

KNOWLEDGE MEDIA

KMi

I N S T I T U T E

Modelling Discourse in Contested Domains: A Semiotic and Cognitive Framework

Technical Report KMI-06-14
July, 2006

Clara Mancini and Simon Buckingham Shum

Mancini, C. and Buckingham Shum, S.J. (2006, in press). Modelling Discourse in Contested Domains: A Semiotic and Cognitive Framework. *International Journal of Human-Computer Studies*.

[PrePrint: <http://kmi.open.ac.uk/publications/pdf/KMI-TR-06-14.pdf>]



Modelling Discourse in Contested Domains: A Semiotic and Cognitive Framework

Clara Mancini* and Simon J. Buckingham Shum

Knowledge Media Institute & Centre for Research in Computing
The Open University, Milton Keynes, MK7 6AA, U.K.
Email: c.mancini@open.ac.uk / sbs@acm.org

* Corresponding author, Tel.: +44-1908-652165, Fax.: +44-1908-652140

Abstract

This paper examines the representational requirements for interactive, collaborative systems intended to support sensemaking and argumentation over contested issues. We argue that a perspective supported by semiotic and cognitively oriented discourse analyses offers both theoretical insights and motivates representational requirements for the semantics of tools for contesting meaning. We introduce our semiotic approach, highlighting its implications for discourse representation, before describing a research system (ClaiMaker) designed to support the construction of scholarly argumentation by allowing analysts to publish and contest ‘claims’ about scientific contributions. We show how ClaiMaker’s representational scheme is grounded in specific assumptions concerning the nature of explicit modelling, and the evolution of meaning within a discourse community. These characteristics allow the system to represent scholarly discourse as a dynamic process, in the form of continuously evolving structures. A cognitively oriented discourse analysis then shows how the use of a small set of cognitive relational primitives in the underlying ontology opens possibilities for offering users advanced forms of computational service for analysing collectively constructed argumentation networks.

Keywords

Argumentation, Discourse Representation, Coherence Relations, Ontologies, Semantic Annotation, Semantic Web, Semiotics

1 Introduction

Intellectual work in domains as diverse as scientific research, academic study, mission operations and scenario planning revolves around *sensemaking* (Wieck, 1995): the interpretation of potentially ambiguous signals from the environment, in order to take action. In Weick's framework, stakeholders have to literally "make sense": give form to their interpretations in one or more modalities in order to share and negotiate its meaning with others, possibly defending their interpretations and challenging others. The form that this takes will obviously vary with context: real time mission operations bear some similarities to real time scientific experiments, but are faster paced and (arguably) less reflective than the construction of scientific arguments in a journal paper, or a competitor analysis to underpin an organisational strategy. The common characteristics across these contexts are that *perspectives are contested* (one cannot assume a consensus to which everyone subscribes), and *the domain's content and boundaries are shifting* (depending on environmental events, the discovery of new relevant material, or a change in perspective).

Our long-term research objective is to understand the challenges and opportunities for sensemaking-support tools. As a specific example of sensemaking in a contested domain, we focus on e-Science, specifically the work of publishing and contesting research knowledge using the conventions of scholarly discourse. Elsewhere, we have described the design rationale, functionality and user evaluation studies around ClaiMaker¹, a software tool that is serving as a research vehicle (Buckingham Shum et al., 1999; In Press; Uren et al, 2003; In Press). This paper explores a different aspect, seeking to explain the system's theoretical foundation.

Our thesis is as follows:

- Sensemaking – expressing and contesting explicit, competing interpretations of the world – is a central activity in many knowledge-intensive tasks, which are conducted increasingly in an asynchronous, distributed manner over networks.
- An approach to augmenting sensemaking is to provide ways for analysts to annotate documents with possibly different interpretations, and structure the resulting discourse in ways that are both cognitively

¹ ClaiMaker: <http://claimaker.open.ac.uk>, from the Scholarly Ontologies project: <http://kmi.open.ac.uk/projects/scholonto>

tractable (augmenting human reasoning) and computationally tractable (augmenting machine reasoning)

- There are established theoretical perspectives for understanding discourse and the role of language in the collective construction of coherence which in principle could inform formal representations in software. We draw on two, namely Semiotics and Coherence Relations.
- These perspectives draw attention to certain requirements if one is to model discourse, which have motivated the design of a system for analysts to conduct distributed semantic annotation and discourse over the Web:
 - Firstly, we adopt a *relation-centric* approach, in which the nodes of a conceptual network are contextually defined by their inter-relationships, as motivated by a semiotic perspective.
 - Secondly, because the meaning of relations is also contextual, we have been exploring the potential of a representation scheme grounded in a *cognitive model of relational primitives*, which have been proposed as cross-modal, and context-independent.

The paper is organised as follows. We begin by framing the challenge of “modelling discourse” (§2). We then introduce the elements of semiotic systems (§3). We then summarise the scholarly discourse software tool that we have been developing (§4), in order to show how it provides a language in the terms of the semiotic analysis (§5). After discussing related research (§6) we conclude by describing ongoing and future work that combines extensions to the modelling in this paper, with user interface design and empirical user studies (§7).

2 Modelling discourse, and the discourse of modelling

“Modelling” provides a way to elicit, and where required harmonise, stakeholders’ conceptions of a problem or domain, that is, those elements that need to be represented for subsequent processing, whether by humans or machines. Thus, we consider designing a new paper form for an organisational process as modelling, as well as ontology-based modelling to enable computational interoperability between services or datasets.

In the context of developing a model-based software tool, consider two senses in which discourse may be modelled:

Discourse is modelled in the software design environment: Typically, the discourse that takes place between stakeholders as they develop a model is not itself modelled. By “discourse” we mean the typical discussions and arguments that take place such as *who should model, what to model, how to*

model, how to populate the model, how to deal with modelling inconsistencies, and so forth. Software design environments may support discourse within the team through instant messaging, voice/video conferencing, wikis, blogs or forums, but the environment typically has no computational access to the structure of this discourse beyond email/forum thread structure.

The ClaiMaker system that we reflect on in this paper could be considered as another tool to support collaborative modeling (e.g. distributed ontology engineering), although this is not our primary applications domain at present; our attention has been on scholarly discourse. However, elsewhere we have reported on the use of real time Dialogue Mapping and Conversational Modelling to mediate and capture the modeling discourse in order to yield argumentation-based design rationale, for instance, in ontology engineering (Buckingham Shum et al, 2002), Year 2000 contingency planning (Selvin and Buckingham Shum, 2002), and work practice modeling (Clancey et al, 2005).

Discourse is modelled in the final software artifact: If the software artifact is intended to be a tool to support either explicit *domain modelling* (which would include all semantic annotation and ontology engineering tools) or *discourse* (which would include all Computer-Mediated Communication - CMC), we can ask to what extent it provides support for users to disagree, either directly with each other (in a CMC context), or specifically, about how to model the world. Furthermore, even if the tool is not for ‘modelling’ or CMC explicitly, we may ask to what extent users are able to discuss, challenge and evolve the tool’s model of the world.

Reflecting on these two senses, if discourse is not modelled either during the design process, or in the software artifact, the implicit claim is that facilitating and possibly capturing such explicit discourse is either of little intrinsic importance (there is no value in adding structure, or preserving a trace), or is considered too difficult (it is best to keep communication informal), or is actually undesirable (politically sensitive discourse should not be captured, or encouraged). Ultimately, there is only one narrative or worldview worth keeping: the final ontology, the derived model that is populated, and the encapsulation of that model in an artifact.

If all stakeholders subscribe to the ontology’s worldview, then this is not problematic, and approaches to modeling and ontology engineering are adequate. The cognitive discipline imposed through the correct application of a modelling approach brings with it rigour, consistency and compatibility. However, this is problematic if the discourse remains ephemeral and invisible when there are different perspectives that stakeholders need to articulate and critique. Supporting discourse is about

supporting the construction of different perspectives on a domain, and not simply allowing the representation of any one of these perspectives. In discourse, the borders of domains may fluctuate and even dissolve, as new connections are made, and because concepts play many roles as the discourse threads in which they are used. In the discourse space, perspectives may change and multiply as speech acts are performed and responded to, and they get stronger or weaker in an incessant process of negotiation between speech agents. The shape of discourse is dynamic and mutable, and supporting its construction requires us to support the representation and development of its flow. This, we suggest, is the representational challenge for modelling discourse, which we will motivate in the following sections.

3 The elements of a semiotic system

Returning to *sensemaking* as construed by Weick (1995), *modelling* is a powerful way to *make sense* because we share our understanding with others through the use of a symbol system. Whenever we represent something, we trigger a complex process of *semiosis*, which a good representation system should take into account. Through this process, symbols take the place of the things that people want to talk about, whether these are real or imaginary, concrete or abstract. It is the connections holding between symbols and the things that they stand for, on the one hand, and symbols and the people who use them, on the other, that determine the meaning of any representation.

Pierce, one of the founders of the “science of signs”, gives us conceptual categories to describe this phenomenon. In Peirce’s terms (Pierce, 1931-35; Eco, 1968), any semiotic act consists of the association between a *symbol* (or *representamen*) and a *referent*, the entity to which the symbol refers, and for which it stands. Such association is possible through the mediation of what Peirce calls the *reference* (or *interpretant*): a concept or mental representation that allows the identification of the symbol as referring to the referent. The triadic relationship between symbol, referent and reference constitutes the sign.

A particular referent may be referred to by different symbols, depending on the context and on the *interpreter’s* perspective on the referent. For instance, in public places, the stylised icon on a restroom’s door of a human figure wearing either trousers or skirt refers to the human gender for which the restroom is intended, the representation of the clothing being the discriminating element. However, in biological or medical scholarly literature, the same concept of gender is expressed by different symbols altogether: a circle with a cross refers to the female gender, and a circle with an arrow refers to the male gender, the cross and the arrow being the discriminating elements. Other times, the same symbol may be used in

different contexts to refer to different things. So, if the stylised icon of a human figure is placed in a triangular street sign, the clothing is no longer relevant, instead the icon refers to possible pedestrians. However, when a symbol is used in a context different to its original, something from the original context remains attached to it. So if the medical circle+cross/arrow icons for gender were used to indicate a public toilet, they might also evoke medical *connotations* (see below).

