

Mapping the multimodal genres of traditional and electronic newspapers

John Bateman, Judy Delin and Renate Henschel

To appear in: Royce, T. & W. Bowcher. eds. (in preparation) *New Directions in the Analysis of Multimodal Discourse*. New Jersey: Lawrence Erlbaum Associates.

1. Introduction

1.1 Introduction: the notions of ‘genre’ and ‘genre mapping’

This paper is concerned with the application of the functional linguistic concept of ‘genre’ to multimodal documents—that is to documents whose total meaning is constructed out of combinations of text, layout and graphics/images. Although the term ‘genre’ has been used with respect to multimodal documents for many years, this use has generally relied on the everyday dictionary sense of the term as inherited, in an increasingly weakened form, from traditions of literary analysis and rhetoric. Here we relate genre to the more formally specified and technical usage of linguistics and, particularly, of functional linguistics. We show how the phenomenon of multimodality requires an extension of the term, not only over and above traditional usages but also with respect to the more technical constructs hitherto employed within linguistics.

Our particular motivation in attempting this extension in the meaning of genre is as follows.

Genre within functional linguistics attempts to provide a theoretical mechanism that is both predictive and explanatory. Membership of some text in a genre allows firm predictions to be made concerning the linguistic details of that text and, conversely, the occurrence of particular combinations of linguistic features can in turn be a strong indicator of genre membership. This, coupled with the now common interpretation of genre as a socially significant activity, provides a theoretically and practically useful link between social context and language.

The existence of this link has itself an important consequence. Both social context and language are not made up of discrete unrelated instances: both context and language must be considered as varying continuously and simultaneously along a variety of dimensions of description. By virtue of the link between them, however, we can characterize the variation best as *co-variation*—when the context changes, the language changes. Moreover, the changes are not random but are systematic. Particular changes in context, e.g., variation along an interpersonal dimension of a more or less formal situation, or variation along a textual dimension of more spoken or more written, each give rise to particular systemically conditioned changes in the language that occurs.

In moving to multimodal genres, therefore, we are led to consider the *system* with respect to which such genres vary and the consequences for realization—in language, layout and graphics/images—that such variation gives rise to. The systematicity involved makes it possible to ‘map’ multimodal genres, to show how diverse genres are related to one another in patterns of similarity and difference. This is to consider multimodal genres not as ‘vocabulary’, but as grammar: not as idiosyncratic arbitrary signifiers, but as instances akin to ‘grammatical forms’ related to one another productively by the system of language.

It is only when we can place individual genres against the general background formed by their possible dimensions of variation, the ‘map’ of their territory, that we can also go on to discuss ‘emerging genres’, genre blends, and even genres that do not yet exist but which might. And this

is of crucial importance in the area of multimodal documents. New genres are now forming very rapidly as new forms of expression are attempted with electronic documents; old genres are changing under the influence of these new genres, while new genres also appropriate and revise existing forms. Such complexities cannot be addressed by frameworks in which individual genres are considered isolated from their context and from each other. This makes a more systematic framework for addressing genres individually and collectively essential.

1.2 The example application: newspaper front-pages

The particular area for genre mapping that we will illustrate in this chapter will be drawn from newspapers and, in particular, newspaper ‘front pages’. We will consider some extreme points in the space of possible newspaper genres, varying the mode of presentation across traditional print-based and web-based products. We will end with a discussion of some perhaps intrinsic weaknesses of some classes of newspapers and news presented via the World-Wide Web, thereby showing the importance of genre considerations.

There are several reasons for adopting newspaper front-pages here as the object of discussion. Front-pages have long been a particular area of study and concern in both design and journalistic research and are commonly called the ‘show windows’ of their respective newspapers; any front-page thus works hard to establish the individuality of its newspaper. At the same time, newspaper designers have to be extremely careful about innovation and experimentation in the front-page as it also acts as the point of reliable identification for its regular clientele. Front-pages are therefore inherently ‘conservative’ and relatively long-lived—their design does change, but not quickly and not without careful consideration. Such pages are then ideal from the perspective of investigating regular patterns of realization and expression at all levels, a precondition for uncovering genres.

In addition, due to how newspapers are constructed and what they have to do, it is possible to see a rather close affinity between newspapers and web-based information offerings. In his study of, and recommendations for, newspaper design, for example, Stephen Ames (1989) cites from an interview with Randy Stano, then graphics designer for the Miami *Herald*, as follows:

“... we’ve got to make it easy for the readers. ... [Most readers] want news quickly, easily and they want a lot of it. .. Yes, you get people who want the New York Times ... type stories, but you’ve got to be able to break these stories up.” (1989:27)

This is to be achieved by breaking longer stories into packages and sub-packages using subheads, pulling out quotes to be set in headline type as graphic elements, making a portion of the story into a sidebar. As Stano summarizes: “The reader needs as many points of entry into a package as are feasible.”

This clearly echoes discussions of the kinds of non-sequentiality commonly targeted in discussions of hypertext in general and the web in particular (cf. Landow, 2001). It suggests that it should be relatively easy to transfer a print newspaper genre to an electronic version. Breaking longer stories into small components that are nevertheless linked via the resources of hypertext sounds an ideal development offering the best of both worlds. Not only are small components achieved, with many points of potential access, but the “[t]he interminable turns from the front page to inside pages in [some US newspapers]” which “would annoy European readers” (Giles and Hobson, 1990:216) could be turned into an advantage as a ready made opportunity for hyper-linking.

We will see below that this kind of adaption to the new media is precisely what many electronic newspapers in fact offer. We will also see, however, that certain inherent properties of the genres involved can compromise this adaptation.

1.3 The goals of the present paper

As we have suggested above, our main goals in this paper all revolve around genre mapping. We will show how, in the context of a particular research project—the ‘Genre and Multimodality’ project GeM, we can set about describing multimodal genres in sufficient detail as to take the next step of looking for empirically motivated characterizations of their ‘generic’ properties. But we also want to go further to suggest something of the utility of adopting a functional interpretation of genre. The multimodal genres that we describe are not simply collections of formal patterns, those patterns are selected in order to achieve the particular social activities necessary for the genre. In long established and evolved genres, the forms of expression are of necessity well suited to their generic purposes; with newer, artificially developed genres this may not be relied upon to be the case. Thus, despite the apparent suitability of electronic publication of newspapers in hypertext form, we will see that there are still problems and that some prominent web-based news offerings cannot be hailed as unmitigated successes, even though they conform well to traditional web design guidelines and usability recommendations. Our final goal in this paper will therefore be to show that this situation arises out of a conflict between the requirements of the newspaper genre as such and the constraints within which the designers of the web offerings have operated. All such design and re-purposing must then as a consequence thoroughly understand both the genre that is having its mode of presentation changed and the semiotic potentials of both the original and the new delivery methods—in short, what genres those methods can support and how.

