Where No Mind Has Gone Before: Ontological Design for Virtual Spaces

Nancy Kaplan
Stuart Moulthrop
Institute for Publications Design
The University of Baltimore
1420 North Charles Street
Baltimore, Maryland U.S.A. 21201-5779
nakaplan@ubmail.ubalt.edu
samoulthrop@ubmail.ubalt.edu

ABSTRACT

Hypermedia designers have tried to move beyond the directed graph concept, which defines hypermedia structures as aggregations of nodes and links. A substantial body of work attempts to describe hypertexts in terms of extended or global spaces. According to this approach, nodes and links acquire meaning in relation to the space in which they are deployed. Some theory of space thus becomes essential for any advance in hypermedia design; but the type of space implied by electronic information systems, from hyperdocuments to "consensual hallucinations," requires careful analysis. Familiar metaphors drawn from physics, architecture, and everyday experience have only limited descriptive or explanatory value for this type of space. As theorists of virtual reality point out, new information systems demand an internal rather than an external perspective. This shift demands a more sophisticated approach to hypermedia space, one that accounts both for stable design properties (*architectonic space*) and for unforeseen outcomes, or what Winograd and Flores call "breakdowns." Following Wexelblat in cyberspace theory and Dillon, McKnight, and Richardson in hypermedia theory, we call the domain of these outcomes *semantic space*. In two thought experiments, or brief exercises in interface design, we attempt to reconcile these divergent notions of space within the conceptual system of hypermedia.

KEYWORDS: Spatial hypertext, interface design, information mapping, navigation

1 INTRODUCTION: NODE, LINK, SPACE

Hypermedia designers and theorists have tried repeatedly to move beyond the directed graph concept, which defines hypermedia structures as aggregations of nodes (containers of information) and links (dynamic connectors). Halasz [11] folds these primitives into node/link "composites." DeRose [7] distributes link functions over a taxonomy of "intensional" and "extensional" links, while Nanard and Nanard [23] envision a type scheme for link anchors as well. Parunak [27] replaces the concept of linked nodes with set-theoretic groupings, while Marshall et al. [20] propose a frame-based model. Stotts, Furuta, and Ruiz [32] present a document-browsing automaton that operates on link structures alone. Theorists interested in the graphic dimension of hypermedia structure, such as Lai and Manber [15], Gloor [10], and Noik [24], propose formalisms for surveying and analyzing spatial relationships within a text.

These conceptions are explored most profitably in Marshall and Shipman's work [21] on the expression of implicit structures through spatial arrangement. In their approach, nodes and links acquire meaning in relation to the space in which they are deployed. Accordingly, some theory of space seems essential for any advance in hypermedia design; but the type of space implied by electronic information systems, from hyperdocuments to the "consensual hallucinations" of virtual reality, requires careful analysis. Bolter [4] suggests that hypertext creates a new "writing space"; writing about space in virtual reality systems, Wexelblat [34] invokes the term "semantic space." The nature of this space resists easy definition. Familiar metaphors from physics, architecture, and everyday experience have only limited value here. As Dillon et al. [8] observe, the psycholinguistic or semantic space of a text (electronic or otherwise) can never be represented with perfect accuracy by any physical system. A radically subjective element necessarily comes into play. Novak [26] points out that contemporary information systems demand an internal rather than an external perspective. Virtual space is an interactive continuum defined by our committed participation as well as our simple attention. In the space of electronic media, we are always more than distanced observers.

Proceeding from this starting point, we offer a critical review of two issues in hypermedia design: transition ("navigation") and information mapping. Our approach to these issues is guided by Winograd and Flores's notion [33] of "ontological design" — a design philosophy that addresses human-machine interactions in terms of complex environments instead of simple, end-directed functionality. In response to Winograd and Flores's demand that new design be based on conceptual complication or "breakdown," we offer two thought experiments that might help redefine our understanding of navigation and mapping.

2 ARCHITECTONIC AND SEMANTIC SPACES

Marshall and Shipman's work on VNS, NoteCards, and Aquanet represents a crucial advance in thinking about the spatiality of hyperdocuments. In place of "explicit" links they concentrate on the arrangement of nodes in a graphic representation, permitting their system to address an "implicit" level of structure. This approach escapes the constraints of the traditional directed-graph model without sacrificing the visual dimension of hypermedia. However, Marshall and Shipman's conception reflects only one aspect of the complex phenomenology of virtual space. In Marshall and Shipman's work, as in other major systems such as Writing Environment, SEPIA, and Storyspace, the general idea of *space* tends to collapse into the much narrower domain of *screen real estate*. The user's manipulation of objects within a graphic representation implies some related transformation in a mental or linguistic space, but that space is accessible only through the representation. "Space" comes to be defined in terms of the active window on a display screen.

Graphic display is of course indispensable to the concept of virtual space, since connecting directly to the visual cortex remains for the moment the stuff of science fiction. However, screen real estate is not the only sort of space involved in hypermedia design. The space of the screen is defined in terms of pixels and other coordinate systems. It represents *architectonic* space. This is a regular, mathematically precise space in which absolute principles apply: objects are uniquely identified, have exclusive locations, and obey rules of geometry and perspective. Most important, the architectonic space of screen display is binary. Absence and presence, or the relation of figure to ground, are immediately and unambiguously apparent. A pixel is either devoted to a graphic element or it is not. Elements may overlap or include each other, but in doing so they complicate the visual field, introducing an appearance of depth. Space on the screen imitates space in the physical world, where architecture involves manipulating stable objects according to regular principles. Along with this space of geometries and built structures, hypermedia also requires us to operate in a very different environment. We might call this *semantic* space, since it is deeply connected to the production of meaning, interpretation, and other activities involving symbols. Marshall and Shipman invoke architectonic space in the context of writing; semantic space emerges more clearly in the act of reading or reception — though since hypertexts tend to blur the roles of reader and writer, these distinctions cannot be absolute. To understand semantic space, we need to follow Novak's prescription and imagine ourselves within an information environment: for instance, a hypertext. Hypermedia researchers have expressed considerable concern that higher levels of complexity in interactive documents might impair effective communication (see Raskin [29], Lesk [19], Carlson [5], Wright [35], Charney [6], and for a critical review, Dillon et al. [8]). These reservations are justified. Hypermedia texts are undeniably more complicated and cognitively challenging than conventional texts. As some studies have indicated (Egan et al. [9], Jones and Spiro [14]), this complexity may have positive value for interpretive tasks requiring multiple frames of reference. But encountering these multiple frames, or heterogeneous streams of language, always invites confusion. The conceptual or semantic space of a hyperdocument is inevitably more complicated than that of a comparable linear text.