These dynamics between symbols and the meanings that are attached to them are described by Saussure, the other founder of semiotics. In Saussurian terminology (Saussure, 1922; Barthes, 1964), the symbol corresponds to the *signifier* and its reference (the mental representation that allows the identification of the connection between the symbol and the referent) corresponds to the *signified*. Signifier and signified are the two 'faces' of a sign. In the examples above, we have two different signs. In the first one, the stylised shape of a human figure constitutes a signifier whose signified is respectively the concept of male or female, and whose referent is respectively a generic male and a generic female (the potential users of the restroom). However, in a different context, the same concept is signified by a different signifier (circle+arrow/cross), which also refers to the reality of a generic human male and female. As these different signifiers have different uses, a number of additional signifieds will be associated to each of them, that derive from their specific context of use. For instance, the two human icons used in public places are so strongly associated to the concept of restroom that they can also be used to signpost the facility.

In general, different interpreters associate additional signifieds to a signifier, based on their culture, experience, perception and purposes. For instance, when I use the word "dog" I think of a four-legged barking animal, and my referent can be a specific dog or the many dogs I have known in my life. Likewise the same word would commonly be associated to the same generic concept of a four-legged barking animal and referred to a specific dog or to the many dogs ever known by my interlocutor. So, although my interlocutor and I may both think of a four-legged barking animal, our respective mental representations of the same abstract concept are likely to differ as they will be based on the personal knowledge that each of us has about the referent (dogs). If my interlocutor has ever been attacked by a dog, fear and pain might be evoked, whereas these concepts may not be associated with my mental representation.

In semiotic terms, the more crystallised and conventionalised aspects of a mental representation constitute a sign's primary signified – its *denotation* – while its more specific and contextual aspects constitute the secondary signifieds – its *connotations*. The association between the signifier and the denotation is codified in the linguistic system and has conventional validity within it. The association between the signifier and its connotations is

largely grounded in the more specific context in which the sign is used, having more cultural value and rhetorical character.

More often than not, a word's *denotation* can be translated into another linguistic code through its association with the corresponding word. However, part or all of the original word's *connotations* are likely to get lost in the passage: in the new language those connotations are not necessarily associated with the corresponding sign, because different cultures are likely to have different perspectives on the same sign's referent. For instance, I can translate the English word 'dog' into the Korean word '개', but whereas in the English culture dogs are considered pets, in the Korean culture they are considered livestock, which evokes a different set of connotations. That is, in the two cultures the meaning of the word that refers to the same four-legged barking animal is different, because in each culture the 'set' of its associated signifieds is different, as a reflection of the fact that in each culture dogs play different roles. The concepts introduced above are summarised in

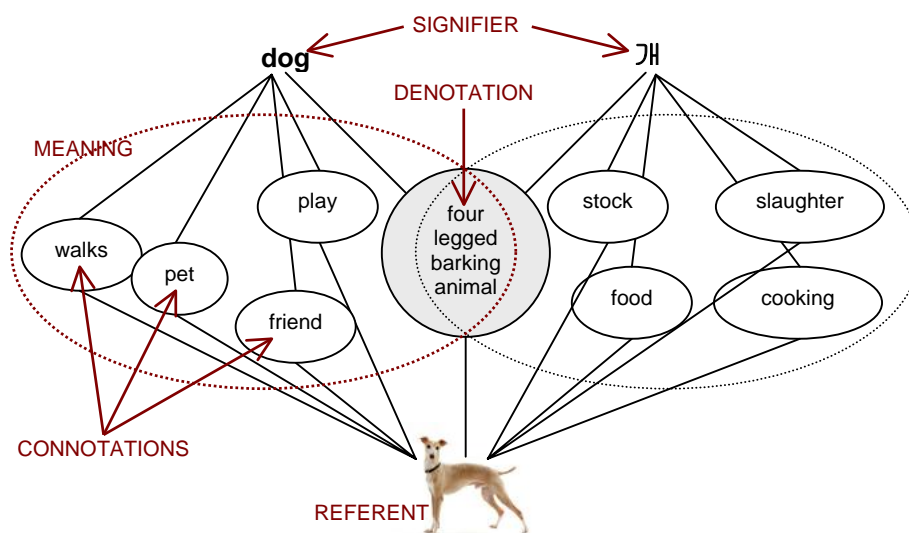


Figure 1. The components of a sign system.

This phenomenon, which concerns the relation between countries' or civilisations' macro-cultures, as well as the relation between small communities' micro-cultures, can be described by looking at any language (natural or not) as a system. In Saussure's terms, the signs of a linguistic system relate to one another along two axes, the *syntagmatic* being the axis of the relations 'in presentia', and the *paradigmatic* being the axis of the relations 'in absentia'. In this bi-axial reference system, the syntagma accounts for the combination of signs in the discourse chain, in which each sign plays a particular role. On the other hand, the paradigm accounts for the possible substitutions between the signs that could play the same role in each particular slot of the discourse chain, and that are associated by

semantic affinity. For instance, in the sentence “*the lion caught the antelope*”, instead of “*the lion*” there could be “*the predator*”; instead of the “*antelope*” there could be “*the prey*”; instead of “*caught*” there could be “*captured*”: substitutions that express different connotations in reference to the same situation. These possible substitutes express the range of signifieds (connotations) that can possibly be associated with that sign, that is, that sign’s meaning within that linguistic system.

These complex dynamics and levels of signification are at work in any natural communication process. However, they are seldom taken into account by ontology based knowledge representation systems, few of which are concerned with explicitly accounting for the distinction and dynamic interaction between the denotative and connotative dimensions of meaning (although for an example which does account for the connotative dimensions, see Liu (2000; 2005). This is where ClaiMaker differs from other ontology based systems, as we explain in the following sections.

4 ClaiMaker

ClaiMaker (Buckingham Shum et al., 1999; In Press; Uren et al, 2003; In Press) is a hypertext system designed specifically to represent discourse as a process of semiosis within the scholarly domain. In order to achieve this, ClaiMaker takes a semiformal approach, making use of constrained base relational classes to assist computational services, but imposing no constraints on how these relations are rendered, how nodes in the network are expressed, or how nodes are classified. ClaiMaker’s ontology allows users to establish as many referential relations between concepts and sources and connective relations between sources as they want to. This is how, through multiple connections between concepts and referent sources, different interpretations and perspectives can be expressed, and meaning is ‘implicitly negotiated’ every time that a new claim is produced in the discourse space².

² By ‘implicit negotiation’ we refer to the process by which meaning evolves due to the different ways in which a concept, a claim (or a word, or any sign or ontological entity) is used, not that users explicitly debate that concept – although in scholarly discourse this is of course also very important.

ClaiMaker enables scholars to annotate research publications with their interpretations of the significance of contributions. The system allows the user to associate documents with concepts (*primary claim*), and to then make connections between these concepts (*secondary claim*) to create a semantic triple (Figure 2).³ However, unlike more traditional ontological knowledge representation systems, the focus is primarily on relational semantics: there is no formally defined concept class hierarchy (although one might ‘plug in’ a domain taxonomy or ontology to constrain concept typing if desired).

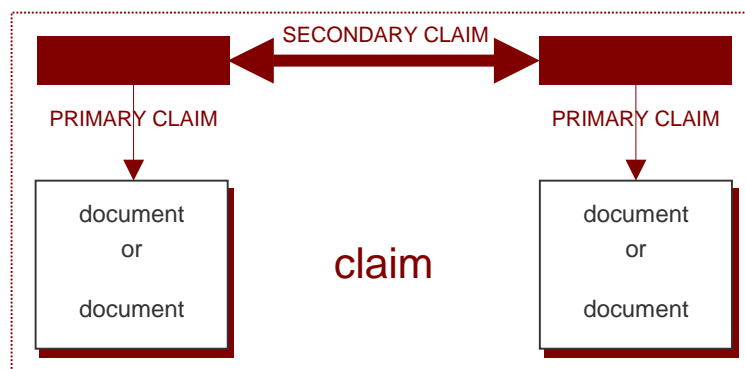


Figure 2. Schematisation of a ClaiMaker’s claim structure.

In order to link concepts via secondary claims, the user is provided with a taxonomy of argumentative relations, consisting of base relational classes articulated into more specific ‘dialect’ links identified by labels that scholars and researchers understand and use. Links are defined by their *type* (the base relational class to which they belong), and *polarity* (indicating whether the relation is positive or negative) (Table 1).

The motivation for this approach is that all scholarly fields make use of the base relational categories, but the ‘dialect’ of the field, that is, the language in which claims are made, substantiated and contested varies across discipline. Table 1’s dialect is the default provided, but can be modified without affecting computational services that operate on the base relational classes.

We see the process of annotating research documents with concepts (nodes in the network quoting, paraphrasing or interpreting a piece of literature – see §5.1 for examples), as a process of making claims – connections between ideas – about the different roles that a publication’s contributions should play within the community’s discourse. In this sense, ClaiMaker is

³ In ClaiMaker, a secondary claim can also connect sets. Sets can contain arbitrarily large numbers of member concepts (which may be concepts, sets, or claims). Each set can be treated by the system as a single, although complex, object and related to any other object via a secondary claim.

designed to mediate the ongoing process of scholarly discourse itself, not by modelling any particular scholarly domain, but by modelling the relatively stable discourse moves by which knowledge claims are made in most research fields.