2. Extending the notion of genre to ‘multimodal’ genre

Although the term ‘genre’ has an everyday sense and is regularly and widely used informally to refer to some recognizable kind of text—including inherently multimodal documents such as newspapers, film and so on—it also has a number of more technically defined meanings which are of more concern to us here. Genre has always been concerned with capturing or classifying regularities in language use. As Martin writes:

“Genres are how things get done, when language is used to accomplish them. They range from literary to far from literary forms: poems, narratives, expositions, lectures, seminars, recipes, manuals, appointment making, service encounters, news broadcasts and so on. The term genre is used here to embrace each of the linguistically realized activity types which comprise so much of our culture.” (Martin, 1985:250).

Correspondences between ‘activity types’ and language have long been noted and form the basis of the study of register; more recently, genre has been found by many researchers to be a useful construct in that it goes further than register and introduces *structure* into the relationship. Activity types involves various components—i.e., an activity is generally staged, it needs to be accomplished in sequences of particular sub-activities, each performing or continuing to perform various communicative tasks. These stages also leave their mark on the linguistic forms by which they are constructed. Thus, while moving through a genre, quite different constellations of linguistic selections may be found: both their internal constitution and their occurrence together in particular sequences provide evidence as to which genre is being employed. This is observable very clearly in early linguistic approaches to narrative, such as that of Labov and Waletzky

(1978), which decomposed narrative along lines familiar to traditional rhetoric while also raising the challenge of identifying distinctive groupings of linguistic features for each of the stages proposed.

Genre has therefore come to be defined as a linguistically distinctive, staged activity carried out for some recognized social purpose that is achieved by using language; the following definition by Swales in this vein has become almost standard and is often cited:

“A genre comprises a class of communicative events, the members of which share some set of communicative purposes. These purposes are recognized by the expert members of the parent discourse community, and thereby constitute the rationale for the genre. This rationale shapes the schematic structure of discourse and influences and constrains choice of content and style. ... In addition to purpose, exemplars of a genre exhibit various patterns of similarity in terms of structure, style, content and intended audience.” (Swales, 1990:58)

Several genre theorists have subsequently built on this basic model. Genre can then be used to explain and predict both the overall macro structure of some linguistic product—e.g., the unfolding of a narrative, or of a service encounter, or of a news article—and particular collections of linguistic features that may occur in each stage of that product. And all of these properties are in turn related to, and motivated by, the communicative purpose of the linguistic product as a whole.

This kind of definition already establishes genre as a useful construct for explaining linguistic regularities across a broad range of levels of description, but it can also be criticized as presenting too static a view of genre, one which merely provides a set of text building blocks that must be put together in a particular sequence in order to get a ‘well-formed’ exemplar. Deviation from the specified sequence then results in ‘less prototypical’ exemplars. To counterbalance this interpretation, genre can also be viewed in terms of *process* rather than in terms of product. According to this view, genre presents a repertoire of generic strategies for constructing discourses (cf. Martin, 1985; Kress, 1993; Lemke, 1999); language users select and combine such strategies from the available stock. In Kress, Jewitt, Ogborn and Tsatsarelis’s (2001) discussion of teaching and learning multimodality in schools, for example, the authors describe how students go about re-presenting in written textual form information that was originally presented multimodally. The students’ decisions can be seen in terms of flexible selections of genre, and thereby of generic strategies, for communicating the required information.

Most interestingly for us here, Kress *et al.* draw attention to the fact that distinct genres in fact have different potentials for creating various kinds of meanings: not all meanings can be expressed in all genres. This leads to a consideration of the respective “semiotic affordances” of the genres selected (Kress *et al.*, 2001: 144). ‘Affordance’ is used here in the sense of Gibson’s (1977) ‘ecological psychology’ and refers to the inherent guidance for use that *artifacts themselves* provide and in terms of which their users perceive them: thus, a door handle may be perceived in terms of the activities of turning and pulling and door-opening that it supports, while a door may be perceived in terms of the activities of passing through it or not that it supports. Extending this notion to semiotic affordances then includes social semiotic artifacts such as genres: a genre might be defined in terms of the social communicative purposes it allows to be achieved, the meanings that it can carry. And differing genres may have differing affordances.

We will employ the notion of affordances below when we come to consider what has been done with electronic presentations of newspaper editions. Important at this point, however, is first to

bring out a corollary of the view of genre as a semiotic potential for making meaning—for this to be the case we must also consider individual genres not in isolation but as instantiations of particular options drawn from the general genre resources of a culture. Describing genre then moves from a description of individual genres to a mapping out of the overall space of semiotic potential that those individual genres occupy. Although various models have been proposed for exploring this space systematically (cf. Lemke, 1999; Martin, 1997), an empirical anchoring for the mapping task is always essential. To achieve this, studies of genre must consider collections of linguistic objects in order to tease apart the possible dimensions of similarity and difference on the basis of the actual variation observed. Our own work is carried out within this tradition.

This can be taken further, however, by extending the range of ‘linguistic objects’ admitted for investigation. In a series of books and articles, Kress and van Leeuwen have shown convincingly that linguistic modes of analysis can usefully be applied to a much wider range of ‘texts’ than is traditionally targeted within linguistics (cf. van Leeuwen and Kress, 1995; Kress and van Leeuwen, 1996, 2001). Images, photographs, pages containing combined texts and illustrations in particular layouts, and even entire documents can all be considered as ‘texts’ in a strict linguistic sense. In terms of functional linguistics, this means that all elements that occur in such multimodal texts can be analyzed, related to one another, and interpreted in terms of their selection/construction from a collection of multimodal semiotic resources and in terms of their contributions to the communicative and social function of the whole. The meaning of such ‘texts’ is then not to be found within any single contributing mode alone, but only in their composition (cf. Lemke, 1998; O’Halloran, 1999).