The inherent complexity of hypermedia structures becomes most apparent in navigation, or selecting and activating links. Even in a highly structured, task-specific application such as technical documentation or argumentation, following a link brings ambiguity and uncertainty into the reading process. If the author has given adequate attention to Landow's rhetoric of arrivals and departures [16], or what Dillon et al. call "landmarks" [8], then we should understand clearly enough where we are going and why. In terms of local linearity, a hypermedia text need be no more challenging than a print document. But this does not neutralize the complexity of hypermedia reception. Even as we follow the locally linear track we have selected, we become aware that we could have chosen otherwise. Every link we follow must be accompanied by a number of alternatives. Some of these may have been explicitly signaled by the system, as in NCSA Mosaic, where link anchors are underlined and color-coded to show whether we have already visited their destinations. Other alternatives, which Moulthrop [22] has called "implicit" destinations, are present to the reader as speculations or mental projections. In this regard, hyperdocuments invite what Hofstadter [13] calls "jumps outside the system."

The concept of space used in this account of navigation shares little of the clarity and unambiguousness of architectonic space. As Dillon et al. [8] see it, "we cannot navigate semantic space, at least not the way we navigate physical environments, we can only navigate the physical instantiations that we develop of the semantic space" (p 187). Harpold [12] discusses hypertextual linking as "detour," not a definitive trajectory from departure point to arrival point, but an elliptical and fundamentally uncertain displacement. The hypertextual detour, he says, "is a turn around a place you never get to, where something drops away between the multiple paths you might follow" (p 173). Perhaps that which "drops away" as we traverse a hypermedia link is indeed our orientation in architectonic space, with its stable geometries and singularity of expression. In the directed graph we can see clearly where we have come from and where we are going; but this is not the case in semantic space. For every point of actual arrival in a hypermedia text, there are an unspecified number of places we never get to, alternative destinations which the system has either disclosed to us, and which we have chosen not to visit, or which may be simply undeveloped or unexpressed in the current version of the text. Semantic spaces are n-dimensional, as Dillon et al. [8] point out, while architectonic spaces have at best three dimensions, and more usually two. Architectonic space is always either empty or filled (see Benedikt [1]). In the absence of an object, the space is empty — or in any event, that is what the assignment of a null or neutral value to a pixel is taken to represent. In semantic space, however, the default condition is not *definitely empty* but rather *indefinitely filled*. A semantic space is a domain of possible expression. It is "semantic" because it is the place where meanings or interpretations come into existence; and in hypertext, as the cognitive critics cited earlier insist, there is always a surplus of meaning.

3 "ONTOLOGICAL DESIGN"

These speculations appear to move hypermedia design a bit out of its usual conceptual range. Figures like Harpold (along with Bolter, Landow, and the authors of this paper) come to hypermedia through Landow's controversial "convergence" between information technology and literary theory [17]. Though all of us also develop hypermedia systems and documents, we became interested in hypermedia partly because of its similarity to certain notions associated with the poststructuralist critique of print and its cultural institutions. Harpold's concept of "detour," for instance, draws upon ideas developed by the psychologist Lacan and the philosopher Derrida. These concepts may have more apparent bearing on philosophy and literature than on the design of information systems; but appearances can be deceiving. It is possible to implement an architectonic space by means of graphic representations, but as Winograd and Flores observe [33], semantic spaces are notoriously hard to capture in software. Yet one kind of space inevitably implies the other. In analyzing strategic behavior, we may speak of "decision spaces," implying the mathematical formalism (or architectonic space) of a specific formula. As Winograd and Flores point out, these spaces never encompass the full range of options available to an agent in the real world. Architectonic constructs like decision spaces and hypertext webs inevitably imply other, less formally tractable structures in semantic space.

In Winograd and Flores's theory of design, this coupling of architectonic and semantic spaces manifests itself as "breakdown," or the moment when the constraints of a particular formal system become apparent by juxtaposition with their alternatives. As a defining instance of breakdown, they cite ELIZA's infamous response to an exasperated subject who announces, "I am swallowing poison." The program replies, "How long have you been swallowing poison?" (p 121). Assuming we are not dealing with a Hannibal Lecter, we would not expect any human psychiatrist to give such a response. ELIZA's imitation breaks down at this point because we are confronted with a case that falls outside its linguistic parameters. We have dropped out of architectonic space (the space of rule-bound knowledge representation) into the more chaotic space of language and cultural assumptions.

Winograd and Flores base their "new foundation for design" on the proposition that advanced information systems will always be susceptible to such awkward transitions — much in the same way that hypermedia texts must inevitably carry heavier cognitive overhead than print or video. Given this limitation, designers might be better off abandoning a hopeless struggle; or in other words, since we can't get rid of the bug, why not call it a feature? We can accomplish this conversion, Winograd and Flores suggest, by embracing a new concept of "ontological" design. "New design can be created and implemented," they note, "only in the space that emerges in the recurrent structure of breakdown. A design constitutes an interpretation of breakdown and a committed attempt to anticipate future breakdowns" (p 78). It is this latter element, committed anticipation, which most often seems missing from software design — even in the field of interactive media, where it might arguably be prominent. Much attention has been given to minimizing complexity and cognitive demands in hypermedia (see especially Charney [6]). These critiques envision better forms of architectonic space, whether in the form of clearer graphical representations or more coherent rhetorics of arrival and departure. Such efforts are undeniably important — architectonic space is the only kind of space we can directly manipulate, after all — but from the perspective of ontological design, they are at least potentially misguided. Any hypermedia document is extended simultaneously in both architectonic and semantic space. It occupies both a domain of function (the reliable, connective operation of the link) and one of breakdown (the susceptibility of the link to "detour" or ambiguity). Accordingly, we might turn our efforts toward designing structures that integrate or mediate between these two varieties of space.

4 DESIGN EXERCISES

We will need to re-think our conception of space in hypermedia, and by extension, the dominant metaphor of "navigation" that we use to describe transactions within it. Such re-thinking can be fully realized only in viable system designs and texts; but it might begin with a few conceptual experiments, two of which we offer here. Our general aim in these experiments is to describe architectonic structures which, though still engaged in precise graphical mapping, are better adapted to the multiplicity of semantic space.

