

VIKI: Spatial Hypertext Supporting Emergent Structure

Catherine C. Marshall, Frank M. Shipman III, James H. Coombs
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, CA 94304
Fax: (415) 812 - 4777
E-mail: {marshall, shipman}@parc.xerox.com

ABSTRACT

The emergent nature of structure is a crucial, but often ignored, constraint on authoring hypertexts. VIKI is a spatial hypertext system that supports the emergent qualities of structure and the abstractions that guide its creation. We have found that a visual/spatial metaphor for hypertext allows people to express the nuances of structure, especially ambiguous, partial, or emerging structure, more easily. VIKI supports interpretation of a collected body of materials, a task that becomes increasingly important with the availability of on-line information sources. The tool's data model includes semi-structured objects, collections that provide the basis for spatial navigation, and object composites, all of which may evolve into types. A spatial parser supports this evolution and enhances user interaction with changing, visually apparent organizations.

KEYWORDS: Spatial hypertext, emergent structure, interpretation, visual structure recognition, composites.

1 INTRODUCTION

When we examine hypertext structures from a systems perspective, we often focus on the finished product -- how the expressive qualities of a particular system's data model are reflected in the structures, how readers navigate through the structures, and how the structures may be interchanged between diverse systems [4]. Less frequently do we question how the structures emerge and how systems support the structuring process. Yet developing hypertext structures and the abstractions that guide their creation is a crucial part of authoring.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copyright is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

We have developed a spatial hypertext system, VIKI, that supports the emergent qualities of hypertext structures, acknowledging many of the difficulties we have observed when people define and use the types of abstractions that are prevalent in many hypertext systems [25]. Spatial or browser-based hypertext has unique expressive qualities that take advantage of the human perceptual system, spatial and geographic memory, and more generally, spatial intelligence. VIKI's visual/spatial methods of constructing hypertext structures facilitate exploration and the gradual development of a visual language through informal interaction.

Instead of requiring authors to construct explicit structures, the prototype tool incorporates structure-finding algorithms based on spatial layout and other visually salient properties of objects; this approach is complementary to other text- or document-based approaches such as those described in [26]. Because people are able to leave structure implicit to be recognized and formalized on demand, VIKI overcomes many of the authoring problems associated with hypertext construction, including the need to articulate tacit knowledge or commit to structure prematurely.

VIKI is intended for exploratory interpretation, making sense of a body of electronic information collected in the service of a specific task such as analysis, design, or evaluation. This kind of information-intensive intellectual task has become increasingly important in the light of recent efforts to provide extensive digital resources like libraries, databases, and on-line information services. In essence, spatial hypertext provides an information workspace, the critical vehicle for making public or institutional information resources useful to individuals, work groups, or larger communities of practice.

VIKI uses a simplified hypertext data model that is derived from our experiences with Aquanet [15] and additional empirical studies of spatial information structures [16]. Like Aquanet, VIKI uses semi-structured objects as the atomic entity. But diverging from Aquanet, complex entity-valued relations have been broken down into collections, a system-supported hierarchy of navigable information

spaces, and composites, higher-level structures composed of regular spatial patterns of objects and collections. The tool's user interface provides users with ways of manipulating the visual and spatial characteristics of individual references to underlying objects, collections, and composites.

In this paper, we first explore the niche for spatial hypertext and look at its role in interpretive tasks. We then describe VIKI's data model and show how it is realized in the user interface. We discuss how spatial parsing is integrated into the tool and how it is used to assist users in their tasks. Finally, we discuss use of the tool in a group setting and future directions for the work.

2 SPATIAL HYPERTEXT AND ITS NICHE

We have argued previously that spatial hypertext fills an important niche in our collective conception of hypertext [16]. If we take a representational stance on structures (as opposed to a purely navigational one), and focus on active creation rather than reading and traversal, we can examine the constraints systems put on authors and the support the systems provide toward meeting these constraints. This perspective highlights the niche for spatial hypertext and the role it may play in supporting structure creation.

We have identified a spectrum of approaches to constraining authoring in browser-based and document-based hypertext systems and will contrast four points along this spectrum: (1) permissive, (2) emergent, (3) descriptive or meta-schematic, and (4) prescriptive structuring.

Permissive linking (or more broadly structure-building) lies at one end of the authoring spectrum. Most early hypertext efforts relied on explicit, permissive links to express structure. Connections were unconstrained and intimately tied to the authors' and readers' perceptions of idiosyncratic relationships between nodes, documents, or unitary pieces of information. Permissive linking generally allows people to express these relationships, but does not provide any way of expressing the regularities, abstractions, or assumptions underlying the structure-building process.

At the other end of the spectrum, prescriptive structure-building imposes constraints that embody a particular methodology or a cognitive model of process (e.g. Issue-Based design methodology underlies gIBIS [5]; Sepia is grounded in a model of the writing process [27]). The types of connections that authors may make using these systems are the results of careful analysis and are embedded in the systems themselves (although most prescriptive system designers have found it necessary to make their base set of types user-extensible). In turn, the system provides various kinds of support for readers and authors based on these types.

Recent efforts have brought us descriptive or meta-schematic systems in which authors characterize their domains of interest, and use (or, more ideally, re-use) these

abstractions to constrain and structure their own authoring. Aquanet [17] and MacWeb [20] are examples of this approach. Meta-schematic systems provide mechanisms for expressing and manipulating abstractions, in addition to providing the means of creating and manipulating instances of abstract types. They appropriate the power of knowledge representation languages or constructs like frames to help constrain and describe hypertextual structures.

Emergent spatial hypertext arises from the authoring paradox first introduced by more prescriptive systems and underscored by meta-schematic systems: as a system's representational capabilities and data model become increasingly rich (e.g. MacWeb's experiment with typed anchors [21]), they become more difficult to use. Similarly, as facilitative models of activity are fleshed out, users find themselves devoting more and more attention to meta-decisions -- identifying what they are doing and why. A pressing need arises to support more free-form expression and the gradual emergence of structure, possibly through recognition of the implicit forms that arise through use and performance of the task at hand.