From this point the move to **multimodal genre** is a natural one. We consider genres to stretch over both traditional ‘monomodal’—although, a point now made increasingly often, there is really no such thing as a strictly ‘monomodal text’—and semiotic products composing contributions across various semiotic modes. And, just as with text genres, it is necessary to consider the overall space of potential within which such products live and which gives each individual product a particular value or meaning. Our hypothesis, to be investigated empirically, is that multimodal genres exhibit just as high a degree of regularity in their relationships with each other and with context as do traditional genres. And, what is more, it is crucial to obtain a sense of the underlying space of possibilities because new genres are now being formed multimodally at an unprecedented rate. As the means of production of multimodal artifacts is opened up to become ever more accessible (via word processors that offer control over typography, layout and images), the range and extent of experimentation (both deliberate and not) increases dramatically. Only by setting out the underlying space will it be possible to make sense of this development and so to characterize and explain the genres resulting.

3. Mapping the genre space: the GeM project

Given our goal to explore empirically the semiotic space of multimodal genres, it was necessary to set out initial layers of description by which our ‘data’, the multimodal artifacts to be analyzed, could be characterized and compared. This was the main task of our UK ESRC-funded project GeM (‘Genre and Multimodality’: <http://purl.org/net/gem>). During its life, the GeM project pursued the definition of suitable layers of description, their formal specification as a basis for constructing a corpus of multimodal documents conforming to the most recent recommendations and standards for linguistic corpus design, and an initial collection of documents analyzed in the manner proposed and then stored in the corpus. More detailed accounts of the motivating design decisions for the corpus are given in Bateman, Delin and Henschel (to appear); there we explain

the reasons for providing an annotated corpus, rather than just a collection of selected documents. We will take this aspect of our work for granted here, although it is nevertheless important to note that the examples we discuss are drawn from this corpus and the discussion relies crucially upon the annotations that the corpus maintains. We are now continuing to expand the corpus collection, the layers of description, and the genres covered, as well as performing further explorative analyses of the data gathered. The consideration of newspapers discussed here is an example of one of these latter explorations drawing directly on the corpus.

The design of the corpus is multi-layered; i.e., each document is analyzed and annotated according to several independent descriptions. We then seek generalizations concerning multimodal genres by examining whether particular co-selections from the independent descriptions can be reliably recognized. If patterns of co-selection are revealed, we can assume that the theoretically independent choices are in fact constrained so as to preclude or restrict independent choice. The controlling variable in such cases is sought in a notion of genre. The purpose of the multi-layered annotated corpus is therefore to support more focused qualitative and quantitative research into the systematically recurring features and patterns of multimodal genres.

Our selection of description layers draws both on the linguistic tradition and on accounts of document production and design.

For the former we incorporate the functional linguistic accounts of genre mentioned in the previous section, in particular Martin's and Lemke's functional generalizations of genre to take in genre spaces and their social interpretation. For completeness, we also assume a more or less complete linguistic annotation as already commonly adopted for annotated linguistic corpora, tree banks and the like. Moreover, we include a rhetorical layer of description drawing on Mann and Thompson's (1988) Rhetorical Structure Theory; this is motivated by the fact that discussions of design often include consideration of 'rhetoric' and communicative purpose and so it is useful to have an explicit representation for this aspect of a document's meaning. With this in place we can ask questions concerning the selection of particular forms of expression for particular rhetorical configurations.

For document production and design, we were particularly inspired by the model of document description developed by Waller (1988). Although centered broadly on typography and design, Waller draws extensive parallels and connections with linguistic models of communication and so already goes a considerable way towards providing a model compatible with the other layers of linguistically motivated description that we include. In general, we have extended and replaced those of Waller's layers concerned with linguistic description, while adopting those layers concerned with design and layout. This results in two broad areas: features concerned with the properties inherent in the document being analyzed and features describing aspects of the conditions of production/consumption of the document. These areas are introduced in more detail in Delin, Bateman and Allen (2002/3); here we simply list the layers used in the current corpus along with brief descriptions. This list is not intended to be exclusive—there are many further aspects of multimodal documents that it would be useful to include—we do claim, however, that any account that does not include *at least* these kinds of information will be unreasonably restricted in the range of multimodal phenomena that it can address.

Content structure

the content-related structure of the information to be communicated—i.e., the propositional content or, using register theory, the 'field' of discourse;

<i>Rhetorical structure</i>	the rhetorical relationships between content elements: i.e., how the content is ‘argued’, divided into main material and supporting material, and structured rhetorically;
<i>Layout structure</i>	the nature, appearance and position of communicative elements on the page, and their hierarchical inter-relationships;
<i>Navigation structure</i>	the ways in which the intended mode(s) of consumption of the document is/are supported: this includes all elements on a page that serve to direct or assist the reader’s use of the document; and
<i>Linguistic structure</i>	the linguistic details of any verbal elements that are used to realize the layout elements of the page/document.

The second area, drawn almost directly from Waller, are:

<i>Canvas constraints</i>	constraints arising out of the physical nature of the object being produced: paper or screen size; fold geometry such as for a leaflet; number of pages available for a particular topic;
<i>Production constraints</i>	constraints arising out of the production technology: limit on page numbers, colors, size of included graphics, availability of photographs; for example, and constraints arising from the micro-and macro-economy of time or materials: e.g. deadlines; expense of using color; necessity of incorporating advertising; and
<i>Consumption constraints</i>	constraints arising out of the time, place, and manner of acquiring and consuming the document, such as method of selection at purchase point, or web browser sophistication and the changes it will make on downloading; also constraints arising out of the degree to which the document must be easy to read, understand, or otherwise use; fitness in relation to task (read straight through? Quick reference?); assumptions of expertise of reader.

Taking the conditions of production for multimodal artifacts seriously enables a more realistic appraisal of the precise motivations and reasons for the appearance of documents. Design is often described as a compromise between many competing and sometimes conflicting constraints and these need to be brought into any discussion of the functional motivation of the resulting artifacts. One aspect of the GeM project was therefore explicitly to consult with the designers of the documents included within the corpus in order to make the path from design to product as explicit as possible.

This issue is equally relevant for print and web versions of newspapers. All newspaper versions are subject to stringent production constraints ranging over the provision of material, extreme time pressure for decision-making, conforming to the generic design constraints that identify each issue of the newspaper as a further example of that newspaper as such, and so on. Such constraints can make themselves felt in the smallest details of a document—for example, page layout can be a result of negotiations between publisher and advertisement companies and not a

matter of strictly functional design at all, or the need to produce a particular length of headline may restrict linguistic choices right into grammar and lexis. Moreover, the technology used to produce a particular multimodal document can also impose its specific footprint on the artifact as a whole—for example, as print technology changed newspapers began to use color, but in different ways as the technology developed; analyzing the use or non-use of color images at that time would need to know this boundary condition. Not being aware of the production constraints always opens up the danger of an over-interpretation because more design freedom, and hence controllable resources for making meaning, are assumed than are actually available. More of the relevant production and consumption constraints for printed and electronic newspaper versions will be discussed explicitly below.