Design Exercise 1: Warp Coefficients in Hypertext

This exercise was inspired by a case of technical breakdown. Several years ago, one of the present authors produced a fairly large literary hypertext in Storyspace. The Storyspace stand-alone reader in its then current version required that both the code for the reader and the entire hypertext, including all nodes and links, be loaded into RAM at launch. The combined textual material and program amounted to about 1.2 megabytes, a memory demand large enough to put the product beyond the reach of consumers with older, less capacious computers. In response to the author's request, Bolter designed an experimental Storyspace reader which loaded only part of the hypertext into RAM. This reader module maintained a 500-kilobyte memory store, into and out of which it would rotate sections of the text as necessary. The author tried to make these sections as large and as locally coherent as possible in order to minimize memory shuffling.

Though the new reader proved reliable, constraints of hardware and software design prevented it from working effectively as a delivery system for narrative hypertext. Bolter and the author had anticipated that users would usually follow links to material within local sections, where "local" referred to proximity in the architectonic graph of the Storyspace document. Early informal tests strongly refuted this assumption. About one time in four (roughly speaking) users chose links which carried them to distant places in the document, "distant" again defined in terms of the graph. These non-local excursions required disk access and memory operations, which even on relatively swift machines added noticeably to the transition time between narrative places. These time lags made users aware of a difference between "local" places, which could be summoned up instantly, and "distant" places encoded on some remote sector of the drive platter.

The author of the hypertext in question found these effects undesirable, so development of the new reader module was abandoned. It is important to ask why the author made this decision. What made him assume that all transitions in a hyperdocument should be of equal and apparently null duration? The "breakdown" in this case involved the author's assumption that architectonic space (the space of the Storyspace graph) could adequately model semantic space (the narrative dimensions of the story). Like all breakdowns, this one is instructive. The author insisted on instantaneous transitions because in his view, semantic space has no fixed dimensions. Part of his narrative project was to create interconnections between scattered story elements, connections which would overcome the sequential separation imposed by the printed page. To speak not entirely fancifully, his design agenda involved creating a story that moved *faster than the speed of books.*

This last formulation suggests a useful though admittedly far-fetched analogy between the semantic space of hypermedia and the interstellar spaces represented in science fiction. To imagine journeys beyond the solar system, science fiction writers have to invent methods of faster-than-light travel. These inventions often involve discontinuities in the space-time continuum: wormholes, trans-dimensional portals, or the solution most familiar to addicts of American TV, the "warp field effect," in which vessels bend or fold space around themselves. Fortunately, the physics of such inventions need not be considered here. (Someone please explain how the *Star Trek* universe avoids catastrophic collapse with thousands of ships constantly stretching space-time all over the place. For further discussion of *Star Trek*'s implausible spaces, see Benedikt [1].) We are concerned here only with the analogy between Roddenberry's famous "final frontier" and another imaginary space, the semantic domain in hypermedia. Both are in a crucial sense elastic — spaces in which dimensional properties like distance and contiguity can be easily annulled.

If we follow this analogy, however, we must reconceive the familiar concept of navigation in virtual space. Literally speaking, "navigation" refers to the operation of floating vessels; but the spacecraft of *Star Trek* and other science fiction fantasies are not much like ships. Since they move by "distort[ing] the space-time continuum" [26], starships are not propelled through a medium as are aircraft or boats. Instead they warp or wrap space around themselves. If we think of transitions in hypermedia as somehow analogous to this effect, then we need a new conceptual vocabulary — and most important, new representational schemes. The object of this design exercise is to create an architectonic representation for "warp-effect" transitions in semantic space: shifts in textual location which we understand not in terms of motion-toward (propulsion) but of gathering-in (warping).

Figure 1 illustrates a hypothetical variant of the Storyspace hypertext environment. On the tool palette, the standard navigation tool has been replaced by a warp tool. This tool shows all link anchors in the current place with boxed outlines or "wire wraps." It also displays for each anchor an overlaid field containing an integer value from 1 to 9, shown here in Roman numerals. These numbers may be thought of as warp coefficients. They represent, in a general sense at least, the amount of gathering or distortion that must be applied to the semantic space in order to effect an instantaneous transition from the current place to the destination point of the link anchor. In the illustration, material about "coupling and decoupling" lies relatively close to the current place and may be reached with a warp factor of 2. The explanation of "apparent mass reduction," on the other hand, exists in some remote textual quadrant and requires a warp-8 transaction to bring it in view.

