



Knowledge Media Institute

**Cognitive Coherence Relations and Hypertext:
From Cinematic Patterns to Scholarly Discourse**

Clara Mancini and Simon Buckingham Shum

KMI-TR-110

May, 2001

www.kmi.open.ac.uk/tr/papers/kmi-tr-110.pdf

Proceedings of ACM Hypertext 2001, Århus, Denmark, August 2001

www.HT01.org



Cognitive Coherence Relations and Hypertext: From Cinematic Patterns to Scholarly Discourse

Clara Mancini and Simon Buckingham Shum

Knowledge Media Institute
The Open University
Milton Keynes
MK7 6AA
UK

E-mail: {C.Mancini, S.Buckingham.Shum}@open.ac.uk

ABSTRACT

In previous work we argued that cinematic language may provide insights into the construction of narrative coherence in hypertext, and we identified in the shot juxtaposition of rhetorical patterns the source of coherence for cinematic discourse. Here we deepen our analysis, to show how the mechanisms that underpin cinematic rhetorical patterns are the same as those providing coherence in written text. We draw on computational and psycholinguistic analyses of texts which have derived a set of relationships that are termed Cognitive Coherence Relations (CCR). We validate this by re-expressing established cinematic patterns, and relations relevant to scholarly hypertext, in terms of CCR, and with this conceptual bridge in place, present examples to show how cinematic techniques could assist the presentation of scholarly discourse. This theoretical work also informs system design. We describe how an abstract relational layer based on CCR is being implemented as a semantic hypertext system to mediate scholarly discourse.

KEYWORDS: Cognitive Coherence Relations, cinematic rhetoric, scholarly hypertext, semiotics, argumentation, logical and analogical relations

INTRODUCTION

In this paper, we continue the investigation begun in last year's short paper [18], which proposed that cinematic principles behind the construction of coherent space-time 'worlds' have contributions to make to the design of coherent hypertextual narrative. Last year we were able to outline the conceptual links between cinema and hypertext. In this paper, we deepen our analysis with the introduction of a more formal theory of Cognitive Coherence Relations (CCR), derived from linguistic analysis research. CCR seeks, in essence, to

identify what makes texts "coherent", defining an empirically grounded set of parameters (a 'feature space'), which when combined in different ways can be used to define relational types. We extend CCR to hypertext, and, critically, to cinema. The theoretical goal is to have a common relational language to forge a link between the two worlds. This contributes to a strand of hypertext research that has emerged relatively recently, on the connections between hypertext, cinema and transitions [17,22,23].

Apart from building hypertext theory, we also have a practical design goal of designing scholarly hypertext infrastructures to complement the traditional research paper. It has been argued that hypertext offers both interesting new ways to represent certain kinds of argument, as well as posing particular challenges for a genre which has at its heart the construction of coherent narrative [4,14,15]. This applied thread appears in the paper to illustrate CCR's relevance to pressing representation problems germane to scholarly hypertext, but also, we would argue, to the construction of coherent hypertext narrative more generally.

The paper is organised as follows. We begin by introducing the theoretical and design problems we are facing. We then summarise the basics of CCR, in order to show firstly how cinematic coherence relations, and secondly scholarly discourse relations, can be re-expressed in terms of CCR's feature space. We pursue the scholarly hypertext design problem in more detail by describing how a CCR-based 'abstract relational layer' could provide a solution to some representational problems. From a semiotic perspective, we introduce analogical relations as a fundamental class of connection in both cinema and scholarly argument, noting that CCR provides a mechanism to handle similarity/contrast relations. Finally, we present several examples to show how CCR can be used to assist mapping between the worlds of cinema and scholarly discourse, and how scholarly hypertext could be treated in a cinematic manner.

THEORETICAL PROBLEM: MAPPING NARRATIVE PATTERNS BETWEEN MEDIA

Cinematic signification is based on the juxtaposition of shots, by which the film's discourse is generated. The cognitive connection of shots is conventionally based on a set of rhetorical patterns (the result of half a century of cinematic linguistic evolution) which provide coherence to the linear chain of shots, assisting viewers in recognizing the articulation of a discourse. In fact, the essence of the viewer's "reading" ability is in establishing connections between shots, as the meaning of any single shot dramatically changes depending on how it is connected to the others (as we will see later, this *role multiplicity* of semantic units is essential not only to cinematic discourse, but to any discourse type). Due to the iconic and indexical nature of the medium and to the semantic complexity of its minimal linguistic units, cinematic language does not constitute a grammar, it constitutes a rhetoric.¹

The parallel between the cinematic and hypertext medium is precisely based on the fact that the latter also consists of the articulation of rich semantic units whose connection, due to the activation of a link, effects and expresses a strong semantic relationship. Although hypertext units constitute self standing cores of content, the meaning of a

single unit changes depending on how its connections with the others are activated. That is, the role of any hypertext unit depends on the navigational path it happens to be part of according to the user's choices.

This potential multiplicity of roles is responsible for the multidimensionality of hypertext spaces, that is, the contextual variability of units' meaning. However, while on the one hand this potential constitutes the essential richness of hypertext, on the other hand it also constitutes its main weakness, as it exposes hypertext discourse to a lack of cohesion and coherence. This may particularly compromise the use of hypertext for certain functions like, for instance, argumentation or, more generally, scholarly discourse.²

Hypertext rhetorical patterns (e.g. [1]) may have for the user the same function that cinematic rhetorical patterns have for the viewer, helping them (once they have become used to it) to identify coherent connections amongst all possible ones. In fact, if it were possible to take advantage of the insights gleaned through the cinematic medium's technical characteristics and historical evolution, this would provide hypertext with a set of rhetorical patterns, and above all strategies, that could assist the user in navigation.