Authoring Constraints	Representation of Structure	
	Browser-based	Document-based
Permissive	Storyspace	WWW
Emergent	VIKI	HOS
Descriptive (meta-schematic)	Aquanet, MacWeb	OVAL
Prescriptive	Sepia, gIBIS	PHIDIAS

Table 1 - Prototypical examples of hypertext systems intended to fill different niches

Table 1 sets out the spectrum of approaches to support authoring hypertext structures in current browser-based and document-based systems. Note that our examples thusfar have come primarily from browser-based hypertext systems -- systems that allow their users to work against a spatial backdrop. Comparable approaches exist in hypertext systems that take a document-based approach. The World-Wide Web model [1] and Storyspace [12] support simple, unconstrained linking; PHIDIAS uses an IBIS-derived methodology to constrain the structuring of issue bases [18]; OVAL provides a sophisticated rule-based knowledge representation facility for structuring nodes according to their textual content [13]. HOS uses textual analysis methods to support emergent structuring [26].

Naturally, system developers at either end of the permissive/prescriptive spectrum have found that hybrid approaches are the most effective; to insist on a single authoring paradigm is to render a hypertext tool less usable than it might otherwise be. Permissive hypertext systems

often include link-finding mechanisms to support simple kinds of emergent structure implicit in content, such as Storyspace's link apprentice [2]. We have seen proposals afoot to introduce prescriptive authoring constraints into the global community of World-Wide Web contributors. Authors can create Sepia hyperdocuments starting in the informal Dolphin system [9]; gIBIS has an "Other" node type that acts as an escape mechanism; and PHIDIAS allows users to enter untyped material for future structuring.

In VIKI, we focus on visual/spatial hypertext because we feel it has unique characteristics in its ability to support emergent structure. It allows people to make more complete use of their perceptual abilities by providing expressive flexibility, a basis for spatial memory, and navigational control of global context. We direct our attention to a task that can most benefit from techniques to support emergent structure, the interpretation of large collections of digital material.

3 THE TASK: INTERPRETATION

We have focused our efforts on supporting the interpretive process. Interpretation is that part of writing, collaborating, or thinking in which people collect the materials of interest and make sense of them in the light of their task and the background they bring to it. Of course, groups and individuals are neither monolithic nor consistent in their beliefs; to support interpretation, we consider ways of allowing people to record varying interpretations of a body of materials, using a multiplicity of frameworks or structuring schemas which make very different use of the same content.

Why is interpretive hypertext an important avenue of research to pursue? With increasing availability of on-line information sources (including other people, in addition to databases and corpora), people have access to vast quantities of materials relevant to their work. Yet there are relatively few tools that allow people to collect diverse materials and actively explore the various ways the materials may be interpreted -- to pioneer new ways of looking at the material.

For example, a successful product development organization will often collect business intelligence material relevant to the positioning and viability of a new product. This business intelligence material may include electronic announcements and reviews of related products; competitors' home pages from World-Wide Web; relevant patent documents retrieved from an information provider's extensive database; internal, project-related materials (like engineering documents) created during the course of product development; electronic mail exchanges held within the organization; and, perhaps, information mined from Usenet newsgroup discussions. Once these materials have been gathered, members of the organization must make sense of them -- they need to decide which material is important, extract the salient information, evaluate its veracity, and

organize it for their purposes. So it is the understanding that a group develops and shares over time that makes a collection of materials valuable; the materials themselves are of limited utility without this shared understanding. The difficulty lies in expressing and recording this mutually-defined, ever-evolving interpretive structure.

Thus, interpretive structure is not inherent in document content, but evolves opportunistically through work. In practice, interpretive structure may arise through filing, in conversation, in juxtapositions of documents on one's desktop, in notes and annotations, in markings and marginalia. All of these methods of expressing interpretations are necessarily partial (they are not fully articulated); they readily tolerate ambiguity and fuzziness. Relationships among documents, assessments of their contents, and interpretive abstractions can remain tacit through the entire duration of the sense-making process. Indeed, much of human knowledge and skill will always remain tacit [24].

But externalizing these relationships, assessments, and interpretive abstractions is a crucial part of reflecting on one's own understanding of a problem and communicating one's understanding to others during the course of collaborative work. Our experiences demonstrate that structure is difficult to express, even in hypertext, a medium designed to do just that.

We find then that a tension arises between the tacit knowledge of the expert practitioner (the reader/author), and the representational affordances the hypertext tool provides to coax out interpretations as partial, emerging forms.¹

Interpretive structure is often difficult to articulate within the bounds of language, no matter how informal; visual and kinesthetic modes of expression may be far more appropriate for coding pre-articulate ideas [19]. One of Aquanet's unexpected strengths was the ability it gave people to express interpretations -- interpretations that were less than fully formed -- in terms of visual appearance or spatial positioning. Users could create categories or sets without names and note relationships by relative spatial position; they could use the kinesthetic process of "trying things out" (as one might do wiggling molecular models in space or moving a jigsaw puzzle piece into different orientations). In other words, they could use affordances normally provided by manipulable material objects (for example, paper documents).

Analysis of Aquanet's expressive strengths and shortfalls points out three pivotal requirements we used to drive the development of VIKI: (1) interaction must be informal; (2)

¹ We borrow this notion of affordances from Norman and Gaver [22][7]; for people to use a technology effectively, its affordances must be made apparent. So in our design of a visual/spatial hypertext tool, the resources for coding interpretations visually and spatially must be apparent and accessible.

users must have access to a variety of representational modes; (3) additional structure must be available on demand, either through the development of types from examples and prototypes, or through recognition of implicit structure.