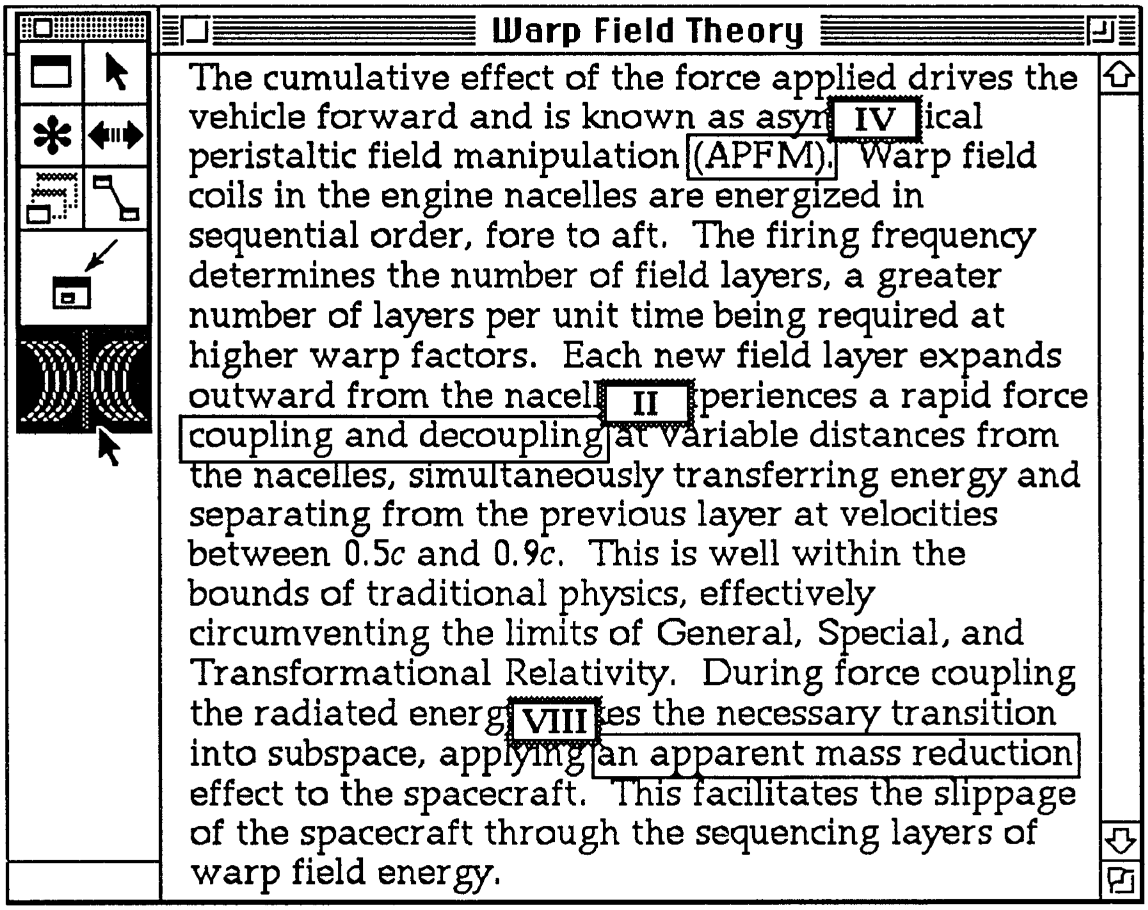

Figure 1: Storyspace, The Final Frontier (with apologies to Bolter, Joyce, and Smith; text from Sternbach and Okuda).

If fantasy physics are not to your taste, it is possible to regard this proposal simply as a weighting scheme for hypertext links, similar to that incorporated in Schneiderman's HyperTIES system [30]. But though not especially original, this design solution nonetheless raises two important questions: how are the weightings generated, and what do they represent? Since this scheme functions in architectonic space, we might generate the "warp coefficient" by referring to the structure graph, as in Marshall and Shipman's work [21] or Bernstein's "Link Apprentice" [2]. That is, the warp number might simply represent the amount of screen space, or the number of nodes and links, that intervene between one place and another. One might object, then, that this scheme emphasizes architectonic certainty over the n-dimensionality of semantic space.

To some extent this must be true; the scheme of weighted links must break down at some point, such as in the case of small differences between weightings. In terms of meaning, what distinguishes a warp-2 link and a warp-3 link? Moreover, what would become of warp geometries in a dynamically evolving, multi-user hypertext application? Why should a deformable information space be mapped in the same way twice? Fortunately from one point of view, though unfortunately for those expecting precision, the metaphor of folded semantic space renders such mathematical difficulties moot. The coefficients proposed here are not precise, digital measurements, but rather evocative descriptions of a continuous, flexible, analog space. The metaphor of warped space, referring as it does to an imaginary enterprise (or *Enterprise*), reminds us that the graphically generated link weightings are partly metaphoric, one might even say science-fictional. They thus allow us to mediate between architectonic and semantic spaces. If these coefficients have any utility in actual applications, it must be as a rough indicator rather than as a strict metric; still, their theoretical significance seems less ambiguous.

Design Exercise 2: A Map of Detours

Design Exercise 1 reveals that proximity has two meanings, only one of which is usually mapped. Thus, in systems like SEPIA, Storyspace, Aquanet, NoteCards, and VNS, where both global and local maps emerge as users inscribe nodes and links, the maps represent chiefly geometric or architectonic notions of proximity. Although Storyspace offers writers and readers three representations or visualizations of the emerging hypertextual structure — storyspace view, chart view, and outline view, all three representations in fact map only the architectonic dimension of the hypertextual space. All three occlude the semantics of cross-hierarchical linkages, privileging whatever semantics the geometries can construct and maintain at the expense of the semantic density and multiplicity of linguistic connections.

Fisheye views represent schemes for mapping both the global architecture of the hyperdocument and one or more regional "contents," what Noik [24] terms the local detail in an area of interest. Noik conceptualizes the representational problem in terms of links from a user's current location that lead to a geometrically distant "region" of the document. The system is designed to visualize three aspects of the totality simultaneously: the architecture of the entire hyperdocument, displayed with a high level of abstraction, the current area of interest, displayed with a high level of enlarged detail, and one or more other (related) areas of interest, displayed with more detail than the global representation but with less detail than the current focal point.

The goal of Noik's fisheye representation is to distort the visual elements of the global map so as to foreground or enlarge the target regions, the areas of interest the reader has identified. In developing this scheme, Noik identifies two important anchors: a current position or focal point (FP) and a desire or query intended to explore possible "next" locations. Possible "next" regions can have variable Degrees of Interest (DOI) which influence the size of their graphic representation. This scheme recognizes that "nextness" or semantic relatedness need not be confined either to geometrically proximate spaces or to those explicitly linked: the conceptual relations can be expressed as collocations or gatherings of nodes according to some other user-defined property.

The problem with Noik's scheme for maintaining a global (architectonic) perspective while users actually attend to local detail is that the global is already well expressed in a hierarchical table of contents. That is, any document so fully structured as the two to which Noik refers — a Free Trade Agreement and an extensive dictionary — can be represented for readers' purposes as an outline. For example, Bolter's hypertext *Writing Space* offers users several representations of the document's structure. Figure 2 shows a geometric map (nested boxes) and Figure 3 a conceptual hierarchy (essentially a dynamic table of contents). Neither represents the semantic space of the text, the text a user encounters when she reads the document through the EasyReader interface, where the links represent multiple opportunities, multiple proximities, only some of which are "near" the current location in architectonic terms. As long as some sort of table of contents is always available to a user of such a highly structured document, a graphical representation of that structure serves chiefly as an orienting landmark for "users who set out to understand the *overall structure* of a large hypertext" (Noik [24], p 192, emphasis added), where "structure" seems to mean the nested hierarchy or architectonic space.