Relation Classes	Relations	Polarity
GENERAL	Is-about	+
	Uses/applies/is-enabled-by	+
	Improves-on	+
	Impairs	-
PROBLEM RELATED	Addresses	+
	Solves	+
SUPPORTS/CHALLENGES	Proves	+
	Refutes	-
	Is-evidence-for	+
	Is-evidence-against	-
	Agrees-with	+
	Disagrees-with	-
	Is consistence-with	+
	Is inconsistent-with	-
TAXONOMIC	Part-of	+
	Example-of	+
	Subclass-of	+
SIMILARITY	Is-identical-to	+
	Is-similar-to	+
	Is-different-to	-
	Is-the-opposite-of	-
	Shares-issues-with	+
	Has-nothing-to-do-with	-
	Is-analogous-to	+
	Is-not-analogous-to	-
CAUSAL	Predicts	+
	Envisages	+
	Causes	+
	Is-capable-of-causing	+
	Is-prerequisite-for	+
	Is-unlikely-to-affect	-
	Prevents	-

Table 1. ClaiMaker’s link taxonomy with the category of ‘polarity’.

Let us now return to the semiotic perspective introduced earlier, and analyse ClaiMaker as a semiotic system, that is, as a ‘language’.

5 ClaiMaker as a ‘language’: a semiotic and cognitive perspective

Recall that a ‘claim’ in ClaiMaker is either a primary or secondary claim. The secondary claim consists of the relation established between two concepts. The secondary claim presupposes two primary claims, which in turn consist of the relation between each of the two concepts and their respective sources (whether this is an electronic document or a claim previously created in the system). For instance, in one’s interpretation, the paper *a* by *A* may express the concept “*Beauty responds to universal parameters such as the Golden Ratio*”, while the paper *b* by *B* may express

the concept “*The growth of most natural organisms appears to be regulated by the criterion of the Golden Ratio*”. Provided that they are explicitly referred to their respective sources, identified by bibliographic references (primary claim), these two concepts can be related to each other (secondary claim). In this way, the secondary claim {“*Beauty responds to universal parameters such as the Golden Ratio*” ‘is-consistent-with’ “*The growth of most natural organisms appears to be regulated by the criterion of the Golden Ratio*”} does not simply relate two unanchored concepts. It relates two identified sources from a particular point of view, for which a specific user has taken responsibility. This interconnection between specific sources and individual interpretations is fundamental for the construction of the collective discourse that ClaiMaker was designed to represent. The following semiotic analysis shows how, in ClaiMaker’s ‘language’, primary claims, secondary claim, and their related objects (the two concepts and their respective sources) constitute the essential components of a minimal syntagmatic unit (Figure 1 and 3): ClaiMaker’s basic discourse act.




Figure 3. Example of a ClaiMaker claim.

5.1 ClaiMaker and primary claims

In the following sections we describe how primary claims work both at the level of referential relations and at the level of concept classification. We explain what function they play in the system and why they are so important from a semiotic point of view.

5.1.1 Primary claims and referential relations

Let us start by analysing the primary claims. In semiotic terms, a primary claim consists of establishing a referential relation between a concept and a referent. The concept is constituted by a *signifier* – the text by which the concept is expressed – and its *signified* – the conceptual content expressed by the text. The referent is constituted by a part or an aspect of an electronic source – whether a document or a pre-existing claim. From a semiotic point of view, the part or aspect of the source that is being referred to constitutes what the concept stands for, according to the interpretation of the user who created the concept, which in turn expresses the user’s perspective on the

source itself. At the same time, while providing *evidence* in support of the primary and then the secondary claim, the source as a whole constitutes the *context* in which the part or aspect that is being referred to is originally used. Referring to our introductory example, just as dogs play different roles in different cultures and even in different people’s experience, likewise throughout scientific communities and between different scholars, literary sources are perceived and used in different ways, depending on cultural backgrounds and research interests. This means that different scholars – at different times, with different backgrounds, research goals and perspectives – may describe different things in the same literary source, as they perceive or focus on different aspects of it. That is, the concept that they ‘extract’ from the source stands for the source from a particular perspective, and is the expression of a particular contextualised reading (Figure 4).

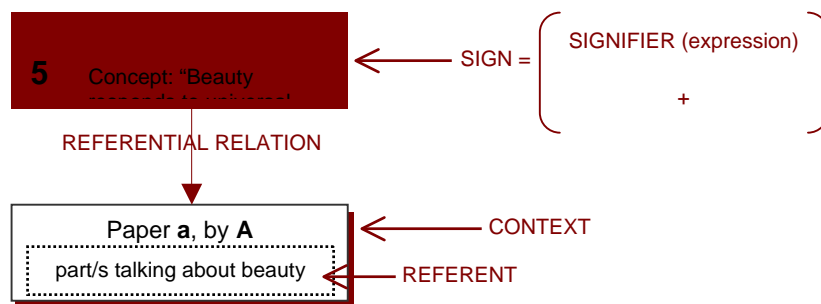


Figure 4. Semiotic analysis of a ClaiMaker’s primary claim.

In other words, making a primary claim is a semiotic operation, which consists of creating a *sign* (the concept) that refers to (stands for) a particular referent (the document or pre-existing claim) in the virtual reality (the ClaiMaker repository), in some respect (in a specific context). In our example, as our user has claimed that **B**’s paper expresses the concept “*The growth of most natural organisms appears to be regulated by the criterion of the Golden Ratio*”. However, another user might claim about the same paper that it expresses the concept “*The development of many living beings shows how the Golden Ratio plays a special role in nature*”. Finally, a third user might claim that the paper expresses the concept “*There is no definite evidence that the Golden Ratio regulates the development of natural organisms*”. In fact, all three concepts are compatible with the source, but while the first and the second concepts have a similar signified, the third one seems to signify something else, expressing a quite different perspective on the same source. And this difference is such that, while – in a secondary claim – the first and the second concepts can both easily be claimed to be *consistent-with* “*Beauty responds to universal parameters such as the Golden Ratio*” (from **A**’s paper), the third concept seems rather inconsistent (Figure 5).




Figure 5. An example of different and even contradictory claims anchored in the same sources (referents).

Consequently, the meaning of a source changes from one primary claim to the other, depending on the perspective from which it is looked at, that is, depending on the part or aspect of it that is directly referred to, and on the way in which this part or aspect is interpreted within its context (the source itself).

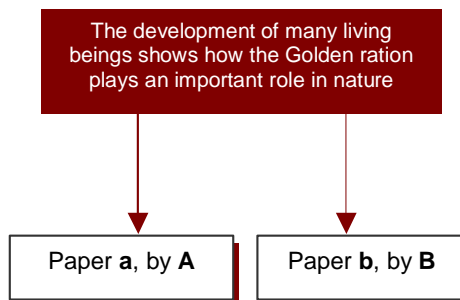


Figure 6. Example of how the same concept can be anchored in different sources.

On the other hand, a user might find that the concept “*The development of many living beings shows how the Golden Ratio plays a special role in nature*” can be annotated on (anchored in) *A*’s paper, as well as *B*’s paper (Figure 6). Just as the same human icon – always referring to the human being – can be used to signify either “*rest rooms*” or “*pedestrians*”, depending on the context in which it is used, so can a ClaiMaker concept have different meaning, if it is referred to different sources. The user who created the concept must have referred it to a part or an aspect of *A*’s paper, but the user who decides to re-use it may find that it can be also referred to a part or an aspect of *B*’s paper. That is, this second user might want to use the same concept to create a new primary claim, and as a result there would now be two claims using the same concept but in reference to two different sources. This sets the concept, its signified and its direct referent within two different contexts, triggering a different series of connotative connections.

This means that the overall meaning of the concept changes from one primary claim to the other, depending on what source its direct referent (a part or aspect of the source itself) is taken from, that is, depending on what context the referent belongs to.

In other words, since in different primary claims different concepts can refer to the same source, and the same concept can refer to different sources, the referential *relation* between the concept and its referent is extremely powerful in determining its meaning. On the one hand, the referential relation expresses the fact that different primary claims may involve the same source in different ways, that is, the same source is looked at from different perspectives. On the other hand, the referential relation expresses the fact that the same concept may refer to different sources, that is, the same concept signifies different things in different contexts.

So, in ClaiMaker, the role and the meaning of each *source* within the system's discourse space are first of all defined at the level of the primary claims by the relations that the source itself holds with all of the concepts that refer to it. Only in the second instance are the role and meaning of a source defined at the level of the secondary claims, by the relations that the source itself holds with other sources through the discourse relations that users establish between concepts.

Similarly, both the role and meaning of each *concept* within the discourse space are first of all defined at the level of the primary claims, by the referential relations holding between the concept itself and all of the sources that it refers to. And only in the second instance are the role and meaning of a concept defined at the level of the secondary claims, by the relations that users establish between concepts.

As we have seen, in any semiotic system a sign is made of a signifier and a number of signifieds: the denotative signified, plus a number of connotations grounded in the interpreter's particular knowledge of the referent and in the context in which the sign is used. In ClaiMaker, a concept is like the Saussurian sign whose signifier corresponds to the concept's *expression* (namely, the sentence that expresses it), and whose signified(s) corresponds to the concept's *content* (namely, the content of the sentence) (Hjelmslev, 1959).

When reading the sentence written by a user to express a concept, the members of his own scientific community will tend to recognise in it the same denotative signified as well as the same set of connotative signifieds. This is because they are familiar with the same literature and share the same perspective on it. However, users from different communities or from different disciplines altogether, while reading into the same sentence the same denotation, will most likely read into it a different set of connotations. In some cases, they may not even be able to go beyond the mere

interpretation of the wording. Since it is impossible to codify all the possible connotations that a concept may be associated with, the only way in which the various connotations of a concept can be represented at all in the discourse space is through the referential connections between the concept itself and all its referents. At the same time, the way in which each piece of literature is interpreted and used can only be represented in the discourse space through its referential connections with the set of concepts by which it is referred to.