The relationships between social activity, production constraints and our extended and revised notion of multimodal genre can be depicted as shown in Figure 1.




Figure 1: the relations between social practice and multimodal genre

Social practice determines that certain kinds of artifacts, constructed with particular technologies, will be used semiotically: examples of this are the papyrus scrolls used by the ancient Egyptians, bound books, and websites. All genres relying on some particular types of artifact are therefore constrained by their material basis: certain choices are no longer free. Subsequently, however, social practice can further restrict how the constructed artifacts will be deployed: for example, texts written on a scroll can be written down the scroll following its length so that it forms, as it were, one long page, or they could be written in self-contained segments, corresponding to pages oriented at right angles to its length; both options are compatible with the technology of papyrus scroll production but which option is taken up is a matter of convention. Many examples of a similar variability can be found in newspapers. With the advent of printing presses that did not require that the type be prepared in narrow columns, it was then technologically possible to have wider columns or no columns at all. However, although some newspapers used this new freedom when it became available, many did not. Thus, an option may be employed by a culture as a meaningful genre choice but it need not be.

This slippage between what is technologically possible and the use that is habitually made of some artifact gives rise to what we will term ‘virtual artifacts’—that is, artifacts that are treated by an artifact-producing group as a single technological product although in fact the properties of that artifact are a blend of technological imperative and cultural habit. Most artifacts are in fact of

this kind: although the technological basis of an artifact may change, the use made of that artifact is strongly determined by social practice and so may be quite resistant to the changing technology. The ‘canvas’ seen by the designer is then already a technologically and socially constructed artifact, a virtual artifact rather than a purely material one. It is quite possible nowadays to produce books in which the print runs down the page instead of (as on this page) top-to-bottom, but English books are not going to take up this option while Japanese books generally will. This can have a significant impact on the layout options taken up on a page (cf. Waller, 1988:232). Such layout decisions are then being taken with respect to distinct virtual artifacts—it is not a free designer choice which direction of print is adopted, but it is no longer constrained by the technology (and so may become controlled by genre differently in the future). This then constitutes an important component of the generic constraints imposed by a genre, one which extends beyond form, structure and content, to include uses of the material substrate.

Our inclusion of the conditions of production among the levels of description and variation that must be included in genre, and particularly in multimodal genre, also finds parallels in Kress and van Leeuwen’s (2001) recent discussion of multimodality *per se*. Whereas the analysis of ‘purely’ linguistic products has managed by and large to bracket out issues of the materiality of the carrier of the product—even in the case of spoken language relegating many aspects to ‘para-verbal’ accompaniments to the linguistic sign proper—it seems that this strategy cannot be maintained convincingly for multimodal products. To address the materiality of the sign, and the fact that this can itself be made to do ‘semiotic work’, Kress and van Leeuwen now also extend the traditional linguistic levels of description; in particular, they extend the notion of *linguistic stratification* in two steps:

“The basis of stratification is the distinction between the *content* and the *expression* of communication, which includes that between the signifieds and the signifiers of the signs used. As a result of the invention of writing, the content stratum could be further stratified into **discourse** and **design**. As a result of the invention of modern communication technologies, the expression stratum could be further stratified into **production** and **distribution**. ... [W]e argue that production and distribution produce their own layers of signification.” (Kress and van Leeuwen, 2001:20)

Although space precludes a detailed comparison of this articulation of semiotic resources with the account we have sketched in this section, the similarities are striking. It is increasingly clear that issues of production and consumption are necessary components for any serious account of multimodality.

Thus, we suggest that genre be defined as regularly occurring selections of features—linguistic, layout, rhetorical, content, navigation, canvas (virtual artifact), production, and consumption—that may be linked with particular social practices. Genres show how a collection of semiotic resources pattern together in order to fulfil recurrent, recognizable social goals. Moreover, in Lemke’s (1999) terms, “Co-generic texts are privileged intertexts for each other’s interpretation. When we make meaning with any text, we construe meaning relations that also depend on these other texts.” We take this co-genericity to arise out of all the modalities present in a text and assume intertextual interpretation to apply equally across documents that are richly multimodal.

4. An application to newspapers and newspaper design: sketching some properties of the genre space

When we wish to establish the details of multimodal genres, rather than setting out their theoretical status, we are faced with a very practical and strictly empirical concern: only when sufficient signifying features have been located will it be possible to be confident that true ‘genres’ have been isolated. This process of investigation is intended to gradually reveal more detailed aspects of the ‘landscape’ within which genres are situated: rather than seeing genres as isolated constructs, such as found in lists of unrelated text types, they are turned instead to instances or examples drawn from a known landscape of possibilities. When we have mapped more of this landscape, we can make more reliable predictions about new or emergent genres that occupy nearby, but as yet unvisited, locations.

For our discussion of newspaper front-pages and their allocation to genres, we will focus mainly on two examples that lie at extreme ends of a dimension of variation. This will make the comparison easier but it should not be forgotten that there are also a range of further related genres in between these extremes. We will describe some of these in passing but cannot show them with the same degree of detail as our main examples. The primary examples are the print and electronic versions of the *Guardian* newspaper, one of the sources of documents for the original GeM corpus.

Adding documents to the GeM corpus for subsequent analysis requires that the documents be analyzed and coded so that there is an explicit representation of the results of analysis along each layer of description the corpus requires. The basic preparatory steps to be followed are always the same. First, it is necessary to identify all the distinct visual elements units occurring on each page of the document to be added; this defines a set of uniquely labeled ‘base units’. The labeling of base units enables any unit to be cross-classified according to the other layers of description as appropriate for the document at hand. Next, independent analyses are carried out of each page according to its visual decomposition, i.e., its layout structure. And finally, the rhetorical organization and navigational structures employed by the document as a whole are analyzed. Both rhetoric and navigation are not in general restricted to individual pages, although they are carried by elements that occur as layout elements within pages. Both electronic and print documents are treated in the same way. For websites, each web-page is treated as a single unit just as is the case with print documents.

For the present discussion, we will focus primarily on the layout and navigational structures. The layout structure for a page is created as follows:

- (1) Working visually from the page, decompose the objects on the page in terms of their visual unity and relative mobility/independence,
- (2) Transform the page decomposition into a hierarchical structure, by seeing which collections of layout units group together and form components of larger visual elements,
- (3) Specify presentation information for units: e.g., font size, type, color, image type, resolution, etc.