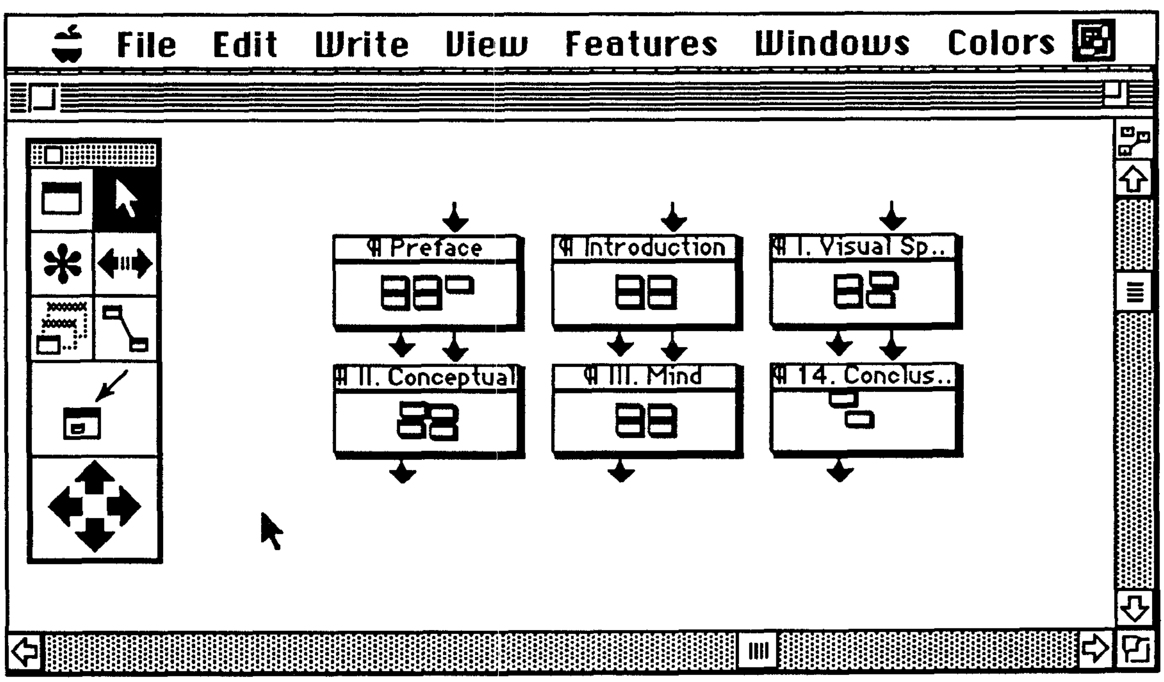

Figure 2: Storyspace view (geometric map) of a portion of Bolter's Writing Space.

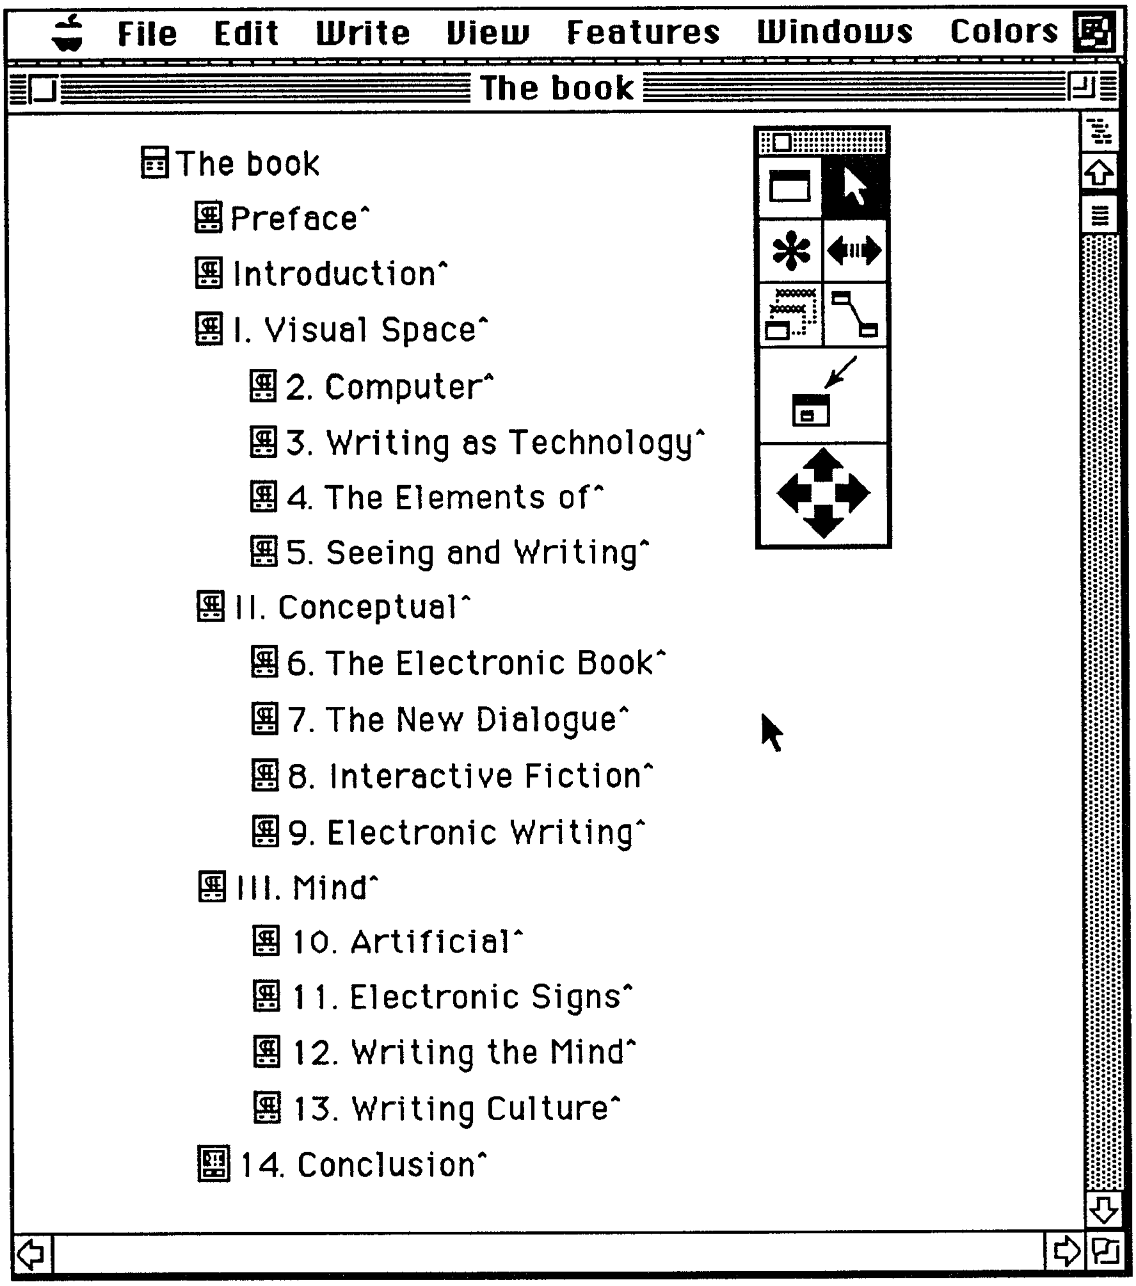

Figure 3: Outline view (conceptual hierarchy) for the same portion of Writing Space.