¹ The minimal linguistic unit of natural language is the phoneme, a *symbolic* non-signifying differential element, whose combination generates morphemes successively articulated to generate the enunciation. The cinematic minimal linguistic unit is the *shot*, an *iconic* and *indexical* [24] semantically rich element, which, in semiotic terms, is the equivalent of a linguistic enunciation. Because in natural language the nature of the signifier has nothing to do with the nature of the signified, the correspondence between the signifier and the signified is more convention-based. Therefore, concepts can be expressed more explicitly, the user being required to build a mental representation of concrete elements. In contrast, because in cinematic language the nature of the signifier does have something to do with the nature of the signified, the correspondence between signifier and signified is less convention-based, so that cinematic language can suggest concepts more implicitly, through the representation of concrete elements and events.

Because in natural language, connections can be explicitly enunciated, their correctness is constantly verified against the content of the units that are connected and a connection can in fact be right or wrong. In contrast, in cinema, because connections are performed without being enunciated, any connection tends to be seen as coherent and sensible in one way or another, as viewers seek constantly to make sense of transitions. As Miles underlines [22], in cinema a connection is not in principle right or wrong, but good or bad, effectual or ineffectual, and a cinematic sequence has to be deciphered as a structural whole, the single element (unit or transition) having no specific meaning in itself. This is why cinematic language, unlike natural language, cannot be considered as a grammar, but as a rhetoric [20].

² To clarify our perspective we suggest the distinction between hypertext and hypertextuality. By *hypertext* we refer to the medium, the concrete form which can only work in an interactive medium. It can follow different registers or working modalities, that we distinguish according to two different parameters: the visibility on the contents' organisation, on the one hand, and the definition of connections between hypertext units, on the other hand. On the one hand, the reader can have access to the contents through a global, *aerial level* view (node maps) or through local, *ground level* discovery ('link by link paths'). On the other hand, the connections between content units can be *enunciated* (explicitly indicated and/or classified as in most semantic hypertexts), or their identification and classification can be *evoked* by visual/kinetic features (like in spatial hypertext). For instance, we see *ScholOnto* [3] as aerial with enunciated connections; *VKB* [28] as aerial with evoked connections; the *Assembly 3/1* hypertext [16] as both ground level and aerial, with indicated but unclassified connections; and *HyperCafé* [27] as ground level with enunciated but unclassified connections.

By *hypertextuality* we mean a quality that any form of text (literary, film, etc.) may possess to different degrees. It refers to the possibility of reading paths or patterns that cross the linear or sequential distribution of text contents. Although a book or a movie is potentially hypertextual, due to the technical characteristics of the medium, hypertextuality finds itself reified only in interactive hypertext. A hypertext's hypertextuality is proportional to its interactivity, that is, to the freedom that the readers have to reify the hypertextuality of the text and to 'perform' hypertextual thinking.

The theoretical challenge is to formulate the basis for a rigorous and systematic “transformation” of rhetorical patterns from cinema to hypertext. The two media have some important differences, mainly relative to the linguistic codes that they use. Although cinematic language includes different codes like oral speech, written text, music and sound, it is essentially iconic, being based on moving images. In other words, visual processes are at the basis of cinematic text reading.

Hypertext is also a visual medium, in that navigation develops on a computer screen and its discourse can make use of images, as well as of sounds. Nevertheless, written text typically prevails, especially in scholarly discourse. The written word still constitutes the backbone of scholarly literature discourse and production as it lends itself to the explicit expression of abstract concepts. This means that symbolic processes are prevalently activated in (scholarly) hypertext reading, similar to literary text reading.

We will try to demonstrate that this difference in the working mechanisms of the two media does not affect the types of cognitive connections that make their respective discourse patterns coherent.

DESIGN PROBLEM: EXPRESSIVE FLEXIBILITY IN SCHOLARLY HYPERTEXT

This theoretical work also informs our approach to a particular design problem, namely, the design of *ScholOnto*, a digital library system to support the dissemination and analysis of scholarly research literatures [3]. We are adopting a discourse-oriented approach to enable researchers to make claims about the contributions of their own, and others’, publications, and to engage in debate about these. A conceptual network will grow as researchers add new claims. The server will manage the network’s complexity, and make it worthwhile to contribute to it by providing services such as agents, filters, visualizations and structural pattern matching.

Of relevance to this paper are some particular design challenges we face:

1. It is in the nature of research that scholars may not agree on the status or classification of a concept. In hypertextual terms, this normally corresponds to the assignment of types to nodes. As a first principle, we need to manage multiple interpretations of concepts, and allow concepts to play multiple roles depending on research perspective and purpose.
2. We want to enable researchers to make semantic connections between concepts in the literature. We can provide them with a predefined schema which expresses our view of the most important relationships, but this will inevitably prove inadequate—we simply cannot anticipate expressions of all possible relationships that all researchers might

wish to express. We want to do more than simply allow users to invent their own relational types, which could result in an explosion of new types and a weakening of any support the system can provide to manage the concept network.

3. One of our goals is to help researchers find relevant work across the literature of multiple disciplines. However, different disciplines speak different languages: terminology, vocabulary, modes of argument all vary. How can we make representations of literatures ‘interoperable’ in an appropriate way? We shall return to these representational problems later, to show how our approach to the theoretical problem set out above is also assisting us in design.

