# Requirements, Characteristics, and Issues for an Information Structures Spatial Hypermedia Environment

**Rosemary Michelle Simpson**
**Brown University Computer Graphics Group**
**rms@cs.brown.edu**

# Introduction

Current research in spatial hypertext has focused on handling emergent structure in a textual or multimedia content authoring environment [Marshall et al. 1994, Shipman et al. 2001, Reinert et al. 1999], or in providing limited gather (agents)/scatter (list explode) and spatial structure identifiers (neighborhoods) in a URL-management environment [Simpson 2001].

What I am proposing differs from previous spatial hypertext work in that it is not a text or multimedia authoring system that has associated link structures; it is a link structure authoring and exploration environment that may have associated text and multimedia data. In a way, it flips the traditional view of hypertext use on its head: the links are the content, the relationship is the semantics.

Building on my experience with Web Squirrel and with the design and implementation of the LMI Gateway hypertext project in the early 1980s, I'm designing a spatially-oriented development and user environment for the representation, exploration, and manipulation of multi-level and multi-dimensional information structures. These information structures may be modeled as graphs, sets, expert systems, constraint-based systems, etc. Their screen representation must correctly reflect the underlying knowledge model and yet be flexible enough to allow users to visually manipulate the space to reflect their purposes.

By "correctly reflect the underlying knowledge model" I mean that visual representations generated by the system from the model must not contain factual untruths, such as showing a person as author of a document that according to the knowledge model she did not author. Visual representations generated by the user in the process of exploration and manipulation, however, may contradict the underlying knowledge model. Reconciliation of these disparate versions occurs when the model is updated, and these changes are kept in a changes file until the model update occurs.

By 'user manipulation flexibility' I mean that, in addition to model correctness, the spatial representation must both reflect a serious concern with user-centered design tools and the ability to simultaneously display multiple points-of-view (POV) and levels-of-detail (LOD) with a visible language that clearly delineates the different domain dimensions.

In this position paper I will first describe three diverse driving applications, then discuss some of the requirements, characteristics, and issues an environment that implements these applications entails, and finally briefly describe a prototype environment currently under development called ConceptLab. I plan to demo this prototype with a small subset of the Memex Domain (described below) at the workshop.

My hope is that the combination of description and demo will spark both a discussion of the issues surrounding my proposal and interest in working on it, either with me or as part of an established group. This vision, which includes the "Beyond the Plane" [Simpson 1997] scenario, is very big indeed and I am under no illusion as to how much one 61-year old history major working part time can accomplish.

# Driving Applications

These applications introduce the kind of environment I'm proposing for the immediate future. In the long term, the online scenario "Beyond the Plane: Spatial Hypertext in a Virtual Reality World" [Simpson 1997] presents some possibilities in an office VR environment.

## Memex Domain

Visualization of the relationships among the people, institutions, publications, projects, and concepts reported on in the ACM Hypertext conferences from 1987 through 2001. Timestamps on all the objects allow simple time-based relationship representations.

## Concept EcoSystem

A time-interval-based world that represents the evolution of hypertext concepts through the years, showing the contextual modification of homonyms and synonyms, including how related concepts were spawned, evolved, modified, and perhaps died. The presentation is an interactive animation in which the user can select different POVs and subsections of the domain. The user may also alter the presentation both to make changes that reflect in the model and to explore, via 'what if' style questions, other scenarios, such as what would the world look like if a particular project or publication had never existed. This scenario would be quite useful for education, for cognitive science exploration, for related work exploration, for the selection of research topics in a PhD program, and for virtual museum exhibitions.

The original Gateway intent for this application was to provide a tool in which functional objects were alive in that they were able to (1) recognize context and changes in context, (2) to assimilate, accommodate to, and modify the context as do biological organisms, and (3) incorporate time intervals into their behavior. This would permit the development of the sorts of things we've seen in Vivarium and, more recently, in NearLife. I would like to be able to have such capabilities in ConceptLab to provide richer knowledge modeling and evolution tools. As with the "Beyond the Plane" scenario, knowledge here includes the full range of multi-modal experience as well as pre- and non-symbolic information.