5.1.2 Primary claims and concept classification

Within a specific discourse community, members tend to use concepts in a more consistent manner than one would expect between communities. For instance, the concept of “*ontology*” means one thing to a philosopher and something rather different to semantic web researcher, although a minority of members in both communities uses, and explicitly discusses, the term in both senses. Similarly, the concept of “*system*” tends to mean one thing to computer scientists, something broader to a systems analyst, and something different again to a systems theorist. If a majority of users use a concept in the same way, that particular use expresses the class’s primary signified within that community’s discourse space, and it constitutes the basis for its collective usability within the community itself. However, users who come from different disciplines, do interdisciplinary work, or do not share the same backgrounds, research goals and perspectives as the rest of the community, may well use different classificatory criteria. Therefore they may associate with the same class concepts that for most users belong to different classes. These uncommon uses express the class’s secondary signifieds, that is, alternative ways of interpreting the concept class (Figure 7).




Figure 7. A concept class can be used in different ways across different scholarly communities: the darker areas represent the particular use that is likely to be made of a concept class in different contexts.

ClaiMaker seeks to reflect these semiotic processes of denotation and connotation in the way in which it handles concept classification. Concepts can optionally be classified (for instance, as *problem*, *method*, *theory*).

However, in contrast to normal ontology engineering, we are not concerned to codify – and hence freeze – the meaning of these classes with slots, subclasses or constraints on use.⁴ Instead, the meaning of each conceptual class is defined by the way in which it is used to classify different concepts, each of which derives its meaning from the ‘set’ of the referential relations in which it participates (both primary and secondary).

Thus, the concept “*There are six versions of String Theory*” might be classified as a *problem* to be solved by one person, and treated as *evidence* by someone else (or even by the same person), depending on the roles played by this concept in the discourse move being made. This is a fundamental requirement in a modelling approach which seeks to recognise the process of semiosis that takes place in discourse. In natural language, as in other languages whose evolution happens in a relatively spontaneous way (subject to cultural, psychological and physical conditioning), meaning is continually negotiated through use. If a particular linguistic expression starts to be used in a new way by a sufficient number of people, that expression acquires a new connotation⁵. Eventually, the new signified may become so widely accepted and predominant that it may end up establishing itself as a new denotation⁶. Obviously, the more influential the group of people who start to use a linguistic expression with a new connotation, the more negotiating power they will hold to establish that connotation in the collective discourse space, which may in turn establish new denotations, in a never ending generative flow. For this to happen, though, the relation between signifier and signifieds, as well as the relation between signifieds and referent, needs to be relatively open to new interpretations and uses, while favouring the stability of existing interpretations as are already used by the majority.

In ClaiMaker, these dynamics find expression in the first place at the level of primary claims, where meaning is continually negotiated via establishing referential relations between referents and concepts, and via defining concepts according to different classes. This constitutes a decisive difference with respect to other ontology-based systems. Let us turn now to the affordances of secondary claims, with respect to the process of semiosis.

⁴ The user is can choose to leave concepts unclassified, can select from a predefined menu, or can label them as ‘Other’ and enter his own class. The menu of classes offered is entirely customisable (although currently only by a system administrator). Search services depend only on base relational classes and in some cases on dialect terms, but not on concept classification.

⁵ Metaphor is a very good example of this phenomenon. Think of the word ‘dog’ in the expression “The soprano was a dog. She kept singing false notes!” (Eco, 1968).

⁶ Catachresis is a very good example of this phenomenon. Think of the expression “table legs” commonly used to refer to a table top’s supports (Metz, 1977).

5.2 ClaiMaker and secondary claims

Based on the way in which primary claims work, secondary claims play a fundamental role in the construction of dynamic and computable discourse structures in ClaiMaker. It is at the level of secondary claims that different document sources are semantically connected⁷ and scholarly perspectives take shape. In the next sections we explain the theoretical motivation for the system's relational scheme, and the way in which users are allowed to make use of it. We also describe the functionalities and services that this enables. Finally, we explain how, in ClaiMaker, discourse relations can work at two distinct but interconnected levels – a user interface level and a computational level – and the advantages this offers to the user.

5.2.1 Discourse relations and coherence

As introduced above, a secondary claim consists of establishing a discourse connection between two concepts, each in turn anchored to one or more respective referent documents. We have described how ClaiMaker provides its users with a relational language, a taxonomy consisting of base relational classes for any particular community to articulate into a dialect according to its own discourse conventions. The central idea on which we base this approach is that discourse in all modalities is “coherent” because it makes moves that can be circumscribed by a set of relational primitives that have been shown to have cognitive reality. These can in turn take a number of forms, each of which can be expressed in different ways in different communities.⁸

This approach has theoretical and empirical grounds in linguistic research (Mann and Thompson, 1988; Knott and Dale, 1994; Knott and Mellish, 1996; Sanders et al., 1992; 1993; Pander Maat, 1999; Sanders and Noordman, 2000; Sanders and Spooren, 2001), which takes a cognitive approach to the problem of text coherence. From ClaiMaker's perspective

⁷ Research documents are of course already explicitly connected via citations, and can be further connected via secondary means such as co-citation analysis, but these have no semantics (although it is known that citation is generally positive in nature). Other techniques in scientometrics and information retrieval are complementary to the semantic annotation paradigm, able to suggest clusters for human validation. However, these are unable to extract argumentation structures, and a distributed semantic annotation paradigm such as ours is unique in being able to mediate abstract interpretations, and the forging of connections between documents which none of the original authors may have seen.

⁸ In SDRT (Asher and Lascarides, 2003) the distinction between less specific (higher level) and more specific (lower level) relations is also made (Asher et al., 2004). However, this work still importantly relies on the specification of domain semantics, whereas our goal is to facilitate the *modelling discourse* through which the *modelling product* – a consensus worldview – is negotiated.

on discourse construction and representation, what makes the cognitive approach especially appealing is that it seeks to identify principles of cognition capable of accounting for the connections that we perceive to hold between elements, in a ‘coherent’ mental representation. If these principles inform our psychological constructs, they also inform our discourse constructs, and therefore offer candidate relational primitives for ontology construction (Mancini and Buckingham Shum, 2001; 2004). We therefore explore in some detail the basis for what might be framed as an ‘upper level discourse relations ontology’ (for a more detailed discussion see Mancini, 2005).

One of the first steps in this direction, is represented by Mann and Thompson’s (1988) Rhetorical Structure Theory (RST), one of the most cited and applied in the field of computational linguistics and natural language generation. This is an analytical framework, applicable to written monologue, which allows an analyst to describe rhetorical relations between text spans in functional terms, and to identify the text’s hierarchical structure. However, linguists within the cognitive approach have identified important limitations to RST, specifically, that RST does not provide a method for specifying coherent relations, treating them as purely descriptive constructs, unable to provide any real insight into the working principles underlying text interpretation and interpretability (Knott and Dale, 1994; Knott and Mellish, 1996; Sanders et al., 1992; 1993).

In contrast to Mann and Thompson, these linguists tend to use the cohesive devices of language as evidence for a psychological theory of text coherence. The rationale behind this is that the structures of language are optimised for our communicative purposes, allowing the expression of our psychological reality. Language must therefore provide devices to mark our psychological constructs and these devices can be used as evidence for the identification of a set of cognitively-derived coherence relations. This approach allows the authors to define, on an experimental basis, a much more rigorously defined set of relational parameters, which (like Mann and Thompson’s approach) is able to account for the hierarchical structure of texts.

From our point of view, however, the work of Sanders, Spooren and Noordman (1993) – which directly informs Knott, Dale and Mellish’s work (Knott and Sanders, 1998) – seems to be more promising. On the line of Hobbs’ pioneering work (Hobbs, 1985), the authors take a cognitive approach to the problem of text representation and they also treat coherence relations as psychological constructs. Like Hobbs, but in contrast to Knott et al., Sanders, et al’s approach is theory driven, rather than data driven: they advance a psychologically plausible hypothesis which they are then able to support with experimental data. Like Knott and Mellish, they propose a set of parameters describing cognitive coherence relations, but, unlike Knott

and Mellish’s set, theirs consists of a small number of cognitively basic concepts. In principle their set is not subject to proliferation because it is based on primitive cognitive categories, which are properties common to all relations. Their relational scheme is defined through the combination of four parameters (Table 2). These are proposed as dimensions for all relations and can take different values, so determining the characteristics of any relation derived from their combination (Table 3).

Primitive Concept	Possible Values	
<i>BASIC OPERATION</i>	additive	causal
<i>SOURCE OF COHERENCE</i>	semantic	pragmatic
<i>ORDER OF THE SEGMENTS</i>	basic	non-basic
<i>POLARITY</i>	positive	negative

Table 2. Sanders, Spooren and Noordman’s 4 cognitive relational primitives, isolated to describe discourse connections in text (from Sanders et al., 1993).

These four parameters, defined by the authors as Cognitive Coherence Relations, are as follows:

- *Basic Operation.* Two discourse segments can be weakly connected, in which case they are related by *additiveness*, or they can be strongly connected, in which case they are related by *causality*. An additive operation exists if the two discourse segments are simply related by logical conjunction, whereas a causal operation exists if between the two discourse segments an implication relation can be deduced. For example: in “*This morning I went shopping (and) nipped into the hairdresser*” the two discourse segments are related by additiveness. In “*This morning it was pouring with rain (so) I grabbed my rain coat to go out*” the segments are related by causality.
- *Source of Coherence.* A discourse relation can be *semantic*, when the two discourse segments are related on the basis of their propositional content, or *pragmatic*, when the two segments are related on the basis of their argumentative or rhetorical function. For instance: in “*It rained all day yesterday (so) the streets got flooded*” the relation between the fact that it rained and the fact that the streets were flooded is presented as the external observation of a cause-effect phenomenon. In “*The neighbours must be out tonight (because) their lights are all off*” the speaker infers a cause-effect phenomenon based on the assumption that, when people go out, they switch off the lights. In other words, the speaker uses the lights being off as an argument to claim that the neighbours are out (Sanders, 1997; Pander Maat and Degand, 2001).
- *Order of Segments.* Two discourse segments can be related in a *basic* order, which corresponds to the order in which events are

meant to actually happen in the described reality, or in a *non-basic* order, which does not correspond to the order of the described events. For instance: in “*I missed the bus this morning (so) I was late for the meeting*” the order of presentation of the events is basic, because it corresponds to the order in which they verify in the described situation. In “*I was late for the meeting (as) I missed the bus this morning*” the order of presentation is backward with respect to the event described and therefore the order of segments is non-basic.

