

MODALITY THEORY: SUPPORTING MULTIMODAL INTERFACE DESIGN

Niels Ole Bernsen

Centre for Cognitive Science
P.O.Box 260, Roskilde University
4000 Roskilde, Denmark
+45 4675 7711 nob@ruc.dk

ABSTRACT

Modality theory addresses the following general problem of mapping task domain information into interactive multimodal interfaces: given any particular set of information which needs to be exchanged between user and system during task performance in context, identify the input/output modalities which constitute an optimal solution to the representation and exchange of that information. This paper proposes a research agenda for modality theory and presents two steps towards its implementation. The first is a generative taxonomy of output modalities covering the media of graphics, sound and touch. The second is a methodology for carrying out information-mapping in design practice. It is argued that modality theory may provide useful support to contemporary designers of interactive human-computer interfaces who have begun to use a rapidly increasing number of different, and often alternative input/output modalities for the expression and exchange of information between systems and their users.

KEYWORDS: Multimodal systems, interface design, media and modalities, HCI methodology, modality theory.

A RESEARCH AGENDA FOR MODALITY THEORY

In recent years, human-computer interaction (HCI) has entered a new stage of development at which the outlines of the field as a mature applied science are emerging [11]. On top of the 'toolkit' of low-level usability engineering methods which are becoming integrated into design practice [20], a new layer of design support methods have appeared which serve to make the design process explicit in terms of design space structure and

development as well as in terms of the designer reasoning which operates in the design space and drives its development. Examples are the Design Space Development (DSD) framework [3, 4] and several approaches to Design Rationale [19]. However, while promising to add much needed perspicuity and explicitness to design processes, this methodological layer in itself contributes relatively little in terms of basic science. Rather, it provides a series of candidate bridging representations [1] between basic science and practical design. Arguably, the provision of more explicit structure and contents to design processes is a precondition for the systematic application of basic science to the solution of usability problems in computer artifact design. Correspondingly, the developments just described are accompanied by strong pressures on basic science to meet the real needs of design practice [9]. One such demand stems from the advent of entire families of new input/output technologies which will impact design practice in ways that science is only beginning to address.

This paper proposes a research agenda for modality theory for HCI and reports first steps in addressing the agenda in the context of the Esprit Basic Research project GRACE. Modality theory integrates research on the information representation and exchange capabilities of multimodal interfaces with the development of an information-mapping methodology which may serve as a bridge between basic science and practical design. The motivation for developing modality theory in support of usability engineering is the following [5]. Contemporary designers of interactive human-computer interfaces are beginning to use a rapidly increasing number of different, and often alternative input/output modalities for the expression and exchange of information between systems and their users. The interface designer's task can be described roughly as follows: (1) Identify the information to be exchanged between users and the artifact to be built; (2) perform a good match in terms of functionality, usability, naturalness, efficiency, etc. between that information and the available input/output modalities; (3) design, implement and test. Designers have become highly skilled at performing these steps (non-sequentially) on static graphical user interfaces (GUIs) in combination with keyboard and mouse. However, we still lack solid scientific theory that may explain and evaluate current design practices even in the area of GUI/task domain information-mapping. Interfaces increasingly incorporate spoken and written language, sound, touch and gesture in addition to new forms of graphical expression. The term 'Modality Theory' seems apt for characterising research on the corresponding, general information-mapping problem, i.e.: *Given any particular set of information which needs to be exchanged between user and system during task performance in context, identify the input/output modalities which constitute an optimal solution to the representation and exchange of that information.*

Solving the mapping problem requires investigation of the following issues:

1. To establish sound conceptual and taxonomic foundations for analysing any particular type of unimodal or multimodal output representation;
2. to establish sound foundations for analysing input modalities and entire interactive computer interfaces;
3. to develop a practical methodology for applying the results of steps (1) and (2) to the problem of information-mapping in information systems design.

Modality theory is not new in substance. Much work, both empirical and semantical in a broad sense of the term, has been done on the information representation capabilities of

selected graphical modalities, often combined with typed natural language [8, 24, 25], and terms such as 'direct manipulation interfaces' have become standard. However, solid and useful taxonomic work is still in its infancy [2, 16, 26]. Taxonomies of multimodal systems from a system engineering point of view are emerging [12, 21]. The approach which is closest to the one adopted in this paper is the seminal work of Hovy and Arens [14]. This work, however, does not address taxonomy, input modalities or the touch medium and only partly addresses the crucial question of how to establish realistic links with design practice. A shared aim, it is presumed, is to *support and constrain* designer creativity rather than to mechanically replace it.