Instead of pursuing the double representation of Noik's system, we could suppose that an appropriate representation of the semantic space need be only local. The design issue would then be, what constitutes locality: where can we usefully draw the boundaries of the neighborhood? Clearly the current location remains, in Noik's terminology, the focal point. Yet, the focal point of a user's experience is necessarily unstable: it shifts each time a user "moves" to a new location or warps the document space around her. What the user may need to know at each transition may not be where he or she "is" in relation to a geometry of the whole document but where he or she stands relative to some set of potential next focal points. That set constitutes the relevant Points of Interest (POI).

One way to represent the relationship between "now" (FP) and "next" (POI) is to display information about the destinations accessible from the current location. But such a limited representation of locality reduces every link's function to successive replacement — what one mass media critic derisively calls the grammar of "Now... this" [28]. Our previous thought experiment constructs just such a limited grammar. Each link anchor points to a potential POI while the "warp coefficient" signifies the relations between the architectonic structure of the document and a semantic interest the user might pursue. However, displaying a set of possible next nodes fails to tell the user very much about the trajectory or contours of her reading [3]. A more helpful view might situate the user in relation to all POI as well as the set of Subsequent Points of Interest (SPOI). SPOI are all those nodes accessible through the current set POI. The resulting depiction, Figure 4, would look much like a familiar tree diagram.

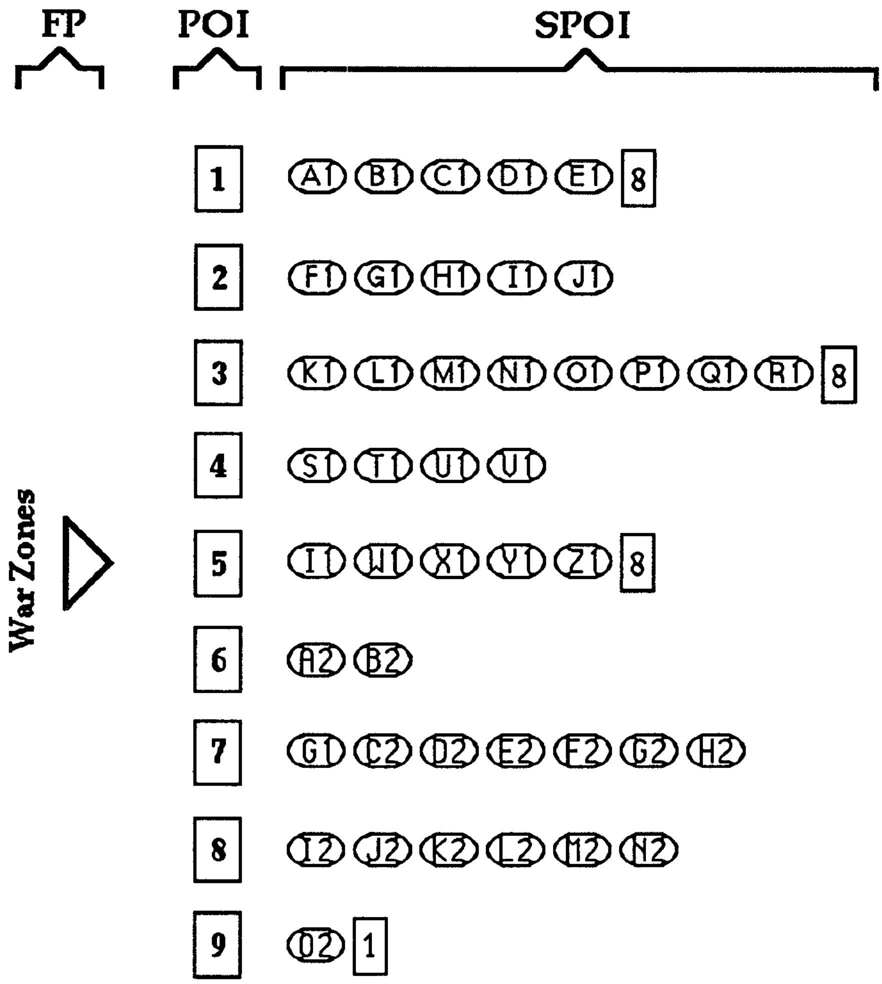

Figure 4: The node "War Zones" from Victory Garden *as FP, with representation of its POI and SPOI.*

The head of the tree, the FP, leads to a set POI which in turn yields a set SPOI. However, this representation obscures any potential redundancy. It is difficult to notice that the node marked G1, a descendant from the node marked 7, also appears as a descendant from the node marked 2. The value of this representation might be enhanced, then, if we collapse redundancies and re-introduce link lines between planes. The resulting graphic, Figure 5, makes available to the user a representation of potential elements for exploration in a more liberally defined neighborhood.