- *Polarity*. The relation between two discourse segments can be *positive*, when the content of the two related segments consistently express the same basic operation, or *negative*, when the content of one of the two segments defies the rule of the basic operation expressed by the other segment. For instance: in “*She did not sleep all night (so) in the morning she looked shattered*” the polarity is positive, since the event described in the second segment is a consistent consequence of the event described in the first segment. In “*She looked as fresh as a rosebud (although) she had not slept all night*” the polarity is negative, as the expected consequence does not follow (Knott, 1998).

From the combination of these four parameters, a relational hierarchy can be derived, as summarised in Table 3.

Basic Operation	Source of Coherence	Order of Segments	Polarity	Class	Relation
causal	semantic	basic	positive	1a	<i>cause-consequence</i>
				1b	<i>condition-consequence</i>
causal	semantic	basic	negative	2	<i>contrastive cause-consequence</i>
causal	semantic	non-basic	positive	3a	<i>consequence-cause</i>
				3b	<i>consequence-condition</i>
causal	semantic	non-basic	negative	4	<i>contrastive consequence-cause</i>
causal	pragmatic	basic	positive	5a	<i>argument-claim</i>
				5b	<i>condition-claim</i>
causal	pragmatic	basic	negative	6	<i>contrastive argument-claim</i>
causal	pragmatic	non-basic	positive	7a	<i>claim-argument</i>
				7b	<i>claim-condition</i>
causal	pragmatic	non-basic	negative	8	<i>contrastive claim-argument</i>
additive	semantic	-	positive	9	<i>list</i>
additive	semantic	-	negative	10a	<i>opposition</i>
				10b	<i>exception</i>
additive	pragmatic	-	positive	11	<i>enumeration</i>
additive	pragmatic	-	negative	12	<i>concession</i>

Table 3. Sanders, Spooren and Noordman’s taxonomy resulting from the combination of the four parameters, and the prototypical relations for which the taxonomy accounts (table from Sanders et al., 1993).

In related proposals, such as Louwrese’s (2001) analytic and cognitive parameterisation of coherence relations, relational concepts such as similarity and contrast are also accounted for, respectively as positive and negative additive relations. That is, *additive* relations connect two events or situations on the basis of some sort of equivalence, but the nature of this equivalence can be either conjunctive or comparative. In the first case, additive relations indicate the joint relevance of two situations, events or objects with respect to a whole or larger picture. For instance, in the sentence “*I need to clean the kitchen (and) I need to dust the sitting room, before the guests arrive*”, the equivalence is set with respect to the joint relevance of the two actions in order for the house to be ready to receive guests. In the second case, however additive relations also indicate the similarity between the connected situations, events or objects (Pander Maat, 1999). For instance, consider the sentence “*The trajectory of a projectile is determined by inertia, which makes it fly forward, and by gravitation, which makes it fall back onto the ground. (Likewise) the trajectory of a planet around another planet is determined by inertia, which makes it move forward, and by gravitation, which makes it deflect from a rectilinear motion*”. Here the equivalence is established by comparison between two situations that present structural similarities.

Type	Polarity	Direction	Examples
causal	positive	backward	<i>A because B</i>
		forward	<i>A so B; because A, B</i>
		bi-directional	-
	negative	backward	<i>A although B</i>
		forward	<i>A nevertheless B; although A, B</i>
		bi-directional	-
temporal	positive	backward	<i>A before B; after A, B</i>
		forward	<i>A after B; before A, B</i>
		bi-directional	<i>A while B; B while A</i>
	negative	backward	<i>A until B</i>
		forward	<i>until A, B</i>
		bi-directional	-
additive	positive	backward	-
		forward	<i>A moreover B</i>
		bi-directional	<i>A similar B; B similar A</i>
	negative	backward	-
		forward	<i>A however B</i>
		bi-directional	<i>A alternatively B; B alternatively A</i>

Table 4. Louwrese’s analytic and cognitive parametrisation of coherence relations, derived from those categories that are represented in most text coherence theories (table taken from Louwrese, 2001).

As far as the primitive *order of segments* is concerned, Louwrese prefers to talk about *directionality* and, unlike Sanders, Spooren and Noordman, he envisages for this parameter three possible values: *forward*, *backward* and

bi-directional. For instance, among the additive relations, conjunctive relations would be forward, whereas comparative relations would be bi-directional; among the causal relations, those in which the cause is presented first are forward, whereas those in which the cause is presented second are backward (Table 4).

5.2.2 Cognitive coherence relations and ClaiMaker's relational ontology

Drawing from both Sanders, Spooren and Noordman, and Louwerse's descriptive frameworks, ClaiMaker's taxonomy derives from a combination of the four primitive parameters that define Cognitive Coherence Relations: *basic operation*, *polarity*, *order of segments (directionality)*, and *source of coherence*. The *basic operation* (or *relation type*, in Louwerse's terms) is expressed by the fact that the links are grouped under different categories in which both causal and additive (conjunctive and comparative) relations are represented. The *source of coherence* (only accounted for by Sanders et al.) is expressed by the fact that within the same group, semantic and pragmatic relations can both be found in the taxonomy. The *polarity* is expressed by the fact that each link is qualified as positive or negative. The *order of segments* (or *direction*, in Louwerse's terms) is expressed by the fact that links between concepts can be forward, backward or bi-directional.⁹ Table 5 shows and motivates the organisation of ClaiMaker's relational scheme according to a CCR parameterisation.

Structuring and using a link taxonomy based on the parameterisation of cognitive primitives offers a number of representational advantages. First of all, this grounds the taxonomy in what – from experimental evidence – appears to be psychological reality, which in principle gives the taxonomy stability and applicability across different disciplines, media and discourse types. For instance, as various analyses have shown, just like the discourse relations holding between hypertext nodes, the discourse relations holding between the shots of cinematic sequences can be described in terms of CCR (Mancini, 2005). In principle, CCR can be used to describe discourse connections as cognitive relations in any language or linguistic form based on the articulation of discreet semantic units (such as text, cinema, hypertext, hypermedia, etc.).¹⁰

⁹ At present bi-directional relations are not implemented in ClaiMaker.

¹⁰ In a future scenario we are developing, ClaiMaker's users will be able to annotate claims directly onto any media document, using the discourse relations within the claim network to generate coherent argumentative sequences across the media (see also Future Work §7). Using a CCR-based taxonomy means that its relations will be able to account for connections holding between discourse units.

Basic Operation	Polarity	Source of Coherence	Order of Segments	Link Name	CCR Equivalent	Motivation	
Additive Comparative	Positive	Semantic	Bi-directional	<i>Is-identical-to</i>	Bi-directional positive semantic additive in the form of similarity.	It is the highest degree of similarity.	
				<i>Is similar-to</i>	Bi-directional positive semantic additive in the form of similarity.	It is a simple similarity.	
				<i>Shares-issues-with</i>	Bi-directional positive semantic additive in the form of similarity.	It is a mild degree of similarity.	
				<i>Is-analogous</i>	Bi-directional positive semantic additive in the form of similarity.	It expresses a similarity and a successful comparison.	
Additive Conjunctive		Pragmatic	Backward	<i>Agrees-with</i>	Bi-directional positive pragmatic additive.	It adds elements to the same view.	
				<i>Is-about</i>	Forward positive pragmatic additive in the form of elaboration.	It indicates elaboration of something being presented or mentioned.	
Additive Comparative		Negative	Semantic	Bi-directional	<i>Is-different-to</i>	Bi-directional negative semantic additive in the form of contrast.	It expresses the negative result of a comparison.
					<i>Is-the-opposite-of</i>	Bi-directional negative semantic additive in the form of contrast.	It expresses the highest degree to which a comparison can fail.
	<i>Has-nothing-to-do-with</i>				Bi-directional 'zero degree' semantic additive in the form of similarity or contrast.	It expresses a case in which two objects are not even comparable.	
	<i>Is-not-analogous</i>				Bi-directional 'zero degree' semantic additive in the form of similarity, or semantic additive contrast.	It is either a non-similarity or the negative result of a comparison.	
Additive Conjunctive	Pragmatic		Backward	<i>Disagrees-with</i>	Bi-directional negative pragmatic additive.	It adds contrastive elements to a view.	
Causal	Positive		Semantic	Backward	<i>Uses/applies/is-enabled-by</i>	Backward positive semantic causal. Backward positive pragmatic causal in the hypothetical form (conditional).	It expresses the result of a cause or condition that makes it possible for such result to happen.
					<i>Improves-on</i>	Backward positive semantic causal.	Similar to solution-problem.
					<i>Addresses</i>	Backward positive semantic causal.	Similar to solution-problem.
		<i>Solves</i>			Backward positive semantic causal.	Similar to solution-problem.	
		Forward		<i>Causes</i>	Forward positive semantic causal.	It causes something to happen.	
				<i>Is-capable-of-causing</i>	Forward positive semantic causal in the hypothetical form (conditionality).	It constitutes a sufficient condition.	