First, informal interaction is crucial to coaxing out emerging interpretations. We have found that if people can't just "point and type" their text into a tool, their trains of thought will be derailed by the tool's interface. Furthermore, if they are required to categorize their input (say, by choosing a type before an instance can be created), they will be unable to use the tool to express partially formed ideas. When we observed Aquanet in use, we saw that people were frustrated by the need to specify what type of instance they were creating before they entered any text.

Second, different representational modes -- e.g. the ability to exploit spatial positioning and visual appearance -- allow people to capture nuances of meaning. We wanted to give people access to retinal and planar variables [3] and allow them to create visually apparent codings [29] to express and remember partial, ambiguous, or inarticulable interpretations.

Finally, emergent structure includes both the particular (relationships among instances) and the general (abstraction of node-internal structure or among-node relationships). To be effective, an interpretive tool must allow both kinds of structure to emerge and become schematic on demand.

4 DATA MODEL FOR SPATIAL HYPERTEXT

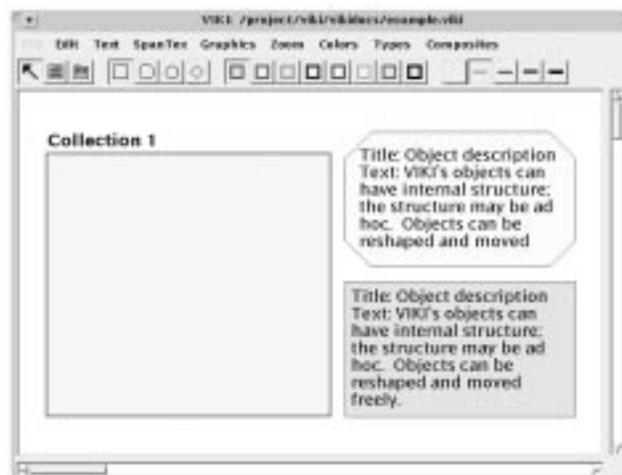
To support emergent spatial structures, VIKI's data model has three fundamental kinds of elements, *objects*, *collections*, and *composites*. Objects are semi-structured nodes; collections and composites, as we discuss below, are specialized kinds of hypertext structure derived from observations of practice, not only in Aquanet, but in other hypertext systems and other media [16]. Objects, collections, and composites may all be abstracted into semi-structured types in accordance with the requirement we have noted for abstraction on demand.

Objects. Objects are the content-holding entities of the data model. Each object is referred to by one or more visual symbols. Visual symbols each have a graphical appearance, determined either by direct manipulation or through the types mechanism described below. Users can manipulate the size, shape, color, line thickness, and font characteristics of individual visual symbols. A new object (and the visual symbol referring to it) is created by a mouse-click.

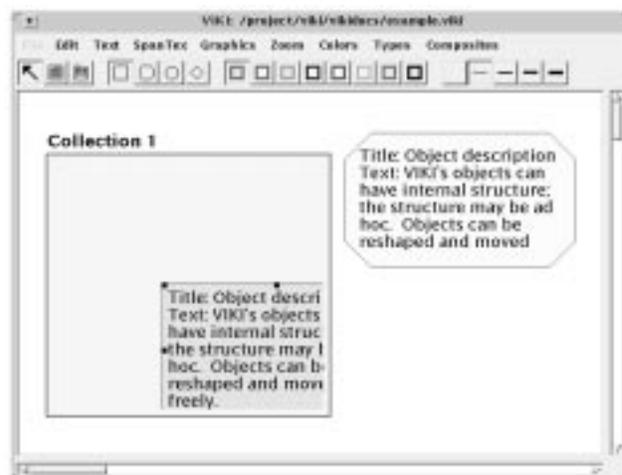
In contrast to Aquanet's tight correspondence between instances of object types and their graphical appearance, VIKI's visual symbols decouple objects and particular references to them in the display space. This decoupling, along with the ability to manipulate the appearance of individual symbols provides users with a far more nuanced

means of expression. In one context, a visual symbol may be a red lozenge with a thick border, signaling, for example, that it's an important source of information; from another spatial context, the same object might appear as a yellow rectangle, indicating that it's a newswire story about a new teleconferencing product.

Figure 2a shows two visual symbols (a light gray lozenge and a darker gray rectangle) that refer to the same underlying object. Note that the content displayed is the same.



(a)



(b)

Figure 2 - Objects and collections in VIKI.

Objects are semi-structured; they may have no internal structure, or they may have an unlimited number of text-valued slots. Users create slots by name; these slots can be added when users encounter or invent content that doesn't fit within the object's existing slots. In other words, slots can be idiosyncratic reflections of the content of an individual object. Thus instances of a type may have structure in addition to that prescribed by the type's abstract description; untyped objects may have any desired internal structure to highlight or extract specific portions of content. The object

in Figure 2a has two slots that are visible on its visual symbols, Title and Text.

Visual symbols are small, manipulable references to an object's content. Indeed users have much control over what is visible on a visual symbol at a given time. They can use the symbol's size to limit the amount of text revealed; they can specify which slots' contents are shown; and they can scroll through content to focus attention on a desired segment of text.

The whole of an object's content may be viewed (and edited) in a separate text-editing window (*maximized*), as is the case in many document-based hypertext systems. Such a maximized object is shown in Figure 3. Because people frequently need access to the content of more than one object at a time (for example, to compare conflicting information), multiple objects may be maximized at a time.



Figure 3 - A maximized object.

Collections. Collections contain an arbitrary spatial arrangement of any number of objects or other collections, allowing them to form a system-supported hierarchy. They too are manipulable visual symbols that refer to large, scrollable 2 1/2 D spaces (spaces in which objects may overlap, but no explicit representation of a third dimension is saved across sessions). Like objects, collections may also be instances of a type, although collection types just preserve a limited number of visual attributes (color, title font characteristics, and border thickness, but not shape). Figure 2a shows a single collection titled "Collection 1." Figure 2b shows the movement of a visual symbol into the collection. The edges of the collection act as a clipping region -- a portal to a subspace. Users may resize or zoom collections to reveal more or less of their contents.