Two contributions to modality theory are presented below. Sect. 2 proposes a generative taxonomy of output modalities. Sect. 3 describes a practical, stepwise methodology for information-mapping which may incrementally incorporate emerging results from modality theory. Sect. 4 concludes and discusses future work.

A GENERATIVE TAXONOMY OF OUTPUT MODALITIES

Attempts to address the first agenda item of modality theory face two problems. The first is that of domain complexity. Literally thousands of output modality combinations are becoming available to interface designers. Theory cannot and should not explicitly address all of these but should provide principles by which any given modality combination can be analysed when needed. This calls for a generative approach from simple elements at the right levels of abstraction. Secondly, the terminology in the field is confusing. We should aim for a terminology that is robust, conceptually clear and intuitively acceptable. If it can be agreed that, e.g., tables, beeps, written and spoken natural language may all be termed 'modalities' (cf. [14]), then the intuitive acceptability of the approach presented below would seem ensured. Modalities in this sense are *representational modalities* and clearly distinguishable from the 'sensory modalities' of psychology.

To gracefully tackle the complexity problem, the generative taxonomy is hierarchical and has two levels, a basic *generic level* and an *atomic type level*. The generic level ensures that the taxonomy is based on a limited set of generic modalities from which any given output modality or modality combination can be generated and analysed. Each generic modality has a number of actual or possible *modality types* subsumed under it which inherit its basic properties and have discriminatory properties of their own. Each generic modality is pure or *uni-modal* and hence elementary as compared with combined or multi-modal representations. Each pure generic modality is characterised by a small set of *basic features* which serve to robustly distinguish modalities from one another within the taxonomy. Each of these features are assumed to have profound implications for a certain modality's capacity for representing information. The features are: *linguistic/non-linguistic*, *analogue/non-analogue*, *arbitrary/non-arbitrary*, *static/dynamic*. In addition, we distinguish between the *media of expression* of graphics, sound and touch which are each characterised by very different sets of perceptual qualities (visual, auditory and tactile, respectively). These media determine the scope of the taxonomy. A pure generic representational modality is thus a complex-property entity characterised by a specific medium of expression and a *profile* constituted by its basic features. For instance, the same linguistic information may be represented in either the graphical, sound or touch medium but the choice of medium strongly influences the suitability of the representation for a given design purpose and is therefore considered a choice between different modalities.

15. Real sound		x	x			x		x		x	
16. Diagrammatic sound											
17. Sound graphs											
18. Real touch		x	x			x		x			x
19. Diagrammatic touch											
20. Touch graphs											
21. Arbitrary diagrams		x		x	x		x		x		
22. Animated arbitrary diagrams		x		x	x			x	x		
23. Arbitrary sound		x		x	x			x		x	
24. Arbitrary touch		x		x	x			x			x
25. Static graphics structures		x		x		x	x		x		
26. Dynamic graphics structures		x		x		x		x	x		
27. Sound structures		x		x		x		x		x	
28. Touch structures		x		x		x		x			x
modality	li	-li	an	-an	ar	-ar	sta	dyn	gra	sou	tou

Table 1. A taxonomy of pure generic modalities. Except for the rows containing the modalities 9-20, each row exclusively represents one single pure generic modality. Four classes of modalities are separated by boldface lines: linguistic, analogue, arbitrary and explicit structures. The table itself is a multimodal combination of modalities 5 and 25.

dimensionality) whereas real-world representations do this to a lesser extent. Given current manipulation possibilities, this distinction (between 9/10, 12/13, 15/16 and 18/19, respectively) seems to have to be prototype-based. Thus, a photograph is a prototypical real-world representation whereas a drawn sketch is a prototypical diagram. The second distinction is between diagrammatic and real-world representations, on the one hand, and graphs (11, 14, 17 and 20, respectively) on the other. Graphs manipulate the representation in specific ways (see below). Creating the taxonomy of Table 1 from permutations on basic properties has been a generative exercise for the author who never before thought about dynamic hieroglyphs (2), animated arbitrary diagrams (22) or explicit sound structures (27), the two latter of which turn out to be in actual interface use.