				<i>Is-prerequisite-for</i>	Forward positive semantic causal in the hypothetical form (conditionality).	It constitutes a necessary condition.
				<i>Impairs</i>	Forward positive semantic causal, associated with opposition.	It causes something not to happen. Or it obstacle its happening.
				5.2.2.1.1.1	Forward positive semantic causal, associated with opposition.	It causes something not to happen.
		Pragmatic	Backward	<i>Proves</i>	Forward positive pragmatic causal.	Because of it, something can be assumed or claimed.
				<i>Refutes</i>	Forward positive pragmatic causal, associated with opposition.	It supports or proves that something cannot be claimed.
				<i>Is-evidence-for</i>	Forward positive pragmatic causal.	It supports an assumption or claim.
	<i>Is-evidence-against</i>			Forward positive pragmatic causal, associated with opposition.	It supports the negation of an assumption or claim.	
	<i>Is-consistent-with</i>			Forward positive pragmatic causal. Forward positive pragmatic causal, associate with enumeration.	It motivates an assumption or claim. Or it adds to something else that motivates an assumption or claim.	
	Forward			<i>Predicts</i>	Forward positive pragmatic causal.	It is an assumption based on present signs.
		<i>Envisages</i>	Forward positive pragmatic causal.	It is an assumption based on present signs.		
	Negative	Semantic	Forward	<i>Is-unlikely-to-affect</i>	Forward negative semantic causal.	It does not cause anything.
		Pragmatic		<i>Is-inconsistent-with</i>	Forward negative pragmatic causal.	It fails to motivate an assumption or claim.

Table 5. CCR-based description of ClaiMaker relational taxonomy.

Secondly, a CCR-based taxonomy ensures that the main categories of this psychological reality are represented in the link set, which in turn ensures a more balanced expression of different kinds of connection. Developing a customised taxonomical dialect from a few primitive but cognitively exhaustive relational parameters provides a framework to ensure that those cognitive dimensions are adequately represented in their taxonomy, to make informed decisions about inclusion or exclusion of certain cognitive dimensions.

Thirdly, at any level of articulation and specialisation, a CCR-based approach ensures that the taxonomy's links are accountable for by a small number of primitives, which allows for consistent discourse modelling, processing and searching at very different levels of granularity. Consider the case in which a user searches for all concepts which hold with discourse unit **X** the rhetorical relation *is-prerequisite-for*. Even if there are no direct matches to this dialect-specific label, the system could search at the CCR-level for discourse units related to **X** by any *causal positive semantic forward* connections, such as *is-capable-of-causing* and *causes* (Table 1). In general, basing the definition of ClaiMaker's taxonomy on CCR parameters means that the same underlying set of relational parameters can be employed by different research communities, speaking different dialects, by changing the labels of the relations, without changing the underlying functionality of the system, based on the same cognitive structure.

Let us now consider ClaiMaker's relation-centric services that have been implemented to date.¹¹ *Lineage* is essentially ancestry and (with its inverse, the descendant) focuses on the common notion in research that ideas build on each other. ClaiMaker's *lineage* tool semantically tracks back from a node to see how it evolved, whereas the *descendants* tool tracks forward from a node to see what new ideas evolved from it. To provide lineage analysis as a ClaiMaker service, path queries are constructed from link-types using a set of primitives. For example, one can search for paths that may be of any length, and which contain (in any order) any of the positive links that have type *similarity* in either direction, or the two general links *uses/applies/is-enabled-by* or *improves-on*, going in the direction away from the target node of the query. The *improves-on* link type is included to reflect the notion of progress implicit in lineage, while *uses/applies/is-enabled-by* has a weaker implication of 'building upon'. In CCR terms these are both *positive semantic causal* relations: in the first case, one phenomenon causes its own improvement by the other in the same way in which a problem calls for being given a solution; in the second case, one phenomenon is a direct cause or condition for the other to take place. The *similarity* links – which constitute *positive semantic comparative additive* relations in CCR terms – are included because if a new node **A** is like another **B** that *improves-on* a **C**, then **A** may well also be an improvement on **C**. *Similarity* links are acceptable in either direction because comparative relations are bi-directional (if **A** is like **B**, then **B** is like **A**). Summarising, from the CCR viewpoint, the functionality of *lineage* and *descendants* needs to always

¹¹ In this paper we concentrate on the theoretical ground on which ClaiMaker is based, therefore we do not go into detail when describing the services offered by the system. For a more detailed description of the system's technical functionality see Buckingham Shum et al. (In Press).

follow positive relations, and they need to be either *causal* or *comparative*: either they denote a step forward along a development line, or a convergence across different lines.

Together with the *lineage analysis*, ClaiMaker also offers the service of the *perspective analysis*, which is based on the contraposition of discourse units related by *additive positive conjunctive or comparative* relations, on the one hand, and by *additive negative conjunctive or comparative* relations, on the other. In the future, more search services can be designed and implemented, which exploit the CCR-based parameterisation of ClaiMaker's relational taxonomy, as discussed in the concluding section on future work.

5.2.3 The two levels of ClaiMaker's relational language

When we started to develop ClaiMaker, we iterated through versions of the link taxonomy through a combination of intuition as professional researchers, informed by empirical data from our literature modelling. This data-driven approach consisted of modelling claims and argumentation as found in a range of research domains, including computer supported collaborative work, text categorisation, and literary criticism. Relations common to several domains were identified that are used in argumentation practice. We found that they could be classified into relation types with similar rhetorical implications: *Supports/Challenges*, *Problem Related*, *Taxonomic*, *Causality*, *Similarity*, and *General* (as shown in Table 1).

As our semiotic and CCR analysis developed in parallel and overlaps in some of the base relational classes emerged (as, for instance, *Causality* and *Similarity*), we started to refer to the CCR parameters more consistently while refining/defining our taxonomy. Consistent with what we have subsequently come to understand through the lens of CCR theory, we found that some relations occurred in pairs of opposites, as *proves* and *refutes*, where one has positive and the other negative implications. At the time, we termed this property "polarity" intending the more everyday meaning, although polarity in CCR terms has a different meaning, including both negative relations and positive relations associated with negative unary relations.

Such examples made us aware of discrepancies between the everyday meanings of rhetorical relations as they appear in our user interface, and relational categories derived purely from low-level CCR-parameters which are the product of academic study, as reviewed. As we now discuss, reflection from a CCR perspective has led us to the view that relational classes defined strictly from CCR parameters are not intuitive for untrained users to work with, hence the need for presenting the taxonomy in a way that is more consistent with a common-sense categorisation.

For instance, all the links defined in the CCR-based categorisation as *pragmatic causal* relations (both positive and negative), as well as two *pragmatic additive* relations (*agrees* and *disagrees-with*) are grouped in the interface set under the category *supports/challenges*. This is because it is more intuitive to think of a concept **A** as supporting or challenging another concept **B**, than it is to think of a concept **A** as holding a pragmatic causal relation with a concept **B**. Likewise, agreement and disagreement are commonly perceived as supportive or challenging phenomena, although they are not: agreement in itself does not constitute support, but it expresses similarity of perspective or opinion. For instance, if my friends agree with me that “*The president of that country is not a good leader*”, strictly speaking, our agreement does not support our opinion that he really is a bad leader. However, if one of us provides good arguments to motivate such a claim, or better, provides evidence that the president is a bad leader, then the claim is supported.¹²

With respect to the causal semantic relations, they are for the most part simply grouped under the category of causal relations, with four exceptions: *solves* and *addresses* (intuitively presented as *problem-related* links) and *impairs* and *is-enabled-by* (presented under the category of *general* links). The motivation for grouping relations as *general* links was again usability and learnability: they could not be intuitively classified elsewhere.

Another discrepancy between a purely CCR-driven organisation of the link taxonomy and the organisation of ClaiMaker’s taxonomy concerns the classification of positive and negative relations. For instance, in the interface set, links like *impairs* and *prevents*, or *refutes* and *is-evidence-against* are classified as negative relations, although in formal CCR analysis they are positive. Saying that **A** *impairs* **B** is equivalent to say that **A** *causes* **B** not to happen, which means that a relation of causation does hold between the two (the negative form would be ‘**A** does not cause **B**’). Saying that **A** *refutes* **B** is equivalent to saying that **A** proves not-**B**, which again means that **A** actually does prove something (different from saying ‘**A** does not refute **B**’). **A** *is-evidence-against* **B**, means **A** is evidence for not-**B** (different from ‘**A** is not evidence for **B**’). Once again, the argument for classifying these relations as negative is that intuitively the user would perceive them as negative based on their semantic content: *impairs* being perceived as the negative of *causes* (when in fact it is its opposite); *refutes* being perceived as the negative of *proves*, etc. A claim that **Data X** *refutes* **Hypothesis Y** is not seen as a positive link to Y in common-sense terms.

¹² In other words, an argument or a piece of evidence can support a claim or a position, however, agreement or disagreement have to be associated with an argument or with a piece of evidence, in order to support or challenge any claim.

Finally, there is no CCR-based equivalent to the *taxonomic* base relational class that appears in ClaiMaker, since it concerns mereological (*part-of*) and taxonomical (*example-of* and *sub-class-of*) relations. These are extremely common in conventional knowledge modelling, where they are used to describe the structure of objects within a domain. However, because they are discourse connectives, CCRs describe the interaction between events and entities. In other words, instead of describing the reality as a given, CCRs refer to it as a process.

To summarise, we have described how the everyday relational language used in the user interface of a practical modeling tool does not always map to the mappings motivated by a linguistic analysis of those relations, such as CCR. While ClaiMaker already has two tiers of relations in the user interface, with *dialects* being grounded in *base relational classes* that the reasoning engine can exploit, we are now investigating how to extend this relational architecture to three tiers:

Dialect(s) (User Interface link semantics)

Base Relational Classes (User Interface link classes)

Cognitive Coherence Relations (Abstract Semantics)

The system would manage the connections between these tiers, and exploit the CCR-layer in the information services provided (see future work).