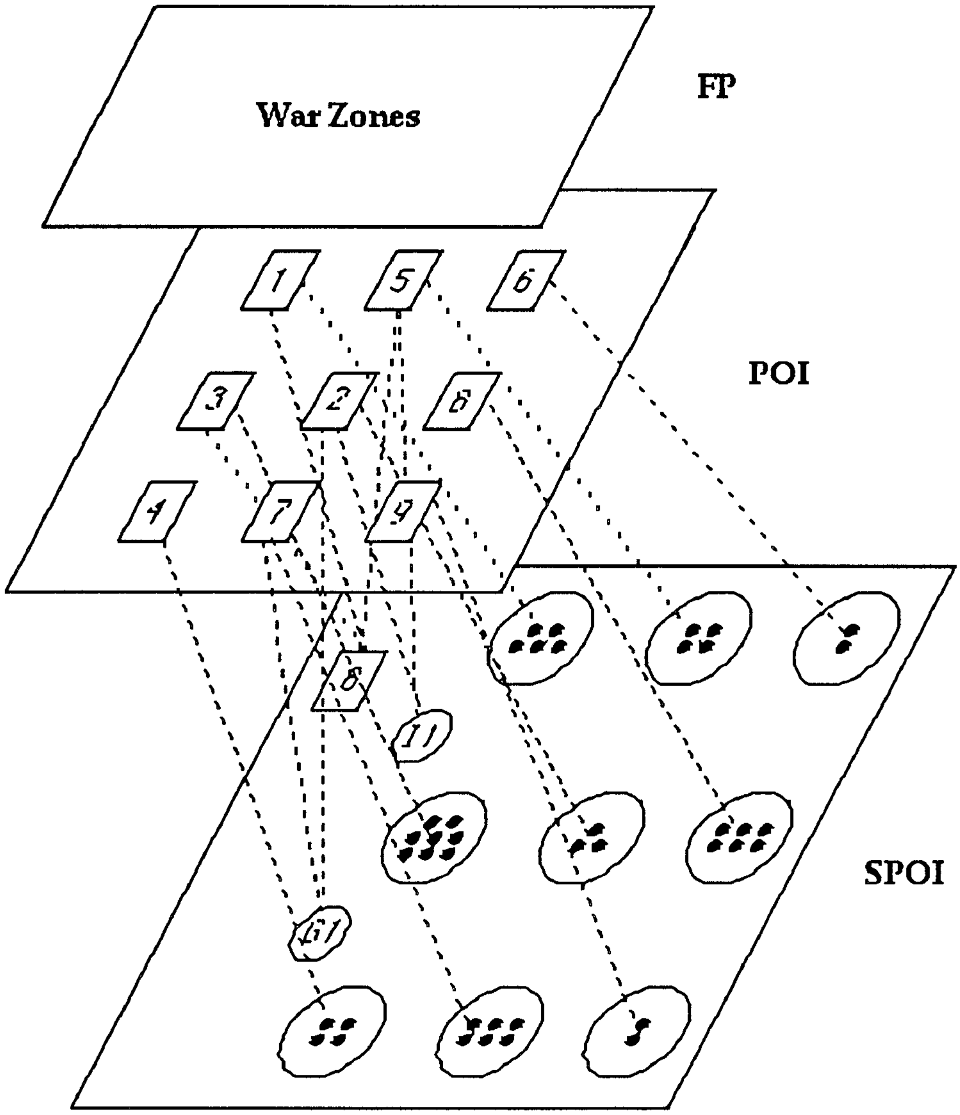

Figure 5: Mapping the neighborhood of the node "War Zones" in Victory Garden.

Though it remains an architectonic depiction, Figure 5 constructs proximity in terms of semantic relation to the current focal point (FP). Moreover, it displays trajectories, some of which may overlap at a later point in the user's interactions. For example, a user choosing the member of set POI marked 5 could subsequently choose the member of set SPOI marked I1. Alternatively, the user could reach I1 through the members 1 or 3 of the set POI.

More important, perhaps, the graphic information displays those avenues *foreclosed* by any given choice. It shows that if the user chooses the member of the set POI labeled 3, she will then be able to reach the unique members of the set SPOI descendant from 3. But she will not then see the unique members of the set SPOI that descend from the nodes marked 8, 6, 1, and so forth, at least not without shifting her current focal point. By declining to display the totality of the document and by suppressing representations of all nodes in the system in favor of only proximate and nearly proximate semantic regions, this strategy attends to the succession of Harpoldian detours that define the user's interactions with the text. It might thus be regarded as another means of mediating between screen real estate and conceptual space.

Yet this mapping scheme is subject to breakdown just like any other. The map shown in Figure 5 may represent the complexity of textual relations within a document. It may even suggest how these relations are, in Nelson's words, "deeply intertwingled." However, it cannot reliably depict Harpold's "place you never get to," since in at least some cases, that place may exist only as a hypothetical alternative in the mind of the user. Like the weighted linking scheme of the first exercise, this relativistic map can *suggest* the dimensions of semantic space, but it cannot exhaustively represent them. Semantic space thus constitutes an inevitable limiting factor for any architectonic representation

5 CONCLUSION: DESIGN LIMITS

Both our design exercises invoke a more problematic definition of space than is usually applied in hypermedia. This is in line with the general emphasis on complex spatiality in virtualizing information systems. Novak [26] describes design principles for virtual reality as "liquid architecture." In our notion of semantic space, we have tried to develop a similar idea for hypermedia. Like Novak, we believe that users of virtual texts must be able to situate themselves within a dynamic information space; hence our interest in warping rather than navigation, and in relativistic local maps instead of global representations. But unlike some proponents of all-encompassing or "seamless" artificial environments (for instance Laurel [18]), we believe that any attempt to represent this internal situation in stable, objective terms must inevitably reach a point of obvious constraint. If our designs are to reflect an intelligent anticipation of such breakdowns, we must understand that the two domains of virtual space, the architectonic space of mapping and the semantic space of conceptual development, do not perfectly correspond.

Having arrived at a very similar conclusion, Dillon et al. suggest that we must refine the conceptual apparatus we use to describe information spaces: "Our view is that the precise nature of the representation is less important to workers in the field of interactive technology than the insights any theory or model of navigation provides. To this extent we propose that a model based on schema theory and including landmarks, routes and surveys as instantiations of basic knowledge is of some utility in considering the design of electronic information spaces" ([8] p 175). This conclusion represents a pragmatic, positivistic approach to what Norman [25] calls the "Gulf of Execution" — the inherent mismatch between architectonic and semantic spaces. As Norman argues more generally, such an approach gives rise to an immensely valuable enterprise, the field of cognitive engineering, with its emphasis on user-centered system design. Though we do not question the importance of this approach for the development of interactive technologies, we suggest that cognitive engineers may benefit from the critical observations of their co-workers in ontological design. Dillon, McKnight, and Richardson are no doubt correct in their call for better spatial metrics and metaphors for hypermedia. In their very limited ways, our design exercises respond to this call. But the obvious limitations of our exercises demonstrate that our pragmatism functions only within clear constraints. Semantic and architectonic spaces cannot be perfectly reconciled. We should aim for systems that harmonize the two as well as possible, but which acknowledge the contingent nature of any such harmony.

REFERENCES

[1] Benedikt, M. "Cyberspace: Some Proposals", in M. Benedikt (ed.), *Cyberspace: First Steps*, pp. 119-224, MIT Press, 1991.

[2] Bernstein, M., J. Bolter, M., and E. Mylonas.. "Architectures for Volatile Hypertext", *Hypertext '91 Proceedings*, ACM Hypertext Conference (San Antonio TX) pp. 243-260, December, 1991.

[3] Bernstein, M., M. Joyce, and D. Levine. "Contours of Constructive Hypertexts",. *Proceedings of the ACM Hypertext Conference*, Second European Conference on Hypertext (Milano) pp. 161-170, November 1992.

[4] Bolter, J. *Writing Space: The Computer, Hypertext, and the History of Writing*, Lawrence Erlbaum Associates, 1991.

[5] Carlson, P. "The rhetoric of hypertext", *Hypermedia* 2, 2, pp. 109-132, 1990.