We discuss how collections act as subspaces to localize context and how they support user navigation in a later section.

Composites. Composites are lightweight structures that consist of combinations of two or more objects or collections in particular visual/spatial configurations.

Composite constituents may be typed and have additional visual/spatial properties that are independent of the composite type definition, once again due to the object/symbol distinction. For example, one constituent of an instance of a composite type may have a thicker border than in the composite's canonical form to encode a specific property of the instance. Composite types may be defined directly, but they can also be suggested to authors based on the results of the recognition algorithms discussed in the next section.

Figure 4d shows an example of two composites which are instances of a composite type. One instance overrides a specific property -- border thickness -- on one of its component objects.

We see composites (especially automatically perceived and maintained composites) as a viable way of building up relational structures in VIKI; they closely reflect the structuring process people went through in Aquanet [15]. Composites provide a simplified mechanism for including entity-valued slots in VIKI's data model. VIKI's composites follow Groenbaek and Trigg's composites, which were derived in part from NoteCards tabletops (where a lightweight structure specified graphical layout of existing nodes) [8][28].

Types. Although tools like Aquanet or OVAL provide their users with powerful inheritance mechanisms for creating types and type hierarchies, in practice these mechanisms can prove to be very difficult to use in an informally structured task; we saw little use of Aquanet's CLOS-inspired mechanism that supported multiple inheritance. Strongly typed hypertext and author-defined type hierarchies are not entirely compatible with the notion of emergent structures, since they presuppose that a user has a clear (and continuous) idea about how the domain he or she is working in is structured and is willing to work hard to articulate it within the expressed framework. We discuss other problems arising from the emergent quality of interpretive structures in [25].

Yet types have some useful characteristics: they encapsulate a particular appearance, set of authoring conventions, or internal structure. Thus we have tried to strike a middle ground in VIKI's type mechanism: it allows users to specify significant aspects of an instance's visual symbol (like color, font properties, and shape) and the internal structure of an object, along with how this internal structure is manifested on the visual symbols that refer to it. We have also anticipated that types will emerge (be designed and defined) gradually through examples and prototypical cases rather than through a separate process.

To introduce flexibility into the types mechanism, VIKI allows any particular visual symbol's graphical appearance to be manipulated independently of its object's type and ad hoc slots to be added to its internal structure. Because some of the visual properties of the object are not specified by the type definition (for example, border thickness), an author

can express a partial interpretation by using these independent properties.

VIKI also helps users recognize instances of types -- if the internal structure of an object matches that of an existing type, the tool asks the user if she intends the object to be an instance of that type.² Other more sophisticated heuristics are used to suggest possible collections and composites based on spatial and visual properties of objects. These mechanisms are described in the next section.

Working within VIKI's data model, authors can use objects, collections, composites, and the visual symbols that refer to them to gradually build up visualizations of the contents of a corpus of collected relevant material, eventually forming types, abstractions, and a sophisticated visual language for performing an analysis of this material.

Noting relationships. To provide users with particular kinds of representational affordances, VIKI focuses on relationships that can be noted visually and spatially. These relationships may either remain implicit, or be recognized by the perceptual heuristics described in the following section.

Figure 4 summarizes the kinds of relationships VIKI's data model supports. First, relationships between objects may be noted very simply, through spatial juxtaposition and placement, as shown in Figure 4a. The visual symbol/object distinction allows users to refer to the same content from many spatial contexts and modify a particular symbol's appearance to suit its use. Graphically noted relationships like drawn arrows are a way this type of relationship can be further formalized. We have opted to let these relationships be indicated very informally, at the possible cost of accidental loss of information, since our observations of practice indicate a preference for easy manipulation of this permissive level of structuring.

A second way of noting relationships (especially categorical relationships and sets) is through visual properties, and -- more explicitly -- through VIKI's simplified object types mechanism. In Figure 4b, the presence of two different visual types is used to express category membership that cross-cuts spatial positioning. In this case, we mean to portray categorization that has taken place after the objects were moved into this spatial configuration.

As we see in Figure 4c, categorical or activity-related connections among objects may also be noted through collections. VIKI's collection mechanism supports two different practices we have observed in Aquanet and NoteCards -- the creation of hierarchical category structures (by means of NoteCards FileBox cards and system-

supported filing links) [10] and the manipulation of objects over time into activity-oriented subspaces [15].

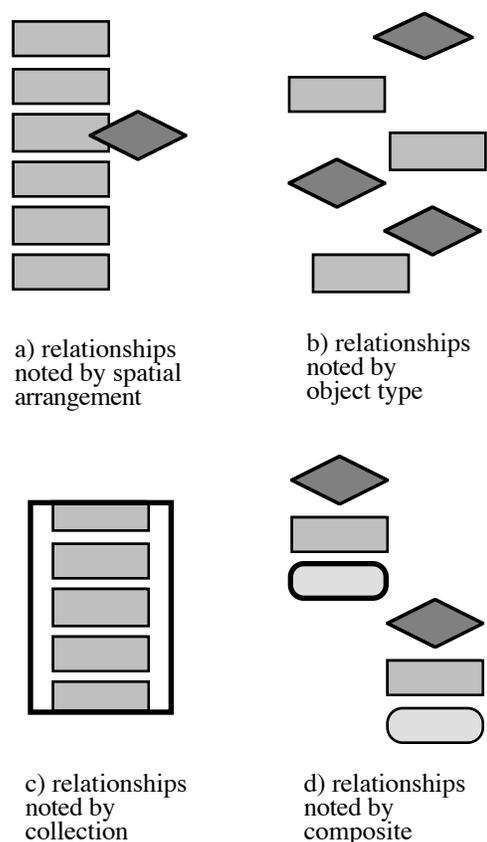


Figure 4- Noting relationships within VIKI's data model.