6 Related work

ClaiMaker can be located in its technological context with respect to two research fields: computer-supported argumentation, and ontology-based annotation. After reviewing these, we briefly summarise other approaches which adopt a semiotic approach to knowledge representation.

6.1 Computer-supported argumentation

There is a long tradition of research into Computer-Supported Argumentation (CSA). The interested reader is referred to (Buckingham Shum and Hammond, 1994) for a review of the empirical evidence relating to the cost/benefit tradeoff of argument-mapping tools, and to (Buckingham Shum, 2003) for an update. A collection of argument visualisation approaches from diverse domains is presented in (Kirschner et al 2003), and also by (Andriessen, et al, 2003) in the specific context of learning. Reed and Norman (Reed and Norman, 2003) present a collection of work at the

intersection of argumentation theory and more formal CSA tools. The diversity of work is reflected in diverse research communities.¹³

Some CSA approaches are relatively informal, with a concern not to disrupt the informal, socially negotiated dimensions of real world discourse by imposing excessive constraints on users. Other approaches seek to introduce a higher degree of formalisation and constraint in order to add clarity to the moves that stakeholders are making, with the additional advantage that it makes possible finer-grained computational processing of the resulting representation.

Extensive though the CSA literature is, we believe that ClaiMaker's ontology is distinctive, based as it is on a core set of relational primitives grounded in semiotics and cognitive coherence relations theory. We are not aware of other systems that support the kinds of scholarly queries we have summarised, and detailed elsewhere (e.g. Uren et al, 2003).

6.2 Ontology-based annotation tools

Within the semantic web research community, there are numerous tools for annotating fragments of documents with terms grounded in one or more ontologies.¹⁴ However, CREAM (Handschuh and Staab, 2002) and similar tools support the annotation (or even automated extraction – (Vargas-Vera et al, 2002)) of *entities and facts* stated in the text; there is usually a 'correct' or 'ideal' result as to how the text should be annotated.

Ontological annotation of such entities complements the work that ClaiMaker is designed to support, which is annotating *interpretations of the significance of the ideas* in the target document. These might, therefore, apply to the whole document or whole sections. Indeed, they might not appear in the text at all if the reader has seen connections and connotations that the author did not intend. ClaiMaker's primary and secondary claims, therefore, enable the annotation of meaningful connections that are extremely difficult to detect automatically (consider a metaphorical connection that a reader makes between two papers in different fields, which have no citation or co-citation connections in the literature).

¹³ CMNA: International Workshop on Computational Models of Natural Argument:
<http://www.csc.liv.ac.uk/~floriana/CMNA6/>
COMMA: 1st International Conference on Computational Models of Argument:
<http://www.csc.liv.ac.uk/~comma/>
CSCL: The International Conference on Computer-Supported Collaborative Learning:
<http://www.isls.org/cscl.html>

¹⁴ Semantic Annotation Portal: <http://annotation.semanticweb.org>

In parallel with the development of ClaiMaker, the Trellis (Chklovski, et al 2005) and SEAS (Lowrance, et al, 2001) systems have been in development in the US, and have some similarities. Both are designed to assist the annotation of documents with statements which are then connected using argumentative relations in order to enable computational services (although the domains are less focused on scholarly discourse, and more on intelligence analysis and decision-support between competing options). However, as we have sought to demonstrate, ClaiMaker's relation-centric approach to reflect the diverse roles that concepts play reflects more deeply the processes of semiosis.

6.3 Semiotic approaches to knowledge representation

Others who take a semiotic approach to the analysis and definition of ontological frameworks for knowledge representation are, for instance, Sowa (2000), Liu (2000), and Gangemi (2004). Their work uses semiotics to enrich the articulation of ontologies in order to reflect the complexity of semiotic systems. However, whether semiotics is used to analyse the domains to be modelled, or to refine the entity and relation classes in the ontologies, this work remains, from our point of view, predominantly object-oriented and fundamentally concerned with the representation of knowledge domains.

Instead, we are concerned with representing the discourse about knowledge domains, rather than the domains themselves. Consequently, our work is predominantly relation-oriented, as it aims at representing the process of meaning negotiation rather than its products. Furthermore, with the choice of CCR as the base relational set, we take a minimalist approach to discourse modelling in order to be able to capture fundamental discourse moves across different media as well as different domains.

This approach acknowledges the importance of structure, which makes it possible to compute and share any representation, and envisages the use of some kind of formal language in discourse construction. However, it also recognises the impossibility of formalising every meaningful aspect of discourse, and instead of trying to capture 'everything' in what would be very complex yet still incomplete representations, it enables the user to take responsibility over the representation and interpretation process.

7 Conclusion and future work

All ontologies are the *product of a meaning negotiation process*, but one in which *the process itself is invisible*. The very purpose of creating an ontology is to articulate a consensus about 'how the world is' in some respect. If there is sufficient consensus this works until our understanding of the world changes. We are ourselves engaged in exactly the same enterprise to the extent that we are proposing a formalisation of how scholars make

discourse moves in the literature. The difference is that our discourse ontology is designed specifically to assist sensemaking and argumentation about how the world is, thus finding application in precisely those situations when systems that freeze a particular consensus are ontologically brittle: when the world changes. Our approach is ontologically brittle to the extent that it cannot model important discourse moves, not when there is a new perspective in a research field to integrate. Indeed, our claim is that scholarly claims about new perspectives are made precisely through the discourse moves that we have modelled.

This paper has reported the representational insights that have been gained into discourse modelling, by approaching the design of an argumentation-based, sensemaking-support tool through two theoretical lenses. Semiotics as a broad framework draws attention to the means by which meaning evolves in a discourse community, motivating modelling decisions that emphasise relation-centric semantics to reflect emergent denotations and connotations. Complementing this, Cognitive Coherence Relations provide a detailed framework that defines parameters for modelling these relationships.

A key part of our future work is to deepen the representational infrastructure. One aspect of this is investigating the modelling of complex constructs such as ‘school of thought’ in order to assist in extracting potentially significant patterns from the claims network. Preliminary results from this work are reported by Benn, et al. (2005).

A second aspect is to extend ClaiMaker’s current search and navigation services by explicitly modelling the parameter space defined by Cognitive Coherence Relations, moving towards a three-tiered relational architecture (§5.2.3). Consider the following two examples of how services could more fully exploit reasoning within a CCR-layer to infer connections that might not otherwise be made:

- The *General* relation *is-about* derives from the more abstract relation of *elaboration*, which establishes a relation between two discourse units (atomic or composite nodes in the data model), one of which has the rhetorical function of explaining, expanding or articulating the content of the other unit. In CCR terms, this is a *positive pragmatic additive* relation which has a lot in common with another positive pragmatic additive relation of a comparative nature: *agrees-with*, whose rhetorical function is reinforcing the content expressed in one discourse unit by adding more content expressing the same perspective. If a user was to search for discourse units that *are-about* discourse unit **X**, the system could suggest discourse units that *agree-with* unit **X** as potentially relevant.

- In another case, the user might be searching for discourse units that *prove* discourse unit **X**. This is a *positive pragmatic causal* relation between two discourse units, one of which has the rhetorical function of substantiating that the content of the other unit is true. However, like *proves*, *refutes* is also a *positive pragmatic causal* relation, but it is associated with a *negative* unary relation having *contrastive* function (**A** proves not-**B**, **A** disproves **B**). The system could reason that if a user wants to know what **X** proves, they may also be interested in knowing what **X** disproves, to expand the picture of the role that **X** plays in the scientific context.

The point of these two examples is that the constrained set of parameters provided by CCR defines a rich set of relationships which could, in principle, be used to generate a wide spectrum of inferences, possibly bridging between different dialects (discourse communities) and modalities. CCR, we argue, points towards an upper level ontology of discourse relations that could find wide application. We are now investigating the possibility of delegating certain classes of computational service to this upper relational level.

Finally, although this paper’s focus has been on discourse-modelling, this is in the service of our applied research goal to design practical, engaging tools. We have already shown how a purely theoretical focus to motivate relation-modelling is insufficient: it must also connect with what the end-users bring to the interaction. We are, therefore, exploring new interaction paradigms for ClaiMaker, and gathering data from user studies. Sereno et al. (2005) describe the rationale behind ClaimSpotter, a document-centric, semantic annotation interface to ClaiMaker, and report an empirical study analysing how untrained researchers make primary and secondary claims using our discourse taxonomy. This work has generated the first significant corpus of claim-authoring data (detailed in Sereno, forthcoming), showing how typical researchers engage with ClaiMaker’s dialects. In addition, Uren et al. (2006) have analysed novice and expert user behaviour with ClaiMaker and associated tools in the conduct of a literature analysis task, comparing this to performance with conventional tools.

8 Acknowledgements

We thank all of our colleagues who have participated in this project, bringing their unique contributions: Victoria Uren, Gary Li, Bertrand Sereno, Enrico Motta and John Domingue. We are grateful to Dnyanesh Rajpathak for useful discussions and Enrico Motta for valuable feedback. The reviewers of this article provided detailed, constructive comments that improved the paper. The Scholarly Ontologies project was funded by the UK EPSRC’s Distributed Information Management Programme (GR/N35885/01).