COGNITIVE COHERENCE RELATIONS

Let us start with the notion of *role multiplicity*. Just as in cinematic discourse construction, a shot can assume different meanings depending on the system of connections it is part of, likewise in scholarly discourse the same object may play different roles (data, evidence, problem, method) depending on the context of connections it is part of, defined by a given author. However, while objects can play different and indefinite roles, the way we infer meaningful *connections* between discourse units seems to be restricted by a finite number of candidates. The core idea is that not only is there a relatively small set of transitions between ‘events’ or ‘semantic units’ in a medium that make sense in a given context, but moreover, these may themselves be configurations of a base set of relationships.

It is here that work on Cognitive Coherence Relations is relevant. Computational and psycholinguistic research on text coherence relations (motivated by goals such as natural language parsing and generation), reports evidence that the articulation of complex text structures is based on a small set of elementary cognitive relations, from whose combination other coherence relations can be derived. These coherence relations are hypothesised to be “universal” in the sense that they are *cognitive*, not specific to a particular domain or genre of text. This line of work critiques and extends Mann and Thompson’s Rhetorical Structure Theory [19].³

Sanders *et al.* proposed a cognitive theory of discourse representation to explain text understanding and construction [26]. They hypothesise that coherence relations are inferred by the reader on the basis of a few primitive cognitive concepts. In this respect, coherence is not an intrinsic property of discourse, but of the reader’s *interpretation*. The authors identify four of these primitive concepts, on the basis of which they propose a taxonomy

³ RST has been considered by Rutledge *et al.* [25] as the basis for constraining the automatic generation of hypermedia sequences.

of coherence relationships. *Basic operation* (including additive and causal relations), *source of coherence* (including semantic and pragmatic relations), *order of segments* (including basic and non-basic order) and *polarity* (including positive and negative relations) combine to generate classes of coherence relations.

With the same goal, Knott, *et al.* [11,12] use the presence of linguistic cue phrases in naturally occurring text as evidence to motivate a set of rhetorical relations, also corresponding to cognitive categories. Like Sanders *et al.*, they conceive coherence relations as *psychological* constructs. Their assumption is that natural language is optimised to express these constructs, so that language's cohesive devices can provide evidence for psychological coherence relations (which are the real object of study). These devices are meant to express the effect that the text is intended to have on the reader, which is that the reader assumes a specific relation to hold between the content of the connected text spans. Through an empirical substitution test (in which connective phrases in a text are substituted to see if the original sense is preserved), they converged on a set of eight binary-valued parameters (later reduced to seven: *anchor*, *pattern of instantiation*, *focus of polarity*, *polarity presuppositionality*, *modal status* and *rule type*), which define a reasonably restricted taxonomy of coherence relation types, organised as a tree.

More recently, the approach has been developed jointly by Knott and Sanders [13], in an attempt to confirm the correspondence between the set of cognitive relations (identified by Sanders *et al.*) and the set of linguistic devices (analysed by Knott and Dale). As a larger source of evidence, they conducted experiments on both Dutch and English subjects, and concluded that the correspondence does exist. In brief, the relational categories identified belong to types like *causal* (A causes B), *disjunctive* (A is alternative to B), *conjunctive* (A coexists with B), *conditional* (A presupposes B) and *sequential* (A follows B). From the point of view of the propositional content of related spans of text, these relation types imply 'defeasible' logical rules, which may either succeed or fail. If the rule succeeds, the relation is considered positive (e.g. A *does* follow B), if the rule fails the relation is considered negative (A *does not* follow B).

We now show how relationships in cinematic language, and then in scholarly hypertext, can be reinterpreted in terms of CCR.

CINEMA AND CCR

The rhetorical patterns that cinema developed over half a century constitute narrative models for conveying logical cause-effect sequences of events in a coherent space-time world. Cinematic avant-garde movements in different periods sought repeatedly to show this is not the only way cinematic language might have developed, but such conceptions of the medium did not prevail. The fact that the medium's evolution led to the selection of the

narrative structures currently used suggests that these patterns reflect the "optimal" trade-off between the medium's expressive potential and the need for cognitive coherence. It is not by chance that narratology has provided the main theoretical framework and tools to analyse the cinematic medium. Literary and cinematic rhetorical structures are comparable to the extent that at different times they have even radically influenced each other (see, for instance, the dialectical interaction between Nouveau Romans and Nouvelle Vague [7]).

This sets the historical background for a step in our argument. When we try to map the set of coherence relations proposed by Knott *et al.* onto Christian Metz's *grande syntagmatique* of 'classic cinema' rhetorical patterns [20], the correspondence is self-evident. For instance, in Metz's *linear narrative syntagmas* of the 'scene' and of the 'proper sequence', we can assume that *positive causal* and *positive sequential relations* connect the diverse shots; in Knott's terms [10], the "intended effect" of the rhetorical pattern is that the viewer assumes those relations to "hold" between the narrative content of the sequence's shots. In the *descriptive syntagma*, *positive conjunctive relations* are intended to be assumed between the content of the shots. Finally, in the *narrative alternated syntagma*, the relations keeping the shots' content together can be read as being both/either *positive conjunctive* and/or *positive sequential*. We contend that this mapping is both plausible, and can be completed without difficulty. We recognize that Metz's rhetorical patterns used in "classic" cinema provide only one possible key to relation interpretation, as the result of a 'standardisation' process towards narrative univocal linearity. Moreover, they can themselves be interpreted, in terms of CCR, in different ways (see footnote 6). The important point, however, is that these dominant cinematic rhetorical patterns can be read in terms of CCR.