[6] Charney, D. "The Effect of Hypertext on Processes of Reading", in C. Selfe and S. Hilligoss (eds.), *Literacy and Computers: The Complications of Teaching and Learning with Technology*, pp. 238-263, Modern Language Association, 1994.

[7] DeRose, S. "Expanding the Notion of Links",. *Hypertext '89 Proceedings*, ACM Hypertext Conference (Pittsburgh PA), pp. 249-258, November 1991.

[8] Dillon, A., C. McKnight, and J. Richardson.. "Space — The Final Chapter *or* Why Physical Representations are not Semantic Intentions", in C. McKnight, A. Dillon, and J. Richardson (eds.), *Hypertext: A Psychological Perspective*, pp. 169-91, Ellis Horwood, 1993.

[9] Egan, D., M. Lesk, R. Ketchum, C. Lochbaum, J.. Remde, M. Littman, and T. Landauer. "Hypertext for the Electronic Library? CORE Sample Results", *Hypertext '91 Proceedings*, ACM Hypertext Conference (San Antonio TX) pp. 299-312, December 1991.

[10] Gloor, P. "CYBERMAP — Yet Another Way of Navigating in Hyperspace", *Hypertext '91 Proceedings*, ACM Hypertext Conference (San Antonio TX), pp. 107-21, December 1991.

[11] Halasz, F. "Reflections on NoteCards: Seven Issues for the Next Generation of Hypermedia Systems", *Communications of the ACM*, 31, 7, pp. 836-855, 1988.

[12] Harpold, T. "Threnody: Psychoanalytic Digressions on the Subject of Hypertexts", in P. Delany and G. Landow (eds.), *Hypermedia and Literary Studies*, pp. 171-184, MIT Press, 1991.

[13] Hofstadter, D. *Gödel, Escher, Bach: An Eternal Golden Braid.* Basic Books, New York, 1979.

[14] Jones, R. and R. Spiro. "Imagined Conversations: The Relevance of Hypertext, Pragmatism, and Cognitive Flexibility Theory to the Interpretation of 'Classic texts' in Intellectual History", *Proceedings of the ACM Hypertext Conference*, Second European Conference on Hypertext (Milano) pp. 141-148, November 1992.

[15] Lai, P. and U. Manber. "Flying through Hypertext", *Hypertext '91 Proceedings*, ACM Hypertext Conference (San Antonio TX), pp. 123-132, December 1991.

[16] Landow, G. "Relationally Encoded Links and the Rhetoric of Hypertext", *Hypertext '87 Papers*, ACM Hypertext Conference (Chapel Hill NC), pp. 331-344, November 1987.

[17] Landow, G. *Hypertext: The Convergence of Contemporary Critical Theory and Technology*, Johns Hopkins University Press, 1992.

[18] Laurel, B. "Interface as Mimesis", in D. Norman and S. Draper (eds.), *User Centered System Design: New Perspectives on Human-Computer Interaction*, pp. 67-86, Lawrence Erlbaum, 1986.

[19] Lesk, M. "What to do When There's Too Much information", *Hypertext '89 Proceedings*, ACM Hypertext Confernece (Pittsburgh PA), pp. 305-318, November 1989.

[20] Marshall, C., F. Halasz, R. Rogers, and W.. Janssen, Jr. "Aquanet: A Hypertext Tool to Hold Your Knowledge in Place", *Hypertext '91 Proceedings*, ACM Hypertext Conference (San Antonio TX), pp. 261-275, December 1991.

[21] Marshall, C. and F. Shipman III. "Searching for the Missing Link: Discovering Implicit Structure in Spatial Hypertext", *Hypertext '93 Proceedings*, ACM Hypertext Conference (Seattle WA), pp. 217-230, November 1993.

[22] Moulthrop, S. "Toward a Rhetoric of Informating Texts", *Proceedings of the ACM Hypertext Conference*, Second European Conference on Hypertext (Milano), pp. 171-180, November 1992.

[23] Nanard, J. and M. Nanard. "Should Anchors Be Typed Too? An Experiment with MacWeb", *Hypertext '93 Proceedings*, ACM Hypertext Conference (Seattle WA), pp. 51-62, November 1993.

[24] Noik, E. "Exploring Large Hyperdocuments: Fisheye Views of Nested Networks", *Hypertext '93 Proceedings*, ACM Hypertext Conference (Seattle WA), pp. 192-205, November 1993.

[25] Norman, D. "Cognitive Engineering", in D. Norman and S. Draper (eds.), *User Centered System Design: New Perspectives on Human-Computer Interaction*, pp. 31-61, Lawrence Erlbaum, 1986.

[26] Novak, M. "Liquid Architectures in Cyberspace", in Michael Benedikt (ed.), *Cyberspace: First Steps*, pp. 225-254, MIT Press, 1991.

[27] Parunak, H. "Don't Link Me In: Set-Based Hypermedia for Taxonomic Reasoning", *Hypertext '91 Proceedings*, ACM Hypertext Conference (San Antonio TX), pp. 233-242, December 1991.

[28] Postman, N. *Amusing Ourselves to Death: Political Discourse in the Age of Show Business*, Penguin, 1985.

[29] Raskin, J. "The Hype in Hypertext: A Critique", *Hypertext '87 Papers*. ACM Hypertext Conference (Chapel Hill NC), pp. 325-330, November 1987.

[30] Schneiderman, B. "User Interface Design for the HyperTIES Electronic Encyclopedia", *Hypertext '87 Papers*. ACM Hypertext Conference (Chapel Hill NC), pp. 189-94, November 1987.

[31] Sternbach, R. and M. Okuda. *Star Trek: The Next Generation Technical Manual*. Pocket Books, 1991.

[32] Stotts, P., R. Furuta, and J. Ruiz. "Hyperdocuments as Automata: Trace-Based Browsing Property Verification", *Proceedings of the ACM Conference on Hypertext*, Second European Conference on Hypertext (Milano), pp. 272-281, November 1992.

[33] Winograd, T. and F. Flores. *Understanding Computers and Cognition: A New Foundation for Design*. Addison Wesley, 1985.

[34] Wexelblat, A. "Giving Meaning to Semantic Spaces", in M. Benedikt (ed.), *Cyberspace: First Steps*, pp. 255-272, MIT Press, 1991.

[35] Wright, P. "Cognitive Overheads and Prostheses: Some Issues in Evaluating Hypertexts", *Hypertext '91 Proceedings*, ACM Hypertext Conference (San Antonio TX), pp. 1-12, December 1991.