The purpose of this analysis is to successively capture the spatial information implicit in the page’s visual design as explicitly represented structural information. Visual perception automatically places visual elements in collections of more or less closely related groups and delivers a range of similarity and difference judgements. The construction of a layout structure hierarchy in step (2) enables this kind of grouping information to be expressed succinctly, while

also providing useful hooks for stating generalizations concerning presentation styles. If, for example, a layout unit consists of a collection of sub-units, all of which have the same type face, the same color, the same indentation, etc., then this information need only be expressed once for the parent layout unit. The sub-units inherit all relevant information from their parent that they do not themselves explicitly state. This corresponds to the formal models of layout used in such standards as the World-Wide Web consortium's Cascaded Style Sheets (CCS) or the Formatting Objects of the extensible style sheet language (XSL: cf. <http://www.w3.org/>). An additional component of the model then provides precise and relative positioning information for the layout elements. Once this layer of analysis is complete, a complete rendering of the spatial organization of the page is available. This can be processed in order to ask very specific questions concerning the position and structuring of layout units. Such questions are typically combined with information from the other layers of analysis in order to frame particular research hypotheses. Below we will do this in combination with the navigation elements on the page.

The navigational structure for a document is constructed quite differently. Here all identified base units that serve a 'linking' function are isolated and classified. Such links can be within particular layout elements internal to a page, or external to a page serving to link more distantly represented portions of a document. The links can also be of a variety of types and serve to link different kinds of material: for example, we may have links introduced in order for a reader to construct a linear reading path from content that is distributed across different layout units—such as occurs in newspapers when articles are broken, or 'turned', across several pages—or, alternatively, we may have links which serve to express additional rhetorical relationships between content units available in a document—such as providing background information in footnotes or hyperlinks. The connectivity of a set of links can then be abstracted out of the document to give a web of navigational paths—this is the navigation structure for the document.

A graphical depiction of the results of performing the layout structure analysis for our two primary examples is given in Figure 2. Here we see two selected pages—one the print version¹, the other electronic—each followed by a diagrammatic rendition of the layout structure. These graphics show the layout elements according to step (1) of the analysis with an indication of hierarchical structure given by embedding. Occasionally we have over-emphasized the space around the borders of the layout units in these diagrams in order that information concerning relative embedding not be lost. The actual layout structure is represented explicitly using the extensible markup language (XML) and so the diagram is purely a convenience for showing the structure involved.

¹ The *Guardian* print front-page shown is from February 2nd, 2000; it differs in some minor ways from the actual page as it went to press on that day however. The rendition used was created directly from versions of the original Quark Xpress files as made available to the GeM project and is not a scanned version of the published newspaper. Both print and electronic versions are used by permission.






Figure 2: A front-page of the Guardian and the Guardian on-line and their layout structures (abridged)

The texts provided in the online version of the *Guardian* generally correspond to the texts found in the print version, although there is considerably less use of photographs. When an electronic version and a print version of a newspaper do not share significant portions of their production processes, converting the photographs from the print version into a form suitable for web-presentation can constitute a significant overhead. What is very different across the two forms, however, is the presentation of their information. The print form should be familiar in style to all readers and so we will not describe it in general terms here. The online version provides clear navigational menus, short summaries of the news articles, and general navigational menus for thematic sections. There are also links to archival materials as well as pointers off-site to related sources of information, such as government and organizational websites. The layout structure shows that the overall organization of the web version is rather straightforward, with repeated use of structurally similar sub-elements. The line lengths are kept very short, also as recommended for online text, and there is little visual adornment. The advertisements are kept to standardized positions: one band across the top and one across the bottom of the web page. Thus, considered in isolation, the *Guardian* website might be considered an appropriate adaptation of the newspaper genre for web use.

In order to assess this more effectively, however, we need to place the document against the other possible forms with which it contrasts. This is one of the tasks that setting out the genre space supports. We noted above that we have deliberately picked two document pages that represent extremes and this is motivated by our goal of mapping out the space of possibilities. A brief look at the websites of the *New York Times*, the *Financial Times*, the *Sydney Morning Herald*, the *Daily Telegraph* and many others, shows a very different electronic identity to that of the electronic *Guardian*, one which in fact lies somewhere between the front-page style of their own printed versions and the electronic form employed for the *Guardian*. Moreover, when we consider this broader range of newspaper web-editions in terms of their layout structures and the presentation of components within those structures, such as headlines, articles, and so on, one of the main variables appears to be how many of the techniques of the print editions are also employed in the web editions. That is: how much the genres of the print version and the genres of the electronic version influence one another. The *Guardian* is minimal in this respect—as Figure 2 shows, its layout organization bears very little resemblance to that of the *Guardian* print edition.

We illustrate this here for the other newspapers only very briefly. Although our discussion will be informal, it should nevertheless be seen as resulting from an analysis of the respective layout structures in precisely the same way as that shown above in more detail for the *Guardian*. The reader is referred to the respective websites of the newspapers discussed for concrete examples of the general descriptions given here; these descriptions apply to the newspaper web pages as served during 2003/4.

First, and in many respects closest of those we consider here to the *Guardian* electronic version, is the web offering of the *New York Times*. This may strike the reader as surprising when first confronted with the *New York Times* web page as it looks quite different to that of the *Guardian*. There is considerably less white space and very different decisions have been made for type faces and font sizes. These combine to create a quite distinct visual identity. But focusing more on the layout and navigation structures yields a rather different picture. The main news content of the page is presented in the center of the page in a single column. Individual articles are clearly separated by headlines and white space as with the *Guardian*. The principal access structure provided for the reader is therefore essentially a vertical list of news article summaries with links

onward to the full articles. Nevertheless, the headlines are significantly larger than those of the *Guardian* and also take over more of the features of headlines from print newspapers, such as by-lines and location information. The page also has a standard navigational column on the left with links to newspaper thematic sections, although a less prominent third column on the right also takes over some of the navigation work.

Somewhat further away from the *Guardian* extreme is the *Sydney Morning Herald*. Although this newspaper also employs a main single column of articles grouped into thematic categories vertically, it adopts within this column large headline fonts of at least three different sizes ranging from the main headline, through several articles with a medium headline, through to final comparatively small headlines. All of these are larger and more striking than those found on the *Guardian* page, however, and shows more variation than the larger headlines used within the *New York Times* webpage. Visually, the headlines form much more prominent layout units than the corresponding elements on the *Guardian* website. Moreover, there is also a column of less significant articles on the right-hand side of the page as well as the usual navigation links on the left-hand side: this is then a deviation from the predominantly vertical access structures that have been seen so far. The reader has a choice of two columns from which news articles may be accessed.