Even an intuitively familiar, exclusive and exhaustive taxonomy of modalities is of limited value unless accompanied by relevant analyses of the basic features whose presence or absence in a given modality strongly influence its capacity for representing information. The representational implications of the linguistic/analogue distinction have been described in [6]. The static/dynamic distinction and the representational implications of using different media need more work. Summarising, *linguistic* representations can, somehow, represent anything. However, linguistic representations lack the *specificity* which is basic to analogue representations [23]. Linguistic representations are *focused*: they focus on the subject-matter to be communicated without providing its specifics. My neighbour, for instance, is a specific person but you won't know much about his specifics from

understanding the expression 'my neighbour'. The presence of focus and lack of specificity jointly generate the characteristic limited expressive power of linguistic representations, whether these be static or dynamic, graphical, auditory or tactile, or whether the linguistic signs used are themselves non-analogue (as in the present text) or analogue. Complementarily, *analogue* representations (also called 'iconic' or 'isomorphic' representations) have the virtue of specificity but lack focus, whether they be static or dynamic, graphical, auditory or tactile. Specificity and lack of focus generate the characteristic limited expressive power of analogue representations. The complementarity noted explains why (multimodal) combinations of linguistic and analogue representations are eminently suited to many representational purposes. Thus, one basic use of language is to *annotate* analogue representations (e.g., a map, a diagram or a dynamic measurement representation), and one basic use of analogue representation is to *illustrate* linguistic discourse [6]. The specificity of analogue representation is related to the fact that analogue representations have 'shape' or dimensionality, i.e. are encoded relative to a system of dimensions such as, e.g., 2-D space [13, 22]. *Graphs* constitute a particular genus of analogue representation in that they represent data in a *graph space* according to one or more dimensions of interest [2]. In graphs, in contrast to real-world representations and diagrams, any 'pictorial' similarity to the represented subject-matter has disappeared but since dimensionality is still represented, graphs remain analogue representations.

Arbitrary representations are selected by designers and others to represent something without relying on an already existing system of meaning whereas *non-arbitrary* representations rely on an already existing system of meaning (cf. Table 1). Arbitrary representations imply an extra cognitive load on recipients who must learn the new representational conventions. The *dynamic/static* distinction depends on whether the temporal dimension is explicitly part of the representation or not. *Explicit structures* (modalities 25 to 28) are used to explicitly mark distinctions and separations among representations. Finally, the different representational properties of *media* depends on the properties of their *information channels* [2, 14].

The central claim embodied in the taxonomy is that it has strong generative power and may predict the information representation capabilities of any type of unimodal or multimodal representation in the media of graphics, sound and touch. However, even a generic taxonomy whose basic features have been spelled out and exemplified much more

Modality	Well-known types
1. Static analogue written language	Hieroglyphs. Rarely used.
2. Dynamic analogue written language	Apparently none. Dynamic hieroglyphs would appear anachronistic.
3. Analogue spoken language	Part of everyday spoken language.
4. Analogue touch language	Apparently none.
5. Written language	Written letters, words, numerals, other written language related signs, text, programming languages, formal logic, logograms such as arrows, musical notation, list and table orderings.
6. Dynamic written language	Moving text, running numerical counters, digital clocks.
7. Spoken language	Spoken letters, words, numerals, other spoken language related sounds, discourse, list orderings.

8. Touch language	Touch letters, numerals, words, other touch language related signs, text, list and table orderings. Example: Braille.
9. Diagrammatic pictures	Pure diagrams, maps, cartoons, sequential, list and table orderings. 1D, 2D or 3D spatial.
10. Non-diagrammatic real-world pictures	Pure photographs, naturalistic drawings, holograms, sequential, list and table orderings.
11. Static graphs	1D, 2D or 3D graph space with geometrical forms. Pure charts (dot charts, bar charts, pie charts, etc.).
12. Animated diagrammatic pictures	Pure animated diagrams, sequential, list and table orderings. Pure standard animations.
13. Dynamic real-world pictures	Pure movies, videos, realistic animations. Sequential, list and table orderings possible.
14. Dynamic graphs	Pure graphs (see 11) evolving in graph space. Sequential, list and table orderings possible.
15. Real-world sound	Single sounds, sound sequences. List ordering possible.
16. Diagrammatic sound	Apparently none, but many possibilities, synthetic or manipulated, exist. Music?
17. Sound graphs	E.g. Geiger counters.
18. Real-world touch	Single touch representations, touch sequences.
19. Diagrammatic touch	Apparently none, but many possibilities exist.
20. Touch graphs	1D, 2D or 3D graph space with geometrical forms. Pure charts (dot charts, bar charts, pie charts, etc.).
21. Arbitrary diagrammatic forms	Pure points, lines, boxes, circles, volumes, etc., sequential, list and table orderings..
22. Animated arbitrary diagrammatic forms	Pure points, lines, boxes, circles, volumes, etc., sequences of such.
23. Arbitrary sound	Single sounds, sound sequences.
24. Arbitrary touch	Touch signals of different sorts.
25. Static graphics structures	Form fields, frames, table grids, line separations, trees, windows, bars.
26. Dynamic graphics structures	Dynamic frames, windows, scroll bars.
27. Sound structures	Apparently none.
28. Touch structures	Form fields, frames, grids, line separations, trees.

