (eds.), *Universals in Linguistic Theory*. New York: Holt, Rinehart & Winston, 1–88.
- Flanagan, T., Eckert, C. and Clarkson, J. P. (2007) Externalising tacit overview knowledge: A model-based approach to supporting design teams. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, **21**(3): 227–242.
- Fried, M. and Oestman, J. O. (2004) Construction grammar: A thumbnail sketch. In: Fried, M. and Oestman, J. O. (eds.), *Construction grammar in a cross-language perspective*. Amsterdam: John Benjamins, 11–86.
- Groenendijk, J. and Stokhof, M. (1991) Dynamic Predicate Logic. *Linguistics and Philosophy*, **14**: 39–100.
- Grosz, B. J., Joshi, A. K. and Weinstein, S. (1995) Centering: A Framework for Modelling the Local Coherence of Discourse. *Computational Linguistics*, **21**(2): 203–225.
- Jaen, M. M. and Basanta, C. P. (2009) Developing conversational competence through language awareness and multimodality: the use of DVDs. *ReCALL*, **21**(3): 283–301.
- Johnson-Laird, P. N. (1983) *Mental Models: Towards a Cognitive Science of Language, Inference and Consciousness*. Cambridge, Mass: Harvard University Press.
- Johnson-Laird, P. N. (2006) *How we reason*. New York: Oxford University Press.
- Kamp, H. and Reyle, U. (1993) *From Discourse to Logic*. Dordrecht: Kluwer.
- Kazanina, N. and Phillips, C. (2003) *Temporal Reference Frames and the Imperfective Paradox*. Proceedings WCCFL 22, Somerville, MA: Cascadilla Press.
- Kohonen, V. (1992) Experiential Language Learning: Second Language Learning as Cooperative Learner Education. In: Nunan, D. (ed.), *Collaborative Language Learning and Teaching*. Cambridge: Cambridge University Press, 14–39.
- Lehtonen, J.-M., Appelqvist, P., and Saranen, J. (2002) *Learning Operations Management with Web-based Manufacturing Simulation*. Proceedings of the European Conference on Information Systems, Gdansk, Poland. <http://is2.lse.ac.uk/asp/aspectis/20020068.pdf>
- LEO GmbH (2011). LEO Deutsch-Englisches Woerterbuch. <http://dict.leo.org/>
- Long, M. H. (1991) Focus on form: A design feature in language teaching methodology. In: de Bot, K., Ginsberg, R. B. and Kramsch, C. (eds.), *Foreign Language Research in Cross-cultural Perspective*. Amsterdam: John Benjamins, 39–52.
- Mann, W. C. (2010) Introduction to Rhetorical Structure Theory. <http://www.sfu.ca/rst/01intro/intro.html>
- Mann, W. C. and Thompson, S. A. (1988) Rhetorical Structure Theory: Towards a functional theory of text organisation. *Text*, **8**(3): 243–281.
- Marcu, D. (1997). *The rhetorical parsing of natural language texts*. Proceedings of the Association for Computational Linguistics (ACL-97). San Fransisco, CA: Morgan Kaufmann, 96–103.
- McCarthy, J. and Hayes, P. (1969) Some Philosophical Problems from the Standpoint of Artificial Intelligence. *Machine Intelligence*, **4**: 463–502.

- Nerbonne, J. (2005) Natural Language Processing in Computer-Assisted Language Learning. In: Mitkov, R. (ed.), *The Oxford Dictionary of Computational Linguistics*. Oxford: Oxford University Press, 670–698.
- Pollard, C., and Sag, I. A. (1987) *Information-Based Syntax and Semantics, Vol I: Fundamentals*. Stanford: CSLI Lecture Notes 13.
- Pollard, C. and Sag, I. A. (1994) *Head-Driven Phrase Structure Grammar*. Chicago: University of Chicago Press.
- Pustejovsky, J. (1995) *The Generative Lexicon*. Cambridge MA: The MIT Press.
- Reichenbach, H. (1947) *Elements of Symbolic Logic*. Berkeley: University of California Press.
- Reichenbach, H. (1971) *The Direction of Time*. Berkeley CA: University of California Press.
- Rutherford, W. E. and Sharwood-Smith, M. (1985) Consciousness-raising and universal grammar. *Applied Linguistics*, **6**(3): 274–282.
- Sanders, R. (1991) Error Analysis in Purely Syntactic Parsing of Free Input the Example of German. *CALICO*, **9**(1): 72–89.
- Schmidt, R. and Frota, S. (1986) Developing basic conversational ability in a second language: A case study of an adult learner of Portuguese. In: Day, R. (ed.), *Talking to learn: Conversation in second language acquisition*. Rowley, MA: Newbury House, 237–326.
- Schmidt, R. W. (1990) The role of consciousness in in second language learning. *Applied Linguistics*, **11**(2): 274–282.
- Schulze, M. and Penner, N. (2008) Construction Grammar in ICALL. *Computer Assisted Language Learning*, **21**(5): 427–440.
- Shimojima, A. (1999) The graphics-linguistic distinction. *Artificial Intelligence Review*, **13**: 313–335.
- Smith, C. S. (1991) *The Parameter of Aspect*. Dordrecht: Kluwer.
- Sowa, J. (1984) *Conceptual structures: Information processing in mind and machine*. Reading, MA: Addison-Wesley.
- ter Meulen, A. G. B. (1997) *Representing Time in Natural Language*. Cambridge MA: Bradford Books.
- Vandeventer Faltin, A. (2003) Natural language processing tools for computer assisted language learning. *Linguistik online*, **17**(5): 137–153.
- Vendler, Z. (1957) Verbs and Times. *The Philosophical Review*, **66**(2): 143–160.
- Warschauer, M. (1996) Computer Assisted Language Learning: an Introduction. In: Fotos, S. (ed.), *Multimedia language teaching*. Tokyo: Logos International, 3–20.
- Xu, J. and Bull, S. (2010) Encouraging advanced second language speakers to recognise their language difficulties: a personalised, computer-based approach. *Computer Assisted Language Learning*, **23**(2): 111–127.