Visual and typed composites provide users with the ability to abstract relationships among types of objects. These abstract relationships may either remain implicit (as we observed in practice with Aquanet) or users may take advantage of the system's ability to recognize these relationships and form explicit composite types to maintain spatial relationships as we saw in Tabletop templates [14]. This level of abstraction helps bring back some of the expressiveness of Aquanet's data model and its entity-valued slots without placing unwarranted constraints on users who prefer to leave the composite structure implicit. Figure 4d shows an example of two instances a visual/spatial composite which VIKI can recognize and, if desired, formalize into a composite type.

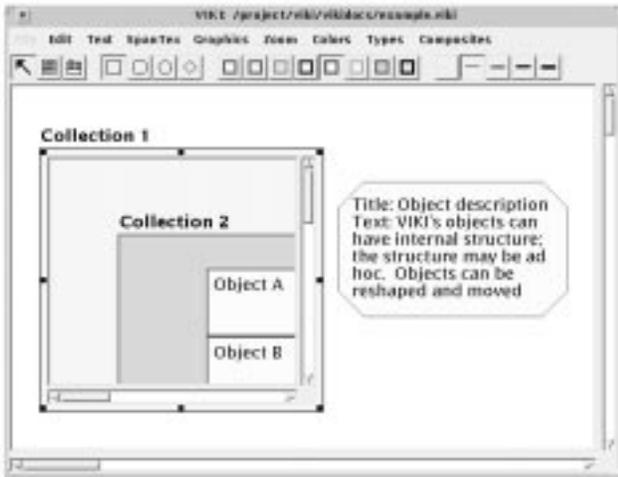
5 SPATIAL NAVIGATION, LOCAL CONTEXT

Collections are a way of localizing meaning -- they give people the ability to immerse themselves in local meaning without losing global context. Context may be defined by representational strategy (to localize a particular visual/spatial structuring schema as we see in the use described in [15]), specific subtask (these items have been brought together because they pertain to an activity, much as was the intent of Rooms [11]), or semantic cohesion

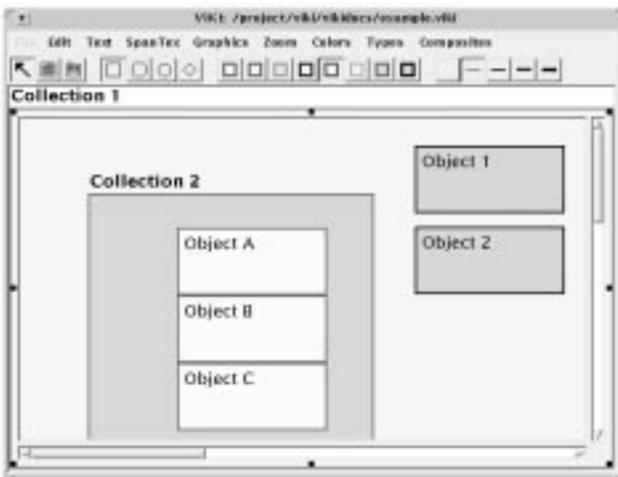
² Of course, if two semantically distinct types have identical internal structures, a user has the option of turning off this suggestion mechanism, and may then resolve the ambiguity by hand.

(these items all belong to the same category or set, in line with the original motivation for relations in Aquanet).

Since collections are acyclic, visually-nested hierarchical structures, we have provided a navigational mechanism that allows people to move between focal collections. Like Boxer [6] and Pad [23], VIKI takes advantage of spatial metaphor for navigation.



(a)



(b)

Figure 5- Navigation between VIKI collections.

Figure 5 shows an example of spatial navigation in VIKI. Figure 5a shows the top level of a VIKI document; two collections are visible, one titled "Collection 1," and a second titled "Collection 2," which is inside Collection 1. A single lozenge-shaped object is also shown at the top level. Traversing into Collection 1 (by double-clicking on its border) causes the selected collection to become the focus -- it takes up the tool's entire available display area. Figure 5b shows the result of this traversal. Now more of the contents of Collection 1 are visible (although remote regions of the collections are always accessible through scrolling, even when they are not the focus). Two more

objects have been revealed in Collection 1 and the three clipped objects in Collection 2 have come more fully into view. To traverse among levels in the hierarchy, a user need only click on the desired collection's border. Multiple levels may be traversed in a single step -- a contrast to the "single level at a time" style of a NoteCards hierarchy (although navigation is necessarily limited by the depth of the hierarchical structure that a user can see).

Because hierarchies of spaces soon grow complex and difficult to traverse, VIKI also includes a Zoom function that -- at its most compressed level -- allows users to grasp the entirety of the space.

6 RECOGNITION

We have integrated a spatial parser (described in [16]) into VIKI to assist authors in a variety of ways. The results of the spatial parser allow VIKI to (1) use implicit structure as a basis for interaction; (2) help people use the object-collection-composite data model; and (3) provide inter-object abstraction on demand. Structure use or abstraction in VIKI is always done at the user's initiative, following our "structure and abstraction on demand" requirement.

The spatial parser performs a bottom-up hierarchical parse using empirically-determined heuristics to decide parse order. The recognition algorithms use three primary attributes of a visual symbol: its x,y location, its width and height, and its object type, which may be thought of in terms of a confluence of visual attributes like shape, color, or font. The parser produces a forest where each branch is assigned a type based on the types of its constituents. If the constituents are of the same type, the branch assumes its type from its constituents; if they are of differing types, the parser creates a new type for the branch.

6.1 Click Selection

The recognition algorithms allow users to interact with implicit structure without requiring the structure to be formally defined or even anticipated. VIKI's click-selection facility is a good example of this strategy.

Click-selection in VIKI works much the same way as it does in a text editor. In a text editor, a single click puts the cursor at a particular point; the next click selects the word; the next, the paragraph; the next the entire document; and the next returns to the single point of selection. VIKI uses a similar technique: each successive click selects the next level of hierarchical structure.