Moving further in the same direction, the *Daily Telegraph* webpage has two full main columns in the center of the page containing news items. Both columns are equally prominent in the layout structure; one is not clearly subordinate as in the previous case. What is more, the columns each use further distinctions of background color, thematic section heads, separating lines and thumbnail photographs in order to demarcate the stories. The headlining news section occupies the top of the left-hand column and is demarcated by a yellow background. A large prominent picture generally occupies the top of the right-hand column. This selection of layout realizational features expresses not only that there is no clear domination relationship between the two columns but also that they are informationally distinct: different kinds of information are being presented.

Next, the web offering of the *Financial Times* can be described as combining certain aspects of the *Daily Telegraph* and the *Sydney Morning Herald* in that it uses different size fonts for headlines and its central left-hand column is more salient (twice as wide) than the right-hand column. It also moves on into a third content column, filling this with further information fragments such as stock market overview, services offered and, as is often the case, advertisements. The narrower right-hand column is not however used for clearly less significant articles in that its headline fonts and presentation style of articles is identical to those occurring on the left-hand news column and it can even be headed by a prominent photograph. This then brings in the print genre tradition of presenting articles over columns of varying widths.

The last two examples move us much nearer their respective print editions. The web version of the *Sun* newspaper attempts to break up the strict column organization visually just as does its print versions. Interestingly, however, the online version contains significantly more information than the print version—which is often made up almost completely of self-advertisement in the form of competitions and prizes. In the web version, these components take up a much smaller position in the layout structure, more comparable to the banner advertisements seen in other online newspapers. In terms of the layout structure involved, there is a table with two columns but the rows are presented visually just as prominently as are the columns. The vertical access structure is very much weaker therefore.

Finally, the feel of the print and electronic versions of the *News of the World* are very similar—they have similar layout structures and adopt similar rhetorical and content decisions. The web version does have its usual left-hand navigation bar but apart from this the front page is composed of pictures, short text fragments and pointers into the rest of the website. This is the least ‘compromised’ adaption of the print newspaper to web-distribution within this kind of document.

This exploratory mapping of certain dimensions of the space of news web edition genres can be summarized as follows. We can see two broad genres that are being combined when forming web editions: one, which we will characterize loosely here as ‘explicitly web-based design’, and another, that of newspaper design. Explicitly web-based design is a new emergent genre that is inherently unstable at this time: it is a product of a virtual artifact whose technological basis is undergoing extremely rapid development. The essence of explicitly web-based design as we consider it here can be easily suggested by looking at web-based documents from the early days of HTML: the technology at that time was so restricting that it was difficult for layout decisions to be made at all (indeed, according to the original intended use of HTML, layout was explicitly to be excluded from the document and moved to be one of the tasks of the browser). Web design guidelines were formulated in order to hold the damage in check—that is, given that the web pages could not use sophisticated layout and that the readers’ web browsers might not even be capable of displaying appropriately the design decisions attempted anyway, certain limitations in design were eminently desirable. Moreover, given the well-known tendency for users of hypertexts to become disoriented, it was considered important that web pages provide as clear as possible indications both of what *is* a link and what those links might lead to.

The most direct way of achieving at least some reasonable layout structure became the HTML table environment. This provides a simplified grid structure analogous to that commonly used in document design (cf. Hurlburt, 1978). Although technologically supporting arbitrarily wide and deep tables, consumption constraints in the form of readability results naturally suggest that columns containing text be narrow rather than wide; in addition, the columns can be quite deep since the user of such a document is expected to scroll down while reading. Information items within the table can be discriminated either by making the row structure of the table explicit (by explicit boundary lines, or color discrimination, etc.), or by employing other textual/visual means within each column as separators (by extra white space, lines within columns, typographical emphasis, differential background colors, etc.). Further grouping can be achieved by employing tables within tables. The virtual artifact resulting from web technology plus table environments is almost precisely the canvas within which the Guardian web edition is played out. This is why the layout structure shown for the example page in Figure 2 can be constructed very straightforwardly from HTML table elements: this is almost how it is constructed (although with considerable use of transparent spacing images for extra white space, explicit lines as demarcating elements, etc.). The result is a very attractive uncluttered webpage, but is it a news webpage?

Comparing the organization of other web pages with similar requirements, i.e., fulfilling the communicative goal of providing access to a wide range of diverse information in a maximally clear and uncluttered fashion supported by strong hierarchies of links, we find many similar documents—not all of them newspapers. There are even strong similarities in the layout structures of the *Guardian* webpage and completely different kinds of documents such as those of magazines, public information sites, and lists of hits from search engines. This supports the allocation of the *Guardian* web edition to a broad family of electronic web-based genres that

overlaps, but is far from co-extensive, with the family of web-based newspapers. This again shows the necessity of considering genres in terms of the space of possibilities—we can move from purely information-driven and largely automatically produced documents, such as search engine results, through more organized and thematically demarcated presentations of information, through to documents that incorporate recognizable features of the printed newspaper genre, right up to web-based documents that mimic print documents. Each of these positions has consequences for the reception and utility of the document genre produced within it.

We have seen that many web editions of newspapers attempt to announce their genre membership as newspapers as well as providing access to structured information: they announce both that they have information and that that information is news. The web editions are still recognizable as web editions however. In contrast to this, there has recently been substantial growth for a further related document type: this is the distribution of online versions of the print editions of newspapers. Customers are offered access to the print edition of a newspaper either by a download service, whereby they receive the newspaper as a file that they can then read at leisure, or by special browsing software, which allows the paper to be navigated online just as it would be in its physical version. The suppliers of the software required for this manner of distribution, such as NewsStand and Olive Software, naturally claim that newspaper readers are much happier with this kind of document, even though proponents of web-based newspaper editions can be critical. Online editions present the most direct solution to the problem cited from Ames above of those “... who want the *New York Times* ... type stories, but you’ve got to be able to break these stories up.” Especially the more sophisticated online edition software, such as Olive Software, which provides active navigational hyperlinks for all the navigational links present in the print version (such as ‘Continued on p3’), support this kind of document very closely indeed.²

It is interesting to consider the reasons for the rise of this kind of newspaper document and its apparently positive reception—especially, contrasting as it does, with the suggestion raised in the introduction that it could be quite simple to export news to a web-based medium. The web edition of the *Guardian* is the document type among those we have considered here which is the most transparently ‘web-based’ with minimal influence from its print version. In a sense, the *Guardian* is most clearly a product of the web; it does very little to impose on the virtual artifact of the webpage strategies that have developed for print newspapers. But very few web editions of newspapers appear to be taking this path. Is the future of online web newspapers then more likely to be simply a return to their print form augmented by active links?