SCHOLARLY HYPERTEXT AND CCR

We have used the set of coherence relations to evaluate the relational set proposed for ScholOnto [3]. ScholOnto's set comprises semantic relations typically used to make scholarly claims, like *agrees*, *disagrees*, *confirms*, *supports*, *addresses*, *envisages*, *predicts*, *is-inconsistent-with*, *is-consistent-with*, and so on. For instance, the relation *disagrees* corresponds to a *negative, actual, non-presuppositional* coherence relation (a particular type of disjunctive relation); *confirms* and *supports* correspond to a *positive, result-driven, actual, non-presuppositional* coherence relation (a particular type of causal relation).

As an empirically grounded taxonomy, CCR provides us with a way to critique the ScholOnto relational set, which until now has been derived from our intuitions as researchers, and our (necessarily incomplete) modelling of various literatures. Firstly, if any scholarly relations map to the same position in the CCR tree, it indicates that their meaning partly overlaps. It was no surprise to find that some relations do fall into the same CCR category,

since we had intentionally designed ‘link families’, for instance, to provide degrees of strength in expressing a difference of opinion with a concept or perspective, from *raises issues with* to *is inconsistent with*, to *challenges* to *refutes*. We describe elsewhere how this enables forms of semantic search and link computation [2]. CCR provides us with a way to define link families in a more principled manner.

Secondly, gaps in the CCR tree—relational types for which there are no ScholOnto expressions—draw our attention to potentially important omissions. For instance, the absence of a *positive presuppositional* relation could motivate the creation of a relational label such as *implies/follows* if this was deemed important enough. The significance of such gaps in the CCR tree depends on the rhetorical needs of the scholars in a given field. It may be that different research fields will require coherence relations that are not represented in the scholarly relation set at present. In the following section we propose the idea of a CCR-based abstract relational layer that helps tackle issues of linguistic variation in the way that individuals wish to express themselves.

Addressing hypertext representation problems

Earlier we introduced three design challenges in designing ScholOnto. The above work on CCR is helping to establish foundations for its discourse-relations ontology. The extension that CCR provides to previously published descriptions of ScholOnto is an ‘abstract relations layer’, of which any particular set of relational labels is but one ‘dialect’. The grounding of relational labels in underlying (CCR) semantics opens up creative possibilities for tackling the problems identified earlier:

1. *Requirement: Concepts may need to play multiple roles.* The ontology is relation-centric. Scholarly concepts need not be classified. All of the semantics are carried in the relationships, which can optionally specify a concept type at one or both ‘ends’ of the link, to clarify the role that the concept plays *in the context of this link*. Other links to that concept may cast it in different roles. In sum, meaning is not intrinsic to concepts, it derives from their use in a context.
2. *Requirement: Expressive nuances for relationships.* We have shown that some common scholarly relationships can be mapped to CCR. However, these particular expressions may not appeal to a given user. Academics are careful about how they position their work, using a specific verbal expression to relate their work to that of others. CCR provides us with a mechanism to handle this. ScholOnto knows only about abstract relational classes grounded in CCR, but each may have many dialects. If researchers using the system need to invent their own ‘link label’ beyond those offered, they can, but it will inherit membership of the ‘link family’ [2] (relational class) they are extending. In this way, the system has a

coarse understanding of the link’s semantics (for computational purposes), but leaving scholars to appreciate the more subtle nuances.

3. *Requirement: Representational support to bridge across disciplinary boundaries around literatures.* Our strategy for this problem is an extension of (2) above. Different disciplines and traditions may use different language, and make use of different styles of argument (corresponding in ScholOnto to structural patterns), but they will still be using CCR. Thus, CCR could assist in handling inevitable variations in language within a field (‘intraoperability’), it could also provide a common layer to make fields interoperable.

We are currently implementing ScholOnto along these lines, but the problem of expressive inflexibility is not restricted to scholarly hypertexts. Any collaborative system that requires users to classify contributions must confront the fact that users’ interpretations of what they are doing may not match the predefined categories [29]. Implementing an abstract relational layer is one way forward.

FROM LOGICAL TO ANALOGICAL RELATIONS

Through the exploration of cinematic and scholarly domains, we have tried to point out the importance of *connections* in discourse coherence and we have identified in CCR a promising candidate for the development of a discourse relation system. Thus far we have focused on *logical* connections between discourse units, which are typically the relational types used in semantic hypertext systems. However, *analogical relations* have yet to be discussed. In this section, we express in semiotic terms the problem of ‘coherent connection’, in order to extend our analysis to analogical relations.

Analogical relations in semiotics

According to De Saussure [5], relations between linguistic terms can develop at two levels, corresponding to two different forms of mental activity. The combination of linguistic signs (concepts) in a linear chain constitutes the *syntagma* (connection *in presentia*). The association of signs (concepts) having something in common constitutes the *paradigm* (association *in absentia*).

In semiotic terms, *logical* relations consist of *syntagmatic* connections (like in linguistic phrases or cinematic sequences), in the development of a logical system. On the other hand, *paradigmatic* associations constitute *analogical* relations. Analogical relations forge a ‘connection’ between two logical systems.