Table 2. Well-known types (if any) of each of the pure generic modalities.

than was possible above, provides an incomplete analysis of the features of output modalities which are relevant to information-mapping in HCI. To complete the analysis and obtain a practical tool, it is necessary to move to the level of *modality types* which, in addition to their inherited basic features, have important properties of their own.

Table 2 represents some well-known types, if any, of the pure generic modalities. These types are, of course, equally uni-modal. To complete work on research agenda item 1 of modality theory we need a more principled inventory of uni-modal types than that of Table 2, each characterised through examples and a set of key representational properties

including its inherited basic features. Developing this inventory as a functional tool for information-mapping in interface design is the subject of ongoing work [18]. In Table 2 some of the well-known types are described as 'pure', e.g., as pure diagrams or pure graphs. This follows from the generative nature of the taxonomy. However, without linguistic annotation, such pure types are of limited use. To increase the usability of the taxonomy, therefore, it seems desirable to develop such types into *minimal* (multimodal) *types* such as standard annotated diagrams or graphs.

AN INFORMATION-MAPPING METHODOLOGY

Agenda items 1 and 2 represent the target scientific foundations of modality theory. What is needed in addition is a practical way of bridging between basic science and interface design practice. The bridge has to carry two-way traffic because the basic science contribution to interface design will probably need revisions due to experience gathered in actually using it to support information-mapping design decisions. Even at this early stage in the development of modality theory, it therefore makes good sense to develop and start applying a methodology for mapping information from task domains into interactive interfaces. The proposed methodology proceeds in five steps [7].

Step 1: Identification of Information and Tasks

The first problem is to identify the information to be exchanged by user and system during task performance in the application domain of the artifact to be designed. So the aim of Step 1 is to obtain the information from the task domain which is needed to select a reasonable and possibly optimal mapping from task domain information to interface input/output representation. The variety of information relevant to this end should not be underestimated (see below). Standard usability engineering methods may be used in gathering this information as part of the requirements specification process. Often, but not always, a central part of the information needed to solve an information-mapping problem is information on users' tasks. However, any reasonably versatile IT artifact can be used for performing a multitude of different tasks and it is obviously not possible during systems design to consider in detail each and every such task as to its information-mapping requirements. In other words, it is necessary during practical interface design to be selective as to the tasks to be analysed in detail. The ideal way to be selective is to identify a limited set of *task scenarios* which are representative of the intended artifact and then carry out the information-mapping analysis on these. The problem is that no guaranteed method for generating an appropriate set of scenarios currently exists in HCI. One proposed heuristic [10] is too weak for this purpose and we are currently testing an alternative method [15]. Let us just assume that the best current methods or heuristics are being applied in identifying representative tasks.

The results of Step 1 would normally be (a) high-level task domain information relevant to the information-mapping problem and (b) a small set of representative tasks which users should be able to carry out on or with the intended artifact. These results constitute an operationalisation of the information-mapping problem. Step 1 is crucial to the success of the methodology as failure to complete it properly means that important information requirements on the artifact have been overlooked.

Step 2: Selective Task Analysis

In Step 2 the representative tasks are analysed in as much detail as possible in order to identify their goals and initial states, the activities and procedures involved, how they

might go wrong, the task (work) environment, the intended users and their experience, etc. The analysis should primarily aim at revealing the input/output information representation and exchange needs of the tasks. That is, while a more or less complete task analysis may be done either formally or informally, not all of the information it produces needs to be explicitly represented in order that the information-mapping methodology will succeed.

Step 3: Information Representation

In Step 3 the relevant information acquired through Steps 1 and 2 is represented explicitly and succinctly, for instance using the Design Space Development (DSD) notation for representing design space structure [3, 4]. In principle, the representation should contain everything which is relevant to the input/output modality choices to be made. The representation should be expressed in terms of modality theory. Step 3 makes explicit the requirements on interactive information to be satisfied by the interface to be designed and concludes the first main phase of the methodology.