The knowledge models that underlie the two scenarios are very similar, but goals, spatial hypermedia representation, and user interaction are all quite different.

## Cluster-based note taking

This application doesn't have a preexisting knowledge model but uses the workspace much as a sketchpad to capture ideas and notes in a free-form manner.

It differs from concept-mapping applications like Inspiration in that the user determines what the semantics of the elements and their relationships drawn on the screen, if any, are. A knowledge model may be created using the palette of structures and relationship specifiers available along with design tools, but the default is a simple spatial visualization with no underlying semantics.

It differs from pure drawing applications like Adobe Illustrator in that it is possible to attach a knowledge model (semantic structure) to the visual representation drawn on the screen. [Mayfield & Nicholas 1994] Its visual support is more like Gabriele Rico's spatial clustering [Rico 2000] than Tony Buzan's hierarchical mind-mapping [Buzan 1996], although a mind-mapping model could be attached to the visual representation.

The debuggers, browsers, inspectors, and profilers that are part of the Information Programming section described below would quite useful for such an emergent system. It could be used as a sketchpad or prototyping system for developing knowledge representations or information visualizations.

# Requirements, Characteristics, and Issues

The requirements presented by an interactive spatial hypertext environment for the representation, exploration, and manipulation of multi-level information structures such as that which supports the scenarios described above involve three major areas: **information programming**, which includes open systems access, programming environment, and structural computing issues; **knowledge representation**, including the dynamic evolution of knowledge models; and **spatial hypermedia user interfaces**. In addition, to implement the Concept EcoSystem application, **time-interval based modeling** must be considered.

# Requirements, Characteristics, and Issues for an Information Structures Spatial Hypermedia Environment

NOTE: this system is a well-defined application that exists within the larger world of other local applications as well as distributed domains such as the Internet. My discussion in this position paper deals almost entirely with the requirements, capabilities, and issues for the well-defined application component. The semi-permeable membrane aspect of the application arises in the open systems interface to other applications and in the import and export facilities.

## Information programming

### Link structure authoring system

By information programming, I mean the treatment of link structure development as a programming task with all that this implies. It means treating links and structures as first class objects with methods and attributes. Thus, a link or a structure may call another structure and use the return value, if any, from the call. It also means creating a full structure-programming environment with debuggers, inspectors, browsers, and profilers for links and structures. Like the LISP Machine environment of yesterday and the Squeak environment of today these tools will be written in the language of the rest of the system and will be accessible/modifiable from the system.

### Framework

Ultimately the system will be a framework, permitting the development of alternative relationship engines, knowledge models, and user interface designs. In order to create an effective and truly modular framework, it is necessary to create at least three preliminary applications in order to understand the abstractions which must be made. I am currently creating the first prototype application, which bears the name ConceptLab. The knowledge model is a simple general graph structure over the domain of the hypertext conferences from 1987 - 1996. The relationship engine is handled by simple string-matching queries over the graph. Most of the development effort in this first prototype is going into the spatial hypermedia user interface, which includes representations of multi-dimensional and multi-level logical groupings, or composites, as well as the representation of single elements and spatial groupings, called neighborhoods. In this first prototype there is no attempt to implement a system level spatial parser; the user defines neighborhoods by shift-clicking or lassoing and optionally names them. These neighborhoods can be retrieved as visual groupings or converted into composites.

### Open System Approach

By open systems approach I mean being able to access other systems on many levels ranging from file formats to application call and return. ConceptLab implements these principles in the following ways: First of all, on the level of import *and* export. Then, in the spatial hypermedia workspace environment, there are windows for external applications such as browsers, word processors, spreadsheets, and the like. Users can access these applications from within ConceptLab and drag-and-drop or cut-and-paste elements into the ConceptLab spatial hypermedia workspace.

### Persistent state

User-defined spatial hypermedia structures must be able to be saved as named components, which can then be incorporated into other structures. This applies to composites, neighborhoods, simple types, and workspaces, all of which are described in the Interface Objects section below.