Literary rhetorical figures constitute a type of analogical connection through the substitution of one discourse element with another belonging to a different semantic domain. This substitution is based on a functional equivalence, and is centred on one or more attributes of the elements involved [8]. Both in natural and in

cinematic language, these mechanisms are at the basis of figurative expression (as in dreams or in poetry).

In contemporary rhetoric theory and literature, the most used and analysed tropes are metaphor and metonym.⁴ In *metaphor* the substitution is based on the similarity of two functionally equivalent objects. In *metonym* the substitution is based on the functional equivalence of one part of an object for the whole object.

Analogical relations in CCR

Turning now to CCR theory, Knott [9] analyses the *comparison* relation as implying a *similarity* defeasible rule, which in the event of failure (an expected similarity between two elements turns out *not* to hold) becomes a *contrast* relation. Similarity and contrast relations are seen as respectively permitting or preventing an inductive rule to fire in the reader's mind, that is, if two objects are similar in one respect, this gives rise to an expectation of similarity in other respects, making it possible for inductive reasoning to take place.

If we apply this mechanism to the way rhetorical figures work in literature, we find that for metaphors the correspondence is quite close: on the basis of a similarity in one respect, two objects may be assumed to be equivalent also in other respects, to the point that the former can be substituted to the second in the discourse. Although in metonym the substitution is made on the basis of a contiguity, the substituting element (part) is chosen for its equivalence to the all (that is, for its functional similarity).

Analogical relations in cinema

As already described, the cinematic medium developed its discourse apparatus to construct logical cause-effect sequences of actions in coherent time-space worlds. However, because of the iconic nature of the language (based on the juxtaposition of visual units), narrative models can easily integrate the analogical with the logical paradigm. Compared to logical juxtapositions, the power of analogical juxtapositions to generate meaning is incomparably higher. Eisenstein used to call this *intellectual montage* (as distinguished from *harmonic montage*), making use of visual metaphors or metonyms in his films. If it is meant to be metaphorical, this type of analogical relation is likely to characterise the connection between an *autonomous insert* and the shots amongst which the insert is located. If it is meant to represent a

⁴ These are also significant elements in psychoanalysis and in certain psychoanalytic theories of cinema. In psychoanalysis these two figures are meant to be involved in the mechanisms of condensing and shifting, characterising what are called the *primary processes*. In cinematic theory, condensing and shifting, either on metaphoric or metonymic basis, are assimilated to figures like, for instance, the *fade* or the *autonomous insert* [21].

metonym, the relation might hold between a *detail* and a *total* of an ordinary *chronological* or *a-chronological sequence* (a common device in murder mysteries).⁵

Analogical relations in scholarly hypertext

Analogical devices have a relevant role not only in poetry, dreams or film, but commonly in scholarly communication. A hypertext system to mediate scholarly discourse should provide tools enabling scholars to establish metaphorical and metonymical connections.⁶

CINEMA — CCR — SCHOLARLY HYPERTEXT

We have argued that CCR could offer a kind of 'relational interchange format' to understand how narrative sequences in two different media such as cinema and hypertext can 'talk' to each other. Preceding sections have shown the CCR mapping to each of these domains. In this section, we take the next step, namely, to consider how to use cinematic techniques to *present* hypertextual arguments with greater force, based on the visual aspects (like shape, color, font, spatial position and juxtaposition of units) and kinetic aspects (like duration and synchronisation of transitions) that characterise cinematic language.

Specifically, CCR could provide the implicit rules to define the type of connection that the hypertext author means to establish between two hypertext units. These connections would be based on, or derived from, CCR to express more specific languages and dialects, relevant to specific scholarly communities. The author could decide whether the hypertext connections should be explicitly indicated and classified, or be left implicit and only used by the system to render the hypertext's structure, which the end user would have to infer on the basis of visual and

⁵ Widely explored by diverse cinematic schools like the Russian, the Dadaist and the Surrealist, the awareness of the analogical dimension constitutes the richest legacy of the avant-garde experience to 'classical' cinema, where the strength of analogical connections is amplified by the fact that they break logical structures. The analogical is the dimension of transversality: against the 'linearity' of the narrative regime, it acquires even more expressive power.

Interestingly, the peculiarity of analogy is also acknowledged by Walton [30] in his treatise on argumentation schemes for presumptive reasoning, where he states that reasoning by analogy is the one scheme which is not possible to define in an univocal manner (pp. 77-80).

⁶ It is precisely within the analogical dimension, and its multiple resonances, that we see a resemblance with Kolb's concept of scholarly hypertext *regions*: "musical compositions [...] create rich and complex temporal and formal and thematic connections to other sections of the piece in many different levels and sizes. Hypertext could do the same with sequences and nodes...there should be large structures, echoes, returning themes, transformations and recapitulations and variations [...]" ([15], p.31-32), richly expressed in his most recent work [16].

kinetic features. Authors would define in advance the visual/kinetic style for connection types, thus establishing a local, consistent, visual language, just as film directors do. Visual and kinetic features could be predefined at different levels to consistently express relational types, at different levels. Table 1 illustrates the notion of distinct, but connected relational levels: a CCR ‘foundation’, with one or more community or individual specialisations as additional layers.⁷

| Underlying Connective Structure | Presentation |
|-------------------------------------------------|--------------------------------------------------------------------------------------------|
| Community/individual relational specialisations | Optional specialisation of visual/kinetic features |
| Cognitive Coherence Relations | Basic visual/kinetic features (shape, color, font, position/juxtaposition, behavior, etc.) |

Table 1. Relational layers in a scholarly hypertext authoring environment. Relational semantics are rooted in CCR, and rendered using presentational constraints inspired by cinematic language.