We conclude our discussion by addressing this question from the perspective of the genres involved. Rather than describing and classifying, we need now to consider what variation across different genres provides for the users of those genres. We will suggest that it is not simply a matter of deep conservatism among the producers and readers of newspapers that the electronic distribution of newspapers that look exactly like the print versions should be well received.

² For both the online editions of print newspapers and the web editions there are important production constraints that also play a significant role. To what extent the production of the online version can build on the existing process of the print edition, the perceived value of the product for the producers—as a news product or a way to increase circulation figures, the degree to which the product can be derived automatically rather than demanding its own design and production team, differences in the targeted audiences for the alternative products: all of these aspects contribute to what ends up as a very complex situation. We cannot do justice to any of these issues here, but simply point them out as necessary for any more complete understanding of the ongoing developments in this area. Here we focus only on the semiotic contributions of genres.

Instead we propose that there are inherent features of the news genre that strongly support the use of more aspects of the print editions. That is, even though there are still substantial consumption constraints that mitigate against the ‘online print edition’—such as download times and costs, low resolution screens that do not yet compete with print in terms of ease of reading, etc.—the genre itself has significant positive aspects that weigh significantly in its favor as the genre of choice for distributing news. The affordances of the print news genre cater to a precise need, one that is not met by the explicitly web-based design.

5. Presentation-specific weaknesses and strengths of the print and web editions when delivering ‘news’

The most striking feature of linked hypermedia as provided by the World-Wide Web is the very fact of hyperlinks. Consumption and access is dramatically and qualitatively altered for documents produced in this medium. However, it is still the case that these possibilities for access to a document’s content be presented to a reader in an understandable form. This is the work that is done by the increasing standardization that is to be observed in the distinct navigation bars and collections of displayed links that make up much of current webpages. The sheer quantity of available links has itself become a dominant organizing feature of webpage design.

We can see this very concretely from our analyses in terms of layout structure and navigational structure by explicitly cross-classifying the two contributions. We can ask how much of the layout structure is taken up with organizing navigation. The results are striking; taking our examples from Figure 2, the print edition of the newspaper devotes 11.1% of its available space to navigation, whereas the web edition devotes three times as much, 33.2%, to navigation. These figures vary somewhat across the related news genres, but the overall bias is clear: much less of the page is used for content, i.e., in the case of newspapers, ‘news’, in the web editions than in the print editions. In terms of the affordances of the genre, we can state that the web edition supports the selection of articles according to various topical classifications (e.g., ‘UK News’, ‘International’, ‘Politics’, ‘TV and Radio’, etc.).

Crucially, however, this is *not* where the print editions of newspapers place their emphasis. The newspaper front-page is designed around the socially constructed concept of *news values* (cf., e.g., Fowler, 1991:12). Where articles are placed on the page and how they are announced to the reader—e.g., with what kind and size of headline, with pictures, across how many columns, etc.—provide the newspaper consumer with a fairly exact statement of just how ‘newsworthy’ the newspaper considers the news in question. And, since newsworthiness is made up of values along several independent dimensions, the combined resources of layout positioning and typographical emphasis allow newspaper print editions to indicate their assessments of an article’s worth very flexibly and without a single strict ordering. For example, on the *Guardian* front-page, we see that there are clear major articles with the largest headlines, but there are also other articles with relatively large headlines, some placed on the left, some of the right, some taking up more space, some less, some with a prominent photograph accompanying them, some without, and so on. The visual presentation does not define a single track for the reader to take. In short, the semiotic resources available construe space and typography as multidimensional indicators of newsworthiness and importance.

Such multidimensional indicators demand that their consumers have flexible access to the material presented. Consumers need to be able to evaluate the competing claims for importance and newsworthiness, rank these according to their own interests and requirements, and jump to the selected articles. This is where the space provided by the newspaper front-page is central: the

articles along with their self-proclamations of greater or lesser importance in various ways are also given clear addresses on the page: they occupy particular layout elements that are immediately accessible via their demarcation from other layout elements. Thus, the cognitive load of access to selected elements is off-loaded to the spatial layout itself; the visual decomposition of the page into its layout structure already functions as an access structure. It is this, we suggest, that is responsible for the extreme effectiveness of layouts commonly found in printed newspaper design.

However, for this to work, it is necessary that the virtual artifact employed does not impose an overly restrictive dimensionality on the information presented. The two dimensional spread of information in a front-page combined with typography works well in this respect and is now, after 200 years of training, well understood by all newspaper readers. In contrast to this, we can postulate a main problem with the *Guardian* web edition style of presentation: there is only one dimension of presentation available, the vertical down-the-page sequencing of one news article after another. This might be used to reflect a single dimensioned evaluation of newsworthiness, were it not for the fact that this column is also broken up into topic areas. This clearly works against the use of the vertical dimension for newsworthiness, although it might be compatible with a rather rigid interpretation of news *relevance*. Thus, although the articles are highly inter-linked and a reader can follow many chains of connection very freely, the basic access structure provided by the combination of layout structure and navigation is completely different to that offered by the print edition.

Each of the other web editions of newspapers considered here can be seen as attempts to increase the dimensionality of their adopted virtual artifacts. Increasing the number of columns, diversifying the headline sizes and type faces, employing color differentially, are all techniques well established over the historical development of print newspapers. But for web editions this strategy is limited in several ways. The canvas constraints imposed on a screen-based medium by low resolution screens effectively limits just how small type faces can become and still be considered usable: the situation is analogous to earlier stages in the development of print and paper technology—if paper is too absorbent, then ink will smudge and render the printed text unreadable. This in turn constrains how many columns are going to be possible for a webpage, which in turn restricts the spatial layouts possible.

A further canvas constraint is generally considered an advantage of web-based documents, but here we see its negative aspects: it is generally obvious *what* one should do with a printed newspaper once we have read a page, that is, we can turn to the next page. With a web document there is no comparable affordance of the artifact itself—there is no physically determined ‘next page’. This means again that the navigational access structure embodied in the physicality of a book or a printed newspaper must instead be presented explicitly to a reader as information—i.e., additional layout units in a space already seriously challenged by the load it has to carry must be allocated to navigation. To counter this problem, web editions have various levels of ‘next’ buttons—for example, it is often possible to follow through the articles in a thematic section simply from one to the next. But here we are then faced with the additional cognitive load of an abstract path unsupported by the spatial organization that comes for free with the printed material object. It quickly becomes unclear just how much one has seen, where one has come from, how much is to go and what other alternatives may have been. It is also unclear to the reader whether the list is being presented with some judgement of newsworthiness, or is being presented with a list sequenced by time of posting, or some other non-apparent ordering criterion.