### Issues

**Versioning, updating, and maintenance:** A critical component both of document and program development is versioning, updating, and maintenance facilities. In ConceptLab this applies to both the knowledge model and the spatial hypermedia workspace elements, such as complex shapes (templates). JavaDoc-style documentation tools that can be used to produced both online and paper reference, help, user guides, indexes, and case studies will be part of the maintenance and updating facilities. The first prototype will not have these facilities, but such facilities are certainly planned for the future.

**Separation of concerns:** When the links *are* the data, as in a semantic net or, more generally, in a knowledge model, what does it mean to say that links are separate from data?

**Structure modules:** What sort of return value would be returned by a link structure module? What is the benefit of such encapsulation?

# Knowledge Representation

In an article entitled "What is a Knowledge Representation?" published in AI Magazine in 1993 [Davis et al. 1993], Davis, Shrobe, and Szolovits describe five roles played by a knowledge representation. As members of the AI community, their perspective is targeted at those who would develop machine-based intelligent behavior. Nonetheless, there are a number of useful ideas for those of us who wish to use a knowledge representation as a relatively passive structural model component. In addition, more active versions, such as expert systems and intelligent agents, could ultimately play a role in applications such as the Concept EcoSystem.

Among the many issues involved in this area, three are fundamental to our context:

When to use passive structure-based models, such as semantic net, vs. active inferencing engines, such as Cyc [Lenat & Guha 1990].

How to integrate 3D semantic net visualizations, such as that shown by SemNet [Fairchild et al. 1988] into other spatial hypertext visualizations?

How, and if, to handle the separation of the knowledge models from the relationship engines.

# Interactive Spatial Hypermedia Interface

## Interface Objects

**Simple type:** domain-specific object. In the Memex Domain application, simple types are Person, Publication, Institution, Project, and Concept. It can be displayed as a simple node. as a named cluster of nodes consisting of its attribute values, or as a structured single node (icon) in which the visible components represent the attributes and are selectable for zooming and rollover text display.

**Composite:** model-based and/or user-specified collection that may include any combination of simple types, composites, neighborhoods, agents, and lists. Composites may be named or anonymous. User-specified collections may be stored as view objects separately from the knowledge model or may be integrated into the knowledge model.

**Neighborhood:** spatially-defined group that may be suggested by a system structure-parser or selected by the user. Neighborhoods may be named or anonymous and may be saved as either as part of a workspace or as separate objects. Neighborhoods are view objects and thus are not part of the knowledge model, but they may be converted into composites, which could then be integrated into the knowledge model.

**Workspace:** a Web Squirrel-style "infinite plane" window in which interface objects are displayed and manipulated. It is part of the system layout which is shown below in the Screen Layout section. Workspaces may be named and saved.

**Agent:** see description in the Behavior section below.

**List:** hierarchical collections of simple types, composites, neighborhoods, agents, and other lists.

**Shape template:** a shape, such as a spiral or a star, that is used to guide the display of interface objects. Shape templates may be rotated and scaled.

**Style sheet:** a specification for display characteristics such as shape templates, fonts, colors, and distance semantics that can be used to identify neighborhoods

**Metadata:** domain-specific information whose display is also domain-specific. Default display and manipulation of metadata occurs in the metadata/notes window shown below in the Screen Layout section.

# Requirements, Characteristics, and Issues for an Information Structures Spatial Hypermedia Environment

## Behavior

**Search and search-and-replace:** includes wildcards, full regular expressions, the ability to know where in an attribute the search string is (beginning/end), and the ability to interactively save-and-refine the results of a search. Searches can be constrained to spatially-defined parts of the workspace (named neighborhoods as well as temporarily identified spaces which are gerrymandered through shift-clicking), selection sets, attributes, or composites. A search can be named and combined with a structuring definition as a script. The results of a search can be used as a selection set specification for printing, exporting, subset saving, and other manipulations.