We present two examples to illustrate this convergence, each comprising four sections: *Cinema*, *CCR*, *Argumentation* and their *Presentation* at *aerial* and *ground* levels (see footnote 2).

Example 1

Cinema. One of the typical cinematic patterns of French school is the *analytic découpage*, a particular kind of descriptive syntagma. It is mainly used in interiors for scenes like social meals and conversations, where a circumscribed environment made of many different elements has to be constructed and presented. The analytic découpage consists for the most part of close-ups, which towards the end enlarge to ‘medium’ and ‘entire’ ‘plan’. The close-ups at the beginning focus on details of the characters, environment and actions, and progressively suggest and reveal the spatial structure and the constituting elements of the scene. The final shots serve to collect all the details together in order to give the global picture and, most important, work as evidence of the fact that all that has been shown till then is actually part of the scene. Without those final collecting shots, the viewer could never be sure that what they have seen is cohesively part of the same environment and contribute to build a coherent space.

CCR. From a coherence relation perspective, the connections between the close-ups can compare to a series of *actual non-presuppositional* (that is, *conjunctive*)

⁷ One can conceive extending *Visual Knowledge Builder* [28] or *Storyspace* [6] with CCR and cinematic presentational features.

relations. On the other hand, the intermediary and the final ‘summarising’ shots would hold a *cause-driven, positive causal actual* (that is, *causal*) relation with all the other shots.

Argumentation. A rhetorical strategy sometimes used is to intentionally hide from the reader the fact that elements being described are connected, introducing them in turn with, at most, hints of interconnections, adding greater force to their final convergence. A variation on this is to examine a few elements, then show their connections (i.e. place them in the same ‘shot’), before repeating the process, cumulatively building a complex whole.

Presentation. This cognitive coherence relation pattern could be represented in hypertext in different ways. At an *aerial* level (see footnote 2), it might be translated as a series of hypertext units appearing in sequence (as indicated by the letters in figure 1) or being distributed on the screen in order to render in spatial terms the rhetorical movement (as shown in Figure 1). The hypertext connections might be or might not be explicitly indicated or classified (see footnote 2). The discourse units A, B and C would start to build the pattern first, followed by a unit addressing and putting together the three of them, providing a first cohesive representation of a relevant part of the argument. The units D, E, F, G would follow in sequence, until a final discourse unit would put them and the A-B-C unit together, providing a complete view, completing the shape of the argument and guiding the user to the conclusions.




Figure 1. Aerial representation (node map) of an argument using conjunctive and causal relations

At a *ground* level (see footnote 2), the relation pattern might be expressed by visual and kinetic elements working as indicators for the relations themselves (see Figure 2). Also in this case, the hypertext connections might or might not be classified. As the user rolls the cursor over an active link, visual indicators might suggest the binary relation holding between the current node (for instance A) and the target node (for instance B or C, or D

or E); kinetic as well as visual elements might express the particular relation characterising the event of the transition; finally, other spatial and visual elements might characterise the resulting juxtaposition of the two connected units.

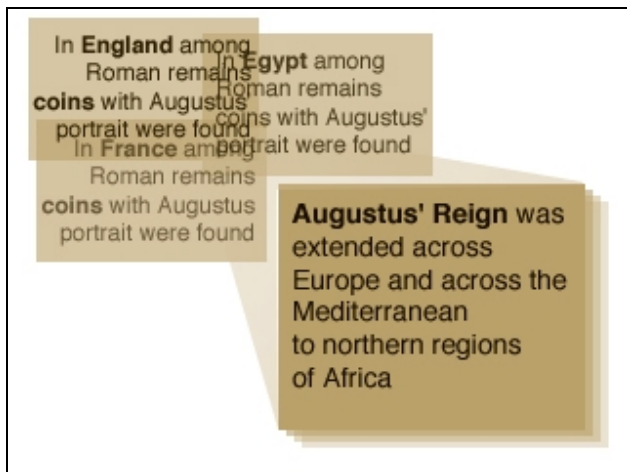


Figure 2. *Ground level* representation of an argument using conjunctive and causal relations. Colour, typeface, translucence, and spatial position have been defined to distinguish the presentation of *conjunctive* and *causal* relationships. (Kinetic/animation features could also be used, but are hard to convey in a static image)

Example 2

Cinema. Analogical coherence relations in cinema may be expressed by what Metz called the *descriptive syntagma* with an *autonomous insert* presenting a scene which has a metaphorical value with respect to the rest of the sequence. (This type of pattern is exploited in the Soviet school, characterised by a strong ideological valence; famous examples are Eisenstein’s *Strike* and *Between the Old and the New*). Metonym is more commonly represented by shots like a *detail* included in a proper *scene* (in Metz’ terms, a sequence representing an action with unity of time and space). The detail shows a meaningful element standing for an entire object or situation.

CCR. From a coherence relation perspective, the relation which holds between the metaphorical/metonymical shot and the rest of the sequence is a *comparison*, where the *similarity* rule succeeds (with respect to the attributes determining the choice of the substitute), allowing (in Knott’s terms) the induction to ‘fire’ that the shot (e.g. an *autonomous insert* or *detail*) provides clues to the meaning of the scene into which it is inserted.

Argumentation. It is not uncommon for speakers/writers to introduce a story or joke into the flow of an argument. At its best, the story/joke paints a metaphorical picture which throws into relief some aspect of the argument.