Figure 6 shows click selection in action. With the first click, the user has selected a single object, the middle item in a three-item list. The user's second click enlarges the selection to include all three similar objects in the list. The third click adds the heading of the list to the selection, and the fourth click extends the selection to include the entirety of a similarly-structured list.

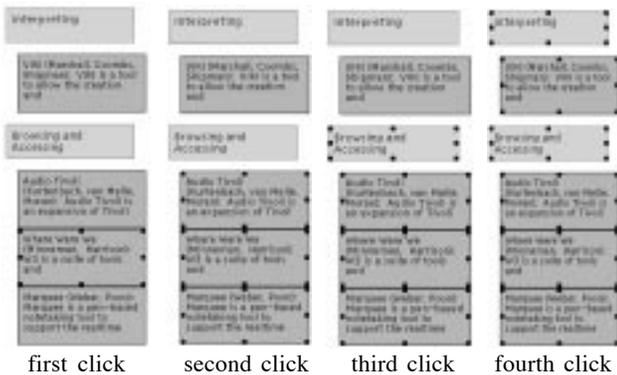


Figure 6- Four steps in VIKI's recognition-based hierarchic click-selection.

Users realize two immediate advantages from hierarchic click-selection. First, users may select objects that are part of partially hidden structures without having to scroll or traverse from the current view. Second, users may select partial structures in areas where objects are too densely packed for sweep-selection.

6.2 Collection Suggestion

The spatial parser also helps users bridge the gap between their activities and the system's data model. For example, at a user's initiative, VIKI will suggest collections -- apparent subdivisions of materials for starting new subspaces. Collection suggestion is an accelerator: it greatly reduces author effort in creating new collections and moving existing, visually structured materials into them.

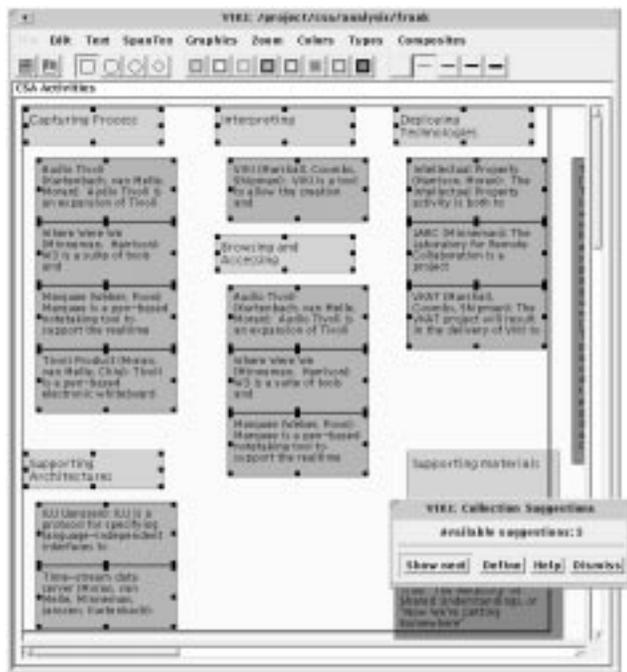


Figure 7- VIKI suggests a new collection at user's request.

VIKI suggests collections based on the highest level structures found by the spatial parser. Collection suggestion uses a standard spelling checker as a model of interaction.

Figure 7 shows the collection suggestion mechanism in action. In the figure, the recognizers have identified an entire region of similarly-arranged materials; a dialog box (in the lower right corner of the screen image) allows the user to define the found region as a new collection. Using this interface, the user can iterate through the list of suggestions and accept only those that are appropriate. VIKI displays these suggestions to the user by selecting all objects and collections that will become part of the new collection and outlining its extent with a dark band. Suggestions can be modified by interactively changing which objects are selected.

6.3 Composite Suggestion

Composite suggestion is an example of how VIKI can use recognition results to provide abstraction on demand. Recurrent spatial patterns of object types are identified by the parser and may become the basis for user-accessible composite types.

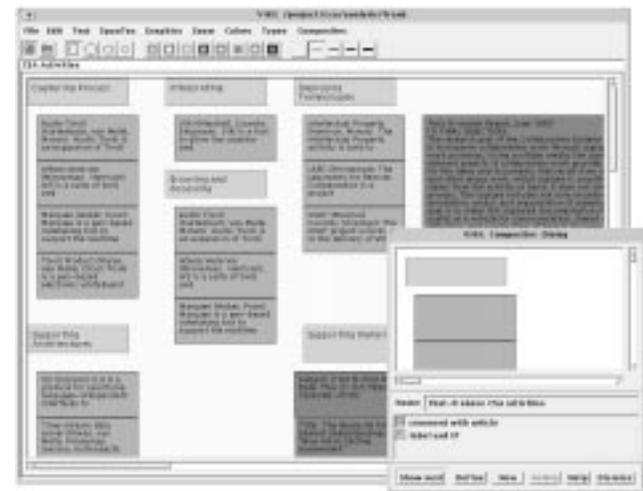


Figure 8- Composite definition dialog with suggestion.

VIKI suggests potential composites based on the parse results whenever a user brings up the composite definition dialog. Users may either use the suggestions as is, modify them, or start from scratch to develop new composite types.

Figure 8 shows the composite definition dialog (in the lower right corner of the figure) with a suggested composite based on a recurrent pattern in the user's work. In this case, the pattern is a list with a heading. In the editable workspace area of the dialog box, a set of abstract objects represents this pattern (a light gray rectangle over a small number of darker gray rectangles). The list of composites that have already been defined appears below the workspace with iconic representations of their graphical appearances.