**Agents**: Web Squirrel-style processes that gather lists of elements matching user-defined Boolean expressions. The list generated by the agent changes dynamically to accommodate additions and deletions of workspace elements. The Boolean expressions are full regular expressions with wildcards. Agents can be constrained to spatially-defined parts of the workspace (named neighborhoods as well as temporarily identified spaces which are gerrymandered through shift-clicking), selection sets, attributes, or composites.

**Selection set specification and use:** specification includes shift-click and lasso in addition to the search, agents, shape templates, and import facilities described in this section. Selection sets may be named and saved.

**Shape template**: operations include pouring a selection set (which includes imported data) into shapes, translation and manipulation of the shapes, sorting shape elements with a resultant change in position, insertion, deletion, and moving individual elements. The contents of shape is by default a selection set and can participate in all selection set-based operations

**Selective focus & context tools:** multiple tools are needed to provide different types of focus & context manipulations that can be simultaneously in use. Examples include pruned and traditional fisheye views, translucent text expansion, translucent detail window overlay [Lieberman 1994], ZigZag [Nelson 2001], rollover text (usually used for metadata but potentially used for detail expansion), optionally persistent translucent marking menu techniques, selection-set constrained smooth and incremental zooming.

**Import and export:** formats that can be imported into the workspace can also be exported. At a minimum these include HTML files, tab-separated fields, and attribute-value pair tab-separated fields. Ultimately a very broad range of text and multimedia formats will be accommodated.

**Save, save as, and print:** these operations apply to the entire workspace, to attribute categories, or selection sets. For example, in the ConceptLab Memex Domain model one could craft a bibliography by creating a script that gathers, combines, and sorts Publication object attribute values and then saves and/or prints the result.

Other facilities include list-explode, agent-to-list, list creation, composite creation, neighborhood creation, and scatter/gather on many levels.

## Display Design Tools

Frank Shipman in his VKB paper [Shipman et al. 2001] comments on the problems that beset collaborative VIKI users, problems that arose from incompatible documentation styles and the subjective and changing nature of visible language evolution. I don't agree with the solutions adopted in VKB because I know from my experience with Web Squirrel how frustrating such shape and color straight-jackets are. The parallel to code documentation problems is obvious and software engineering tools and practices may provide ideas that could be adopted. Once again, viewing hypertext development as information programming provides access to software engineering insights and strategies that have been developed over the last 50 years.

The following list does not attempt to replace Adobe Illustrator but to provide a minimum set of design tools that are critical to effective spatial hypertext design.

**Guides:** include lines, shapes, grids, rulers, and numeric specification, with snap-to

**Drawing tools:** include freehand pens and brushes, shapes such as circles, rectangles, rounded rectangles, and polygons

**Shape alignment tools:** include a variety of element distribution mechanisms according to specified boundaries

**Coloring tools:** for coloring outlines, fills, backgrounds, and text according to a style sheet or specifications such as class or name or shape or shift-click/lasso

**Text:** format specifications applied according to a style sheet or specifications such as class or name or shape or shift-click/lasso. Text may also be flowed along shapes.

## Workspace Display Design Strategies

Control of the display of pre-existing knowledge models can be either specified as metadata attached to the model or can be controlled by user specification. Like Web Squirrel, the workspace is an infinite plane, which gives considerable scope to the user in designing effective structure representations. Some of the different display strategies include:

**Multiple Points-of-View (POV):** in which the same object, such as the VIKI project in the Memex Domain application, appears in several different locations and contexts on the screen. This can arise from the inclusion of the object as an attribute of different domain elements. In the Memex Domain application, the VIKI project, for example, would be referenced by both Catherine Marshall and Frank Shipman Person objects, by Xerox PARC Institution object, by several Publication objects, and by the spatial hypertext Concept object. In addition, the user can place multiple copies of the visual icon for an object on the workspace.

**Multiple Levels-of-Detail (LOD):** in which a composite object could display some of its components as a single icon and other components as a group of icons, while still other components use techniques such as translucent text displays, selective zooms, and pruned fisheyes Other uses of multiple LODs include using different parts of the workspace to manage different LODs on the same object. A paper-based book analogy is the appearance of the same phrase in overview and detailed TOCs, a chapter location, a page heading, and an index.