Such analogical modes of reasoning are used to evoke connotations in the reader’s mind congruent with a particular perspective. Interestingly, a picture or metaphor can be used not only to cast an interpretation on established data/ideas, but also to project into the future. If one can establish that a metaphor works for what is known, it gains credibility as a way to understand what is unknown (“since the earth is like an organism in these respects, then it may also be able to heal itself”). (Naturally, it may be contested that an analogy breaks or is false.) It is not hard to see the correspondence between the introduction of metaphor to argument, and the use of the autonomous shot in cinema to evoke interpretations of what has been seen, or what is yet to come.

Presentation. This pattern could be represented as a series of discourse units building the hypertext space (as shown and indicated by the letters in Figure 3).




Figure 3. *Aerial* representation (node map) of an argument using sequential and analogical relations.

Like in *Strike*, one discourse unit could be substituted with an autonomous unit (external to the logical order of the argument: story, joke, image or other) holding with the rest of the unit series an analogical relation. The conclusions, that should follow the presentation of the “discourse insert”, could even be left to the intuition of the user, being suggested by the insert itself. At an *aerial* level, the relation pattern could take the shape of a node map where the nodes appear in sequence, being appropriately distributed on the screen (Figure 3). Hypertext relations might or might not be indicated or classified (see footnote 2).

At a *ground* level, the same relation pattern might be expressed by visual and kinetic elements to help the users identify the analogical hypertext unit to be juxtaposed to the logical series of the other units. So, when they activate the corresponding link they should be able to recognise the break and the jump into another dimension, and to pay attention to the possible reflections of the analogical dimension on the logical one (see Figure 4). Also in this

case, hypertext relations might be or might not be classified.⁸

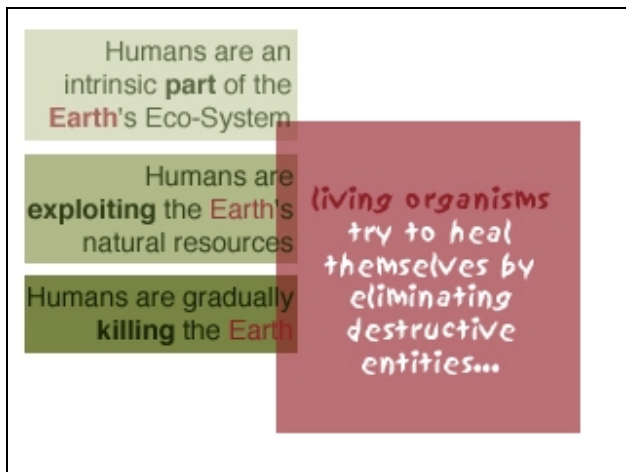


Figure 4. *Ground level* representation of an argument using sequential and analogical relations. Colour, typeface, translucence, and spatial position have been defined to distinguish the presentation of *logical* and *analogical* relationships.

CONCLUSIONS AND FUTURE WORK

We began by formulating the problem as one of knowing how to transfer cinema's strategies for creating coherence in a virtual space-time world to hypertext. If cognitive coherence relations theory holds, then this may not be the right way to think about the problem. If there is in fact a set of coherence relations that we are always using to communicate, and always seeking when interpreting signs, then all symbol systems—whether writing, film, painting, photography—will have evolved genres to express these relationships creatively within the constraints of the medium. In this sense, it is hardly surprising to find echoes of the same narrative devices across media if they derive from the same cognitive coherence relations. As we have already pointed out, there is a strong tradition of mapping between the rhetoric of literature and cinema, and hypertext theorists have explored for some time the relationship of linear and non-linear literature.

What we find intriguing, and offer for consideration by the hypertext community, is the extent to which the precise (formally expressible) formulations of text coherence relations as represented by the work of Knott, *et al.* described here, map to cinematic theory, in

⁸ As visible in Figure 4, the link between the first and second nodes, and especially the link between the second and third nodes, can be interpreted as *sequential* and *causal* at the same time. The difference depends on the writer's intended effect as well as on the readers' interpretation. The point is that CCR still appears to connect the nodes.

turn strengthening its relevance to hypertext. Secondly, we have demonstrated that a set of discourse relationships devised for a scholarly hypertext system can be resolved down to their positions in the CCR feature space. Thirdly, and adopting a more system design-orientation, we have described how CCR (or similar formulations of coherence relations) help to tackle a persistent problem with semantic hypertext, namely, that it constrains users to a predefined set of relational types. An abstract relational layer in a hypertext system could provide a principled mechanism to provide expressive flexibility, a hypothesis that we are now testing. We have described how Knott *et al.*'s formulation of CCR has built into it a potential mechanism to support not only logical, but analogical, relationships. This raises the intriguing prospect of an elegant mechanism to express important forms of analogical cinematic patterns, as well as providing a CCR-based scholarly hypertext system with analogical relationships for argumentation. One convergence of these ideas is in the idea of using the visual and kinetic presentational constraints characterising cinematic language to express rhetorical patterns in hypertexts.

Several strands for future work emerge. We are continuing to implement ScholOnto, a system that embodies these ideas in an *aerial level* type of hypertext. We are also beginning to conceive an authoring environment to design the *ground level* experience for a hypertext user, using CCR to support the creation and the management of connections, on the one hand, and visual/kinetic aspects characterising cinematic language, on the other. Future papers will evaluate the usability, expressive power and performance of these systems, and the theories inspiring them.