7 USE IN A WORK GROUP SETTING

We have used VIKI in a work group setting to perform an analysis task: we have gathered and structured materials that are pertinent to our group's research³. In particular, we have focused on three kinds of materials: (1) project descriptions to help develop and maintain an overview of our group's research activities; (2) materials related to intellectual property management; and (3) materials that help us maintain a shared awareness of related products that are being introduced to the marketplace. Although individuals have used VIKI to maintain personal databases of materials, we will focus on its use in a group setting.

7.1 Connection with World-Wide Web

We have found that this kind of information-intensive application requires strong connections with external sources of materials; VIKI users want to refer to various outside resources. Initially, VIKI supported a standard kind of information import -- materials in a fixed format could be brought into the tool and stored in VIKI's database. However, it was quite clear from the start that this would be insufficient to meet the needs of interpretive activities. Constraining people to use VIKI-internal information caused needless replication of materials stored elsewhere, did not take into account the mutability of these materials, and did not take advantage of other tools that people are using to work with networked information.

To address these issues, we have developed a connection with World-Wide Web that allows VIKI objects to refer (through automatic processing of user-supplied URLs) to Web pages. Thus users can start existing WWW viewers (such as Mosaic) by interacting with VIKI objects. This facility allows people to take advantage of the Web as an information resource, enabling them to interpret and situate Web materials in the context of their work. Many viewers suggest this sort of facility is necessary by including the notion of a "Hotlist;" VIKI provides a more expressive medium for customizing views of the World-Wide Web.

In our analysis task, this connection has become quite useful since members of the group have developed Web pages and HTML documents to describe their research; objects in our "project descriptions" collection now refer directly to these pages.

7.2 Meta Communication

Meta communication (communication *about* content) became increasingly important for several reasons. First, many conventions that people use to structure materials remain implicit. Some notes in our application explain

³ Our work group includes a diverse set of activities that center around technologies to support collaboration. VIKI is one of a dozen small projects that take place in our work group. The figures that have been used in the later portions of this paper are drawn from this application.

these conventions. For example, a participant wrote a note to the other contributors: "The thicker borders mean that I think these are the best articles about InConcert. They're arranged in order of usefulness (most to least)."

Second, any individual's structuring process is invisible to her collaborators. Some notes in our application tend to document process: "moved all history notes into the history collection" or "I changed the types Xerox and commercial to display only titles."

Third, even when a collaborator defines explicit types, some explanation of their intended use is often necessary -- "I created a type called 'Post-It' that's for leaving notes around."

Finally, the notes are used for straightforward communication among collaborators; for example, one note says, "There are several places that need 'cleanup'. I'd like to do this together."

7.3 Utility

Naturally, such collaborative efforts are only successful if they are somehow useful to the various contributors. How has the amassed content been used so far? Our project descriptions have become the basis for our group's semi-annual report, a descriptive summary that every work group must produce. When we have conducted meetings with possible business collaborators, we have been able to gather materials over which this discussion may be held. A shared collection of news stories has become a resource that has enabled us to answer provocative management questions like, "How is what you're developing different from this product I saw at Fry's?"⁴

8 ISSUES AND FUTURE DIRECTIONS

This use of VIKI, along with our previous experience with Aquanet use, suggests three fundamental questions: (1) Are the emergent forms supported by visual/spatial hypertext compatible with other important hypertextual forms as specified in data models such as the Dexter reference model [8]? (2) How may other modes of automated analysis be combined with spatial parsing to yield a richer sense of emergent structure? and finally (3) How does this notion of emergent spatial hypertext and shared partial interpretations scale up to larger collaborative settings (beyond the small work group) and larger document collections?

We have demonstrated the ad hoc expressive capacity of visual and spatial hypertext, but it is also important to examine its compatibility with other hypertextual paradigms. In particular, while VIKI allows people to compose internodal structures easily from visually portrayed objects, it does not support the internal anchoring that we

⁴ Fry's is a local chain of electronics and software supermarkets.

have come to expect from fully-formed hypertext structures. It is readily apparent that in explorations of information objects, people will discover relationships among segments of different materials. We can evade this issue by claiming that users may chunk material into smaller objects so the relationships may be expressed between whole objects, but it is more realistic to extend our model to include conventional modes of linking and anchoring. Expressing hypertextual span-to-span relationships will be challenging within our visual/spatial paradigm.

Spatial parsing reveals much perceptual structure in VIKI layouts, but we also anticipate a need for other modes of structure recognition based on analysis of content or analysis of the temporal properties of objects. Once we have developed other modes of recognition, it will be necessary to merge these text-based or temporally-based assessments of internode relationships with spatial parsing; experimentation will allow us to understand how the various analysis methods interact.

Finally, we must examine the utility of spatial hypertext in a larger collaborative setting. Making meaning in shared spaces is comfortable and not overly daunting within a small work group; as we described in the last section, we have already engaged in such collaborations to produce a usable community memory, a shared, gathered, partially interpreted corpus of materials. But how will systems like VIKI scale up to institutional settings or to more casual kinds of collaborations (the kind we observe taking place in some Usenet discussions)? In a tight collaboration, VIKI spaces act as a means of passing around a collection of relevant documents and a partial interpretation of these documents; as the task duration increases, coordination decreases, and the number of contributors, sub-communities, objects, and types increases, will it be possible to maintain the mutual intelligibility of these spaces? Will one person's use of the color red be understood and used the same way as another's? Will collections' ability to localize visual meaning make the entire hierarchical structure understandable, or will inconsistencies between collections render them incomprehensible? All these questions require use in an extended setting.

It is with these issues in mind that we continue development and search for potential user communities.

ACKNOWLEDGMENTS

We thank Sara Bly for many productive discussions about VIKI and its use and Tom Moran for his input to the tool's design. Our ECHT '94 reviewers supplied helpful comments on the paper's content.