In short, the affordances of printed documents may still be being seriously undervalued in the design of multimodal ‘new media’ documents. When most of a web-page’s basic and most valuable resource, the ‘real estate’ of the page, is being taken to signal explicitly access structures that are self evident in the affordances of more traditional artifacts, the received utility of the new media artifact vis-à-vis the traditional artifact is much reduced. The virtual artifacts currently available for newspaper web editions may just not be able to:

- express multidimensional newsworthiness and importance,
- effectively use layout structure simultaneously as access structure.

These are then significant issues to be considered in the design of multimodal documents of emerging genres of any kind, not just newspapers.

6. Conclusion

We have set out some dominant sources of variability for an area of the multimodal genre space that is partially instantiated in newspaper websites and printed documents. We have further characterized documents in terms of the use that is made of their adopted layout structures. Within web documents of all kinds, a sizeable proportion of the layout elements of a page are devoted to navigational elements that assist their readers to access the document. With newspapers the layout structure and the typographical realizations of its elements is predominantly devoted to functions specific to the genre of news delivery. The emergent genre of newspaper websites is then currently not distinctively one of news delivery at all.

This weakening of the newspaper genre is also effected by the inclusion of other, generally very useful functionalities that are offered by the web editions of newspapers. The presence of off-site links to further information and extensive archival of articles is a substantial addition to the communicative purpose of the document. It is not, however, compatible with the communicative purpose of delivering news ‘as it happens’ or, at least, as soon as possible. Indeed, since archived articles can be reached with the same ease as current articles, it is often not even possible for a reader to ascertain instantly whether they are in current news or in an archive. This again strengthens the generic link of web newspapers to information websites of other kinds.

It remains to be seen whether web-based virtual artifacts that are more appropriate for newspapers will be developed or whether web-based newspapers will migrate to occupy a distinct informational niche. Both technological and design-based developments would be necessary for the former to succeed. This is not limited to the genres of newspapers however. Many genres use sophisticated layout and typography in ways similar to newspapers, i.e., as multidimensional access structures. A real challenge for new media will be to build on this and to create virtual artifacts that also provide this kind of affordance. Until then the genres that emerge will only be able to move within a rather restricted area of the genre space as a whole.

Bibliography

Ames, Stephen (1989) *Elements of newspaper design*. New York: Praeger Publishers.

Bateman, J., Delin, J. and Henschel, R. (to appear) Multimodality and empiricism: preparing for a corpus-based approach to the study of multimodal meaning-making. In: Ventola, Eija, Charles, Cassily and Kaltenbacher, Martin (eds.) *Perspectives on Multimodality*, Amsterdam: John Benjamins.

Delin, J., Bateman, J. A. & Allen, P. (2002/3). A model of genre in document layout. *Information Design Journal*, 11, 54--66.

- Fowler, R. (1991). *Language in the news: discourse and ideology in the press*. London and New York : Routledge.
- Gibson, J.J. (1977). The theory of affordances. In R.Shaw and J.Brandsford, editors, *Perceiving, Acting, and Knowing: Toward and Ecological Psychology*, pages 62--82.
- Giles, Vic and Hodgson, F.W. (1990) *Creative Newspaper Design*. Heinemann. p216
- Hurlburt, A. (1978) *The Grid: a modular system for the design and production of newspapers, magazines, and books*. New York: Van Nostrand Reinhold Company.
- Kress, G. & van Leeuwen, T. (1996). *Reading Images: the grammar of visual design*. London and New York: Routledge.
- Kress, G. & van Leeuwen, T. (1998). Front pages: the (critical) analysis of newspaper layout. In *Approaches to Media Discourse* (A. Bell & P. Garrett, eds.) (pp. 186--219). Oxford: Blackwell.
- Kress, G. & van Leeuwen, T. (2001). *Multimodal discourse: the modes and media of contemporary communication*. London : Arnold.
- Kress, G. (1993) Genre as social practice. In: Cope/Kalantzis (eds.) *The power of literacy: a genre approach to writing*, London: Falmer Press.
- Kress, G., Jewitt, C., Ogborn J. and Tsatsarelis C. (2001). *Multimodal teaching and learning*. London : Continuum.
- Labov, W. & Waletzky, J. (1978). Narrative Analysis. In *Essays on the verbal and visual arts* (J. Helm, ed.) (pp. 12--44). Seattle : University of Washington Press.
- Landow, G.P. (2001) "Reconfiguring narrative". In: *Hypertext 2.0: the convergence of contemporary critical theory and technology*, Baltimore and London: John Hopkins University Press.
- Lemke, J. (1998). Multiplying meaning: visual and verbal semiotics in scientific text. In *Reading science: critical and functional perspectives on discourses of science* (J. R. Martin & R. Veel, eds.) (pp. 87--113). London : Routledge.
- Lemke, J. L. (1999). Typology, topology, topography: genre semantics. MS. University of Michigan. URL: '<http://www-personal.umich.edu/~jaylemke/papers/Genre-topology-revised.htm>' (Last visited: 31.01.2004)
- Mann, W. C, & Thompson, S. A. (1988). Rhetorical Structure Theory: Toward a Functional Theory of Text Organization. *Text*, 8, 243--281.
- Martin, J. R. (1985). Process and text: two aspects of human semiosis. In *Systemic perspectives on discourse, Volume 1; Selected theoretical papers from the ninth International Systemic Workshop* (J. D. Benson & W. S. Greave, eds.) (pp. 248--274). Norwood, New Jersey : Ablex.
- Martin, J. R. (1997). Analysing genre: functional parameters. In *Genre and institutions: social processes in the workplace and school* (F. Christie & J. R. Martin, eds.) (pp. 3--39). London : Cassell.
- O'Halloran, K. L, (1999). Interdependence, Interaction and Metaphor in Multisemiotic Texts. *Social Semiotics*, 9, 317--354 .
- Swales, J. M. (1990). *Genre Analysis: English in academic and research settings*. Cambridge: Cambridge University Press.
- Van Leeuwen, T. & Kress, G. (1995). Critical layout analysis. *Internationale Schulbuchforschung*, 17, 25--43 .
- Waller, R. (1988). *The typographical contribution to language: towards a model of typographic genres and their underlying structures*. PhD dissertation. Department of Typography and Graphic Communication, University of Reading, Reading, U.K.