Step 4: Information-Mapping

Step 4 consists in applying the theoretically developed framework for representing the elementary and generated components of interactive unimodal or multimodal interfaces, i.e., the results of research on agenda items 1 and 2 of modality theory. The framework should eventually contain the elements needed for generating and analysing any specific type of unimodal or multimodal input or output including the effects of combining these into complex interfaces. A mapping is performed of the results of Step 3 into the elements of modality theory. The result will be sets of candidate input/output modalities and modality combinations capable of representing and exchanging the information needed for the representative tasks in context. It is likely that the mapping will often produce several alternative solutions which subsequently have to be compared and traded off against one another.

Step 5: Trade-Offs

In Step 5 a 'higher level filtering' is performed to trade off potential solutions against one another given the results of Steps 1 through 4. The trade-off process may be explicitly represented in some form of Design Rationale representation (e.g. [17]). Step 5 produces a solution to the task domain/interface mapping problem together with its Design Rationale. In some cases, several solutions can be expected to emerge from the trade-off process with identical potential for solving the interface design problem at hand.

Case Studies

We have done two case studies in applying the information-mapping methodology to in-house design projects, i.e., a spoken language dialogue system and a 'water bath' toy control room application [7]. Given the state of progress of our work on modality theory as described above, these case studies of course do not amount to proper testing of the science base. They did, however, provide useful observations which should be taken into account in further developments of the theory. Firstly, *information-mapping problems arise at very different levels of generality during interface design*, from the very first design commitments to be made through to decisions on minute interface details. It follows that the science base should be developed to support information-mapping at different stages of requirements-capture and levels of generality. Secondly, *task domain information relevant to information-mapping derives from many different sources* and not

just from task analysis. In the spoken language dialogue systems case, for example, the fact that Danes do not currently have access to electronic GUI-systems networks such as the French Minitel was important to justifying the selected application domain (flight reservation and information). In the 'water bath' case, the level of control room noise provided a critical parameter for information-mapping. Thirdly, although modality theory adopts a primarily semantic or information-representational approach to the problem of modality choice, this does not mean that cognitive aspects of interface use are judged irrelevant. On the contrary, *questions to the cognitive science base of HCI may arise at any point during the application of the information-mapping methodology* and in particular during the steps involving information-mapping and trade-offs between different candidate solutions.

CONCLUSION AND ONGOING WORK

Modality theory exemplifies the attempt to develop the science base of HCI from the needs of design practice. This paper has presented a research agenda for modality theory and two first steps towards its implementation. The generative taxonomy of output representations provides hierarchical structure and some powerful basic properties to the multitude of actual and potential interface modalities and should be able to incorporate many existing achievements in modality analysis without distorting them. The methodology for information-mapping proposes how to apply the results of modality theory as they appear. It is too early to predict the potential impact of modality theory on design practice. The impact will certainly be less in areas, such as GUI interface design, where designers' craft skills are already well-developed, than in the emerging areas of multimodal and virtual reality systems design. Our current work aims at completing the semantic structuring, identification, property analysis and exemplification of elementary modality types; analysing different media of expression in terms of information channels; and constantly testing and revising results through case studies. Agenda item 2 of modality theory will be next. We would be grateful to receiving contacts with others who are engaged in the broad enterprise sketched in this paper.

ACKNOWLEDGEMENTS

The work described in this paper was carried out under Esprit Basic Research project 6296 GRACE whose support is gratefully acknowledged. Annick Bertels did most of the work on the 'water bath' system. Thanks are due to Michael May, Kenneth Holmquist and John Lee for comments and criticisms.

REFERENCES

1. Barnard, P. Bridging between Basic Theories and the Artifacts of Human-Computer Interaction. In Carroll, J.M., Ed. *Designing Interaction. Psychology at the Human-Computer Interface*. Cambridge, Cambridge University Press, 1991.
2. Bernsen, N.O. Matching Information and Interface Modalities. An Example Study. *Working Papers in Cognitive Science WPCS-92-1*, Centre of Cognitive Science, Roskilde University, 1992.
3. Bernsen, N.O. Structuring Design Spaces. *InterCHI '93 Adjunct Proceedings*, 211-12.