REFERENCES

1. Bernstein, M. Patterns of Hypertext. In *Proc. Hypertext'98*: Pittsburgh, PA, 1998, ACM Press: New York, pp. 21-29
2. Buckingham Shum, S., Domingue, J. and Motta, E. Scholarly Discourse as Computable Structure. In *Open Hypermedia Systems and Structural Computing*, Lecture Notes in Computer Science Vol. 1903, Springer-Verlag: Berlin, pp. 120-128 [www.kmi.open.ac.uk/projects/scholonto].
3. Buckingham Shum, S., Motta, E. and Domingue, J. ScholOnto: An Ontology-Based Digital Library Server for Research Documents and Discourse. *Int. Jnl. on Digital Libraries*, 3, 3, 2000, pp. 237-248 [www.kmi.open.ac.uk/projects/scholonto].
4. Carter, L.M. Arguments in Hypertext: A Rhetorical Approach. In *Proc. Hypertext 2000*, San Antonio, TX, 2000, ACM Press: New York, pp. 85-91
5. De Saussure, F. *Cours de linguistique générale*. Editions Payot: Paris, 1922
6. Eastgate Systems: *Storyspace* [www.eastgate.com]

7. Gaudreault, A. *Du Littéraire au Filmique. Système du Récit*. Méridiens Klincksieck: Paris, 1989
8. Genette, G. *Figure III*. Editions du Seuil: Paris, 1972
9. Knott, A. Similarity and Contrast Relations and Inductive Rules. In *Proc. Workshop on Discourse Relations and Discourse Markers*, International Conference on Computational Linguistics, Montreal, CA, 15th August, 1998
10. Knott, A. Semantic and Pragmatic Relations and their Intended Effects. In: *Text Representation: Linguistic and Psycholinguistic Aspects*, Sanders, T., Schilperoorol, J. and Spooren, W., (Ed.), Benjamins, 1999
11. Knott, A. and Dale, R. Using Linguistic Phenomena to Motivate a Set of Rhetorical Relations. University of Edinburgh, *Technical Report HCRC/RP-39*, 1993
12. Knott, A. and Mellish, C. A Feature-Based Account of Relations Signalled by Sentence and Clause Connectives. *Language and Speech*, 39, 2-3, 1996, pp. 143-183
13. Knott, A. and Sanders, T. The Classification of Coherence Relations and their Linguistic Markers: An Exploration of Two Languages. *Journal of Pragmatics*, 30, 1998, pp. 135-175
14. Kolb, D. *Socrates in the Labyrinth: Hypertext, Argument, Philosophy (A Hypertext)*. Eastgate Systems: Watertown, 1994 [www.eastgate.com]
15. Kolb, D. Scholarly Hypertext: Self-Represented Complexity. In *Proc. Hypertext'97*, Southampton, UK, 1997, ACM Press: New York, pp. 29-37
16. Kolb, D.: *Assembly 3/1* version of the *Sprawling Places* hypertext project, (in press)
17. Luesebrink, M.C. The Moment in Hypertext: a brief lexicon of time. In *Proc. Hypertext'98*, Pittsburgh, PA, 1998, ACM Press: New York, pp. 106-112
18. Mancini, C. From Cinematographic to Hypertext Narrative. In *Proc. Hypertext 2000*, San Antonio, 2000, ACM Press: New York, pp. 236-237
19. Mann, W.C. and Thompson, S.A. Rhetorical Structure Theory: Toward a functional theory of text organisation. *Text*, 8, 1988, pp. 243-281
20. Metz, C. *Essai sur la Signification au Cinéma*. Editions Klincksieck: Paris, 1968
21. Metz, C. *Le Signifiant Imaginaire. Psychanalyse et Cinéma*. Union Générale d'Editions: Paris, 1977
22. Miles, A. Cinematic Paradigms of Hipertext. *Continuum: Journal of Media & Cultural Studies*, 13.2, 1999, pp. 217-226
23. Miles, A. Hypertext Structure as the Event of Connection. In *Proc. Hypertext 2001*, Aarhus, Denmark, Aug 14-18, 2001, ACM Press: New York
24. Peirce, C. *Collected Papers*. The Belknap Press of Harvard University Press: Cambridge, MA, 1931-1935
25. Rutledge, L., van Ossenbruggen, J., Hardman, L., Geurts, J. and Bailey, B. Generating Presentation Constraints from Rhetorical Structure. In *Proc. Hypertext 2000*, San Antonio, TX, 2000, ACM Press: New York, pp. 19-28
26. Sanders, T.J.M., Spooren, W.P.M. and Noordman, L.G.M. Coherence Relations in a Cognitive Theory of Discourse Representation. *Cognitive Linguistics*, 4, 1993, pp. 93-133
27. Sawhney, N., Balcom, D. and Smith, I. HyperCafé: Narrative and Aesthetic Properties of Hypervideo. In *Proc. Hypertext'96*, Bethesda, MD, 1996, ACM Press: New York, pp. 1-10
28. Shipman, F.M., Hsieh, H., Maloor, P. and Moore, J.M. The Visual Knowledge Builder: A Second Generation Spatial Hypertext. In *Proc. Hypertext 2001*, Aarhus, Denmark, Aug 14-18, 2001, ACM Press: New York
29. Shipman, F.M. and Marshall, C.C. Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems. *Computer Supported Cooperative Work*, 8, 4, 1999, pp. 333-352
30. Walton, D.G. *Argumentation Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates: Mahwah: NJ, 1996