REFERENCES

- [1] Berners-Lee, T.J., Cailliau, R., Groff, J-F., and Pollermann, B. World-Wide Web: The Information Universe. *Electronic Networking: Research, Applications and Policy*, Vol. 2, No. 1 (Spring 1992) pp. 52-58.
- [2] Bernstein, M. An Apprentice That Discovers Hypertext Links. In *Proceedings of ECHT '90*, Paris, France, (Nov. 1990), pp. 212-223.
- [3] Bertin, J. *Semiology of Graphics*. Madison, Wisconsin: University of Wisconsin Press, (1983).
- [4] *Communications of the ACM*, Vol. 37, No. 2, (Feb. 1994).
- [5] Conklin, J., and Begeman, M. gIBIS: A Hypertext Tool for Exploratory Policy Discussion. In *Proceedings of the CSCW '88*, Portland, Oregon, (Sept. 26-28, 1988), pp. 140-152.
- [6] diSessa, A., and Abelson, H., Boxer: A Reconstructible Computational Medium. *Communications of the ACM*, Vol. 29, No. 9, (1986), pp. 859-868.
- [7] Gaver, W. Technology affordances. In *Proceedings of CHI '91*, New Orleans, LA (Apr. 27-May 2, 1991) pp. 79-84.
- [8] Groenbaek, K. and Trigg, R.H. Design Issues for a Dexter-Based Hypermedia System. *Communications of the ACM*, Vol. 37, No. 2, (Feb. 1994), pp. 40-49.
- [9] Haake, J.M., Neuwirth, C.M., and Streit, N.A. Coexistence and Transformation of Informal and Formal Structures: Requirements for More Flexible Hypermedia Systems. In *Proceedings of the ECHT '94*, Edinburgh, Scotland, (Sept. 1994).
- [10] Halasz, F.G., Moran, T.P., and Trigg, R.H. NoteCards in a Nutshell. In *Proceedings of the ACM CHI + GI Conference*, Toronto, Ontario (Apr. 5-9, 1987), 45-52.
- [11] Henderson, D.A., Jr. and Card, S.K., Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics*, Vol. 5, No. 3 (July, 1986), 211-243.
- [12] Joyce, M. Storyspace as a hypertext system for writers and readers of varying ability. In *Proceedings of Hypertext '91*, San Antonio, TX, (Dec. 16-18, 1991), pp. 381-387.

- [13] Malone, T.W., Lai, K.Y., and Fry, C. Experiments with Oval: A Radically Tailorable Tool for Cooperative Work. In *Proceedings of CSCW '92*, Toronto, Canada, (Oct. 31 - Nov. 4, 1992), pp. 289-297.
- [14] Marshall C. and Irish P., Guided Tours and On-Line Presentations: How Authors Make Existing Hypertext Intelligible for Readers. In *Proceedings of Hypertext '89*, Pittsburgh, Pennsylvania, (Nov. 5-8, 1989).
- [15] Marshall, C.C. and Rogers, R.A. Two Years before the Mist: Experiences with Aquanet. In *Proceedings of ECHT '92*, Milan, Italy, (Dec. 1-4, 1992), pp. 53-62.
- [16] Marshall, C.C. and Shipman, F. M. Searching for the Missing Link: Discovering Implicit Structure in Spatial Hypertext. In *Proceedings of Hypertext '93*, Seattle, Washington, (Nov. 14-18, 1993), pp. 217-230.
- [17] Marshall, C.C., Halasz, F.G., Rogers, R.A., and Janssen, W.C., Jr. Aquanet: a hypertext tool to hold your knowledge in place. In *Proceedings of Hypertext '91*, San Antonio, Texas, (Dec. 16-18, 1991), pp. 261-275.
- [18] McCall, R., Bennett, P., d'Oronzio, P., Ostwald, J., Shipman, F., and Wallace, N. PHIDIAS: Integrating CAD Graphics into Dynamic Hypertext. In *Proceedings of ECHT '90*, Paris, France (Nov. 1990), pp. 152-165.
- [19] McKim, R.H. *Experiences in Visual Thinking*. Monterey, California: Brooks/Cole Publishing Company, (1972).
- [20] Nanard J. and Nanard M. Using Structured Types to Incorporate Knowledge in Hypertext. In *Proceedings of Hypertext '91*, San Antonio, TX (Dec. 16-18, 1991), pp. 329-342.
- [21] Nanard, J. and Nanard, M. Should Anchors Be Typed Too? An Experiment with MacWeb. In *Proceedings of Hypertext '93*, Seattle, Washington, (Nov. 14-18, 1993) pp. 51-62.
- [22] Norman, D.A. *The Psychology of Everyday Things*. New York: Basic Books, (1988).
- [23] Perlin, K. and Fox D. An Alternative Approach to the Computer Interface. In *Proceedings of Multimedia '93*, Anaheim, California, (Aug. 1-6, 1993), pp. 57-64.
- [24] Polanyi, M. *The Tacit Dimension*. Doubleday, Garden City, NY, 1966.
- [25] Shipman, F.M. and Marshall, C.C. Formality Considered Harmful: Experiences, Emerging Themes, and Directions. Technical report CU-CS-948-93, Department of Computer Science, University of Colorado, Boulder, 1993.
- [26] Shipman, F.M., and McCall, R. Supporting Knowledge-Base Evolution with Incremental Formalization. In *Proceedings of CHI '94*, Boston, Mass., (Apr. 24-28 1994) pp. 285-291.
- [27] Streitz, N., Haake, J., Hannemann, J, Lemke, A., Schuler, W., Scheutt, H., and Thuring, M. SEPIA: a Cooperative Hypermedia Authoring Environment. In *Proceedings of ECHT '92*, Milano, Italy, (Dec. 1-4, 1992), pp 11-22.
- [28] Trigg, R. Guided Tours and Tabletops: Tools for Communicating in a Hypertext Environment. In *Proceedings of CSCW '88*, Portland, Oregon, (Sept. 26-28, 1988), pp. 216-226.
- [29] Tufte, E. *The Visual Display of Quantitative Information*. Cheshire, Connecticut: Graphics Press, (1983).