4. Bernsen, N.O. The Structure of the Design Space. In Byerley, P.F., Barnard, P.J. and May, J., Eds. *Computers, Communication and Usability: Design Issues, Research and Methods for Integrated Services*. Amsterdam, North-Holland, 1993, 221-244.
5. Bernsen, N.O. A Research Agenda for Modality Theory. In Cox, R., Petre, M., Brna, P. and Lee, J., Eds. *Proceedings of the Workshop on Graphical Representations, Reasoning and Communication*. World Conference on Artificial Intelligence in Education, Edinburgh, August 1993, 43-46.
6. Bernsen, N.O. Specificity and Focus. Two Complementary Aspects of Analogue Graphics and Natural Language. *Esprit Basic Research project GRACE Deliverable 2.1.3*, 1993.
7. Bernsen, N.O. and Bertels, A. A Methodology for Mapping Information from Task Domains to Interactive Modalities. *Esprit Basic Research project GRACE Deliverable 10.1.3*, 1993.
8. Bertin, J. *Semiology of Graphics. Diagrams. Networks. Maps*. Madison, University of Wisconsin Press, 1983.
9. Carroll, J.M., Ed. *Designing Interaction. Psychology at the Human-Computer Interface*. Cambridge, Cambridge University Press, 1991.
10. Carroll, J. M., Kellogg, W. A. and Rosson, M. B. Getting Around the Task - Artifact Cycle: How to Make Claims and Design by Scenario. *ACM Transactions on Information Systems*, 10, 2, 1992, 181-212.
11. Carroll, J.M. Creating a Design Science of Human-Computer Interaction. *Interacting with Computers* 5, 1, 1993, 3-12.
12. Coutaz, J., Nigay, L. and Salber, D. The MSM Framework Design Space for Multi-Sensory-Motor Systems. To appear in Moran, T.P. and Carroll, J.M., Eds. *Design Rationale*. Hillsdale, NJ, Erlbaum. In press.
13. Haugeland, J. Representational Genera. In Ramsey, W., Stich, S.P. and Rumelhart, D.E., Eds. *Philosophy and Connectionist Theory*. Hillsdale, NJ, Erlbaum, 1991.
14. Hovy, E. and Arens, Y. When is a Picture Worth a Thousand Words? Allocation of Modalities in Multimedia Communication. Paper presented at the *AAAI Symposium on Human-Computer Interfaces*, Stanford 1990.
15. Klausen, T. and Bernsen, N.O. CO-SITUE: Towards a Methodology for Constructing Scenarios. In Hollnagel, E. and Lind, M., Eds. *Proceedings of the Fourth European Meeting on Cognitive Science Approaches to Process Control: Designing for Simplicity*. Copenhagen, August 1993, 1-16.
16. Lohse, G., Walker, N., Biolsi, K. and Rueter, H. Classifying Graphical Information. *Behaviour and Information Technology* 10, 5, 1991, 419-36.

17. MacLean, A., Bellotti, V. and Shum, S. Developing the Design Space. In Byerley, P.F., Barnard, P.J. and May, J., Eds. *Computers, Communication and Usability: Design Issues, Research and Methods for Integrated Services*. Amsterdam, North Holland, 1993, 197-220.
18. May, M. A Taxonomy of Representations for HCI, Parts 1-3. *Esprit Basic Research project GRACE Deliverable 2.1.5, 2.1.6 and 2.1.7*, 1993.
19. Moran, T.P. and Carroll, J.M., Eds. *Design Rationale*. Hillsdale, NJ, Erlbaum. In press.
20. Nielsen, J. *Usability Engineering*. New York, Academic Press, 1993.
21. Nigay, L. and Coutaz, J. A Design Space for Multimodal Systems: Concurrent Processing and Data Fusion. *Proceedings of InterCHI '93*, 172-78.
22. Slack, J., Fedder, L. and Oberlander, J. Animation: A Representation Framework. *Esprit Basic Research project GRACE Deliverable 4.1.1*, 1993.
23. Stenning, K. and Oberlander, J. Reasoning with Words, Pictures and Calculi: Computation versus Justification. In Barwise, J., Gawron, J.M., Plotkin, G. and Tutiya, S., Eds. *Situation Theory and Its Applications*. Stanford, CA, CSLI, 1991, Vol. 2, 607-621.
24. Tufte, E.R. *The Visual Display of Quantitative Information*. Cheshire, CT, Graphics Press, 1983.
25. Tufte, E. *Envisioning Information*. Cheshire, CT, Graphics Press, 1990.
26. Twyman, M. A Schema for the Study of Graphic Language. In Kolars, P., Wrolstad, M. and Bouna, H., Eds. *Processing of Visual Language* Vol. 1. New York, Plenum Press, 1979.