9 References

- Andriessen, J., Baker, M. and Suthers, D. (2003). (Eds.) *Arguing to Learn: Confronting Cognitions in Computer-Supported Collaborative Learning Environments*. Springer
- Asher, N., Denis, P., Kuhn, J., Larson, E., McCready, E., Palmer, A., Reese, B. and Wang, L. (2004). Extracting and Using Discourse Structure to Resolve Anaphoric Dependencies: Combining Logico-Semantic and Statistical Approaches. *Proceedings of TALN'04 (Conférence sur le Traitement Automatique du Langage Naturel), Workshop on Segmented Discourse Representation Theory, 22 April, Fèz, Morocco.*
- Asher, N. and Lascarides, A. (2003). *Logics in Conversation*. Cambridge University Press, Cambridge, UK
- Barthes, R. (1964). *Elements de Sémiologie*. Paris,: Éditions du Seuil.
- Benn, N., Buckingham Shum, S. and Domingue, J. (2005). Integrating Scholarly Argumentation, Texts and Community: Towards an Ontology and Services. 5th International Workshop on Computational Models of Natural Argument, IJCAI'05: International Joint Conference on Artificial Intelligence 2005, 30 July, Edinburgh, UK
- Buckingham Shum, S. (2003). The Roots of Computer Supported Argument Visualisation. In Kirschner, P.A., Buckingham Shum, S., Carr, Ch.S. (Eds.) *Visualising Argumentation. Software Tools for Collaborative and Educational Sense-Making*. Springer-Verlag, London, pp.3-24
- Buckingham Shum, S. Hammond, N. (1994). Argumentation-Based Design Rationale: What Use at What Cost? *International Journal of Human Computer Studies*, 40(4), pp. 603-652
- Buckingham Shum, S., Motta, E. and Domingue, J. (1999) Representing Scholarly Claims in Internet Digital Libraries: A Knowledge Modelling Approach. *Proceedings ECDL'99: Third European Conference on Digital Libraries, Paris, Sept. 22-24, 1999.* (Eds.) S. Abiteboul and A.-M. Vercoustre. *Lecture Notes in Computer Science Vol. 1696*, Springer-Verlag: Berlin
- Buckingham Shum, S., Motta, E. and Domingue, J. (2002). Augmenting Design Deliberation with Compendium: The Case of Collaborative Ontology Design. *Workshop on Facilitating Hypertext-Augmented Collaborative Modelling. ACM Hypertext Conference, Maryland, June 11th-12th, 2002* (<http://kmi.open.ac.uk/projects/compendium/SBS-HT02-Compendium.html>)
- Buckingham Shum, S., Uren, V., Li, G., Sereno, B. and Mancini, C. (In Press). Modelling Naturalistic Argumentation in Research Literatures: Representation and Interaction Design Issues. *International Journal of*

Intelligent Systems, (Special Issue on Computational Modelling of Natural Argument), (Eds.) C. Reed and F. Grasso.

Chklovski, T., Ratnakar, V., and Gil, Y. (2005). User Interfaces with Semi-Formal Representations: A Study of Designing Argumentation Structures. Proceedings of International conference on Intelligent User Interfaces (IUI 2005), January 9-12, San Diego, USA, ACM Press: NY.

Clancey, W.J., Sierhuis, M., Alena, R., Berrios, D., Dowding, J., Graham, J.S., Tyree, K.S., Hirsh, R.L., Garry, W.B., Semple, A., Buckingham Shum, S.J., Shadbolt, N. and Rupert, S. (2005). Automating CapCom Using Mobile Agents and Robotic Assistants. *American Institute of Aeronautics and Astronautics 1st Space Exploration Conference*, 31 Jan-1 Feb, 2005, Orlando, FL. Available from: AIAA Meeting Papers on Disc [CD-ROM]: Reston, VA, and as Advanced Knowledge Technologies ePrint 375: <http://eprints.aktors.org/375>

Eco, U. (1968). *La Struttura Assente*. Bompiani, Milano

Gangemi, A., Borgo, S., Catenacci, C., Lehmann, J. (2004). Task Taxonomies for Knowledge Content. METOKIS Deliverable D07

Handschuh, S. and Staab, S. (2002). Authoring and Annotation of Web Pages in CREAM. *Proceedings of the 11th International World Wide Web Conference*

Hjelmslev, L. (1959). *Essais Linguistiques*. Travaux du Cercle Linguistique de Copenhague, 12

Hobbs, J.R. (1985). On the Coherence and Structure of Discourse. Technical Report CSLI-85-37, Stanford, CA

Knott, A. (1998). Similarity and Contrast Relations and Inductive Rules. Proceedings of the ACL Workshop on Discourse Relations and Discourse Markers. Association for Computational Linguistics, Montreal, August 1988, pp.54-57

Knott, A., Dale, R. (1994). Using Linguistic Phenomena to Motivate a Set of Coherence Relations. *Discourse Processes*, 18(1), pp.35-62

Knott, A., Mellish, C. (1996). A Feature-Based Account of the Relations Signalled by Sentence and Clause Connectives. *Language and Speech*, 39(2/3), pp.142-183

Knott, A., Sanders, T.J.M (1998). The Classification of Coherence Relations: An Exploration of Two Languages. *Journal of Pragmatics*, 30 (2), pp.135-175

Liu, K. (2000). *Semiotics in Information Systems Engineering*. Cambridge: Cambridge University Press.

- Liu, K. (2005) Requirements Reengineering from Legacy Information Systems Using Semiotic Techniques, *Systems, Signs and Actions. An International Journal on Communication, Information Technology and Work*, 1(1): 36-61
- Lowrance, J.D., Harrison, I.W., and Rodriguez, A.C. (2001) Capturing Analytic Thought. *Proceedings First International Conference on Knowledge Capture*, pp. 84-91, October 2001, ACM Press: NY.
- Louwerse, M. (2001). An Analytic and Cognitive Parametrization of Coherence Relations. *Cognitive Linguistics*, 12 (3), pp. 291-315
- Mancini, C. (2005). *Cinematic Hypertext. Investigating a New Paradigm. Frontiers in Artificial Intelligence and Applications*. IOS Press, Amsterdam
- Mancini, C., Buckingham Shum, S. (2001). Cognitive Coherence Relations and Hypertext: from Cinematic Patterns to Scholarly Discourse. *Proceeding of ACM Hypertext '01*, Aarhus, Denmark, Aug 14-18, ACM Press, New York, pp.165-174
- Mancini, C., Buckingham Shum, S. (2004). Towards Cinematic Hypertext. *Proceeding of ACM Hypertext '04*, Santa Cruz, CA, USA, Aug 9-13, ACM Press, New York.
- Mann, W.C., Thompson, S.A. (1988). Rhetorical Structure Theory: Toward a Functional Theory of Text Organisation. *Text*, 8 (3), pp.243-281
- Metz, C. (1977) *Le signifiant imaginaire*. *Psychanalyse et cinéma*, Paris, Union Générale d'Éditions
- Pander Maat, H. (1999). The Differential Linguistic Realisation of Comparative and Additive Coherence Relations. *Cognitive Linguistics*, 10(2), pp.147-184
- Pander Maat, H., Degand, L. (2001). Scaling Causal Relations and Conectives in terms of Speaker Involvement. *Cognitive Linguistics*, 12(3), pp.211-245
- Peirce, C. (1931-1935). *Collected Papers*. The Belknap Press of Harvard University Press, Cambridge, MA
- Reed, C.A. & Norman, T.J. (2003) *Argumentation Machines: New Frontiers in Argument and Computation*. Kluwer: Dordrecht
- Sanders, T.J.M, Noordman, L.G.M. (2000). The Role of Coherence Relations and Their Linguistic Markers in Text Processing. *Discourse Processes*, 29 (1), pp.37-60
- Sanders, T.J.M, Spooren, W. (2001). Text Representation as an Interface Between Language and its Users. In Sanders, T.J.M., Schilperoord, J., Spooren, W. (Eds.), *Text Representation. Linguistic and psycholinguistic*

- aspects. University of Utrecht, University of Tilburg, Free University of Amsterdam, pp.1-26
- Sanders, T.J.M. (1997). Semantic and Pragmatic Sources of Coherence: On the Categorisation of Coherence Relations in Context. *Discourse Processes*, 24, pp.119-147
- Sanders, T.J.M., Spooren, W.P.M., Noordman, L.G.M. (1992). Towards a Taxonomy of Coherence Relations. *Discourse Processes*, 15, pp.1-35
- Sanders, T.J.M., Spooren, W.P.M., Noordman, L.G.M. (1993). Coherence Relations in a Cognitive Theory of Discourse Representation. *Cognitive Linguistics*, 4 (2), pp.93-133
- Saussure, De, F. (1922). *Cours de Linguistique Générale*. Editions Payot, Paris
- Selvin, A.M. and Buckingham Shum, S. (2002) Rapid knowledge construction: a case study in corporate contingency planning using collaborative hypermedia. *Knowledge and Process Management*, 9, (2), 119-128.
- Sereno, B. (forthcoming). A Document-Centric Semantic Annotation Environment to Support Sense-Making. Doctoral Thesis, Knowledge Media Institute, The Open University, UK.
- Sereno, B., Buckingham Shum, S. and Motta, E. (2005). ClaimSpotter: an Environment to Support Sensemaking with Knowledge Triples. Proceedings of International conference on Intelligent User Interfaces (IUI 2005), January 9-12, San Diego, USA, ACM Press: NY.
- Sowa, J. F. (2000). Ontology, Metadata and Semiotics. Proceedings of International Conference on Conceptual Structures, ICCS'2000, Damstadt
- Uren, V., Buckingham Shum, S., Li, G. and Bachler, M. (2006) Sensemaking Tools for Understanding Research Literatures: Design, Implementation and User Evaluation. *International Journal of Human Computer Studies*, 64 (5), 420-445.
- Uren, V., Buckingham Shum, S., Li, G., Domingue, J., Motta, E. (2003) Scholarly Publishing and Argument in Hyperspace. WWW 2003: 12th Int. World Wide Web Conference (ACM Press, Budapest, Hungary), pp. 244-250.
- Vargas-Vera, M., Motta, E., Domingue, J., Lanzoni, M., Stutt A., and Ciravegna, F. (2002). MnM: Ontology Driven Semi-Automatic and Automatic Support for Semantic Markup. *13th International Conference on Knowledge Engineering and Management (EKAW 2002)*, Eds. Gomez-Perez, A., Springer Verlag
- Weick, K.E. (1995) *Sensemaking in Organizations*. Sage: Thousand Oaks, CA