**Multiple dimensions:** different dimensions can be distinguished by color, shape, and size as well as by metadata displays in the metadata window, by rollover text, and by names.

## Screen Layout

**{ {**

Need a series of annotated screen shots of the ConceptLab layout showing the hot corners, external applications window, metadata and notes window, status window, and "infinite plane" workspace window with composites, structured icons, multiple POVs, multi-dimensions, multiple LOD with different simultaneous focus&context views such as rollovers, translucent text, spatially-constrained fisheye views, and spatially-constrained zooming.

**} }**

## Issues

**Spatial representation:** one of the problems with spatial representations that has been identified by many of the users I work with is that humans automatically attempt to attach meaning to all aspects of the spatial layout. This is a variant of the larger problem of how do you read a figure or a diagram, issues which Edward Tufte has addressed in his books [Tufte 1997] and can be seen in the difficulties found in effective image map design.

The problem becomes more severe, however, when the relationships combine logical and spatial representations and when the spatial relationships are loosely defined. When I demo Web Squirrel to people, I find that neighborhood names are critical to their ability to distinguish spatial from logical groupings and that even then they have problems if the groupings are at all complex.

Grouping ordered (alpha and numeric) material into spirals make sense and people seem like them a lot as a way of seeing and accessing a lot of ordered material in one glance. When I tried to use concentric rings, however, as a way of grouping related groups, I found that people had a lot of difficulty because they thought there were direct relationships along sectors of the concentric ring structure, and an implicit ordering between rings going from the inside to the outside. The topologically identical spiral and line work as a visualization strategy for ordered data, but concentric rings as a way of presenting loosely related clusters of information doesn't.

**Closed vs. open spatial representation:** what is the best way to combine different spatial representations. For some applications a space-filling form, such as is seen in a geometric quilt, conveys a sense of completion. The hypertext fiction community has used this form in the past in such works as Teri Hoskin's "Noon Quilt" [Hoskin 1999] and Deena Larsen's "Samplers" [Larsen 1998]. The disadvantage of this approach, however, arises from the positional constraints imposed by the 2D representation, even when the illusion of overlap is included. The more open and fluid form seen in mandalas, Kandinsky paintings, and David Laidlaw's scientific visualization work [Kirby et al. 1999] provides more opportunity for the simultaneous presentation of multiple points of view and levels of detail.

## Time interval representation, exploration, and manipulation

**{{** TBS **}}**

# ConceptLab Prototype

ConceptLab is spatial hypermedia environment for the representation, exploration, and manipulation of complex information spaces. The goal of the system is to provide an interactive tool for people to explore, manipulate, extend knowledge domains, where the term 'knowledge' can include multi-modal experience as well as pre- and non-symbolic information. For example, in a therapeutic use of a virtual reality-based version of ConceptLab, kinesthetic feedback loops might be incorporated into the knowledge model.

ConceptLab has the notion of a small knowledge base surrounded by a "semi-permeable membrane" through which access to the world of the Internet and other computing resources, both local and distributed, is possible.

## Implementation

The first prototype models the Memex Domain application with an annotated graph structure in a 2D desktop environment. I am programming in Squeak, Alan Kay's Smalltalk successor, because it provides a LISP-machine like programming environment, is cross-platform with the ability to run active documents across the Web, and is designed to allow novice programmers to create extensible multimedia applications. In addition, it provides dynamic binding of variables and the ability to convert objects to classes (templates) at runtime, both of which are critical to effective implementation of an interactive, dynamically evolving concept ecosystem.

The next implementation will take the Memex Domain into the cave and explore the issues involving multi-modal IVR use for the representation, exploration, and manipulation of complex information spaces.

Longer term plans include working towards the system I described as a fantasy in "Beyond the Plane: Spatial Hypertext in a Virtual Reality World", in which the interactive system takes the full human mind-body system into account in providing tools for exploring non-symbolic and pre-symbolic information in the user's head. This version could investigate tools for enhancing therapeutic strategies such as Gene Gendlin's Focusing [Gendlin 1981]and proprioceptive writing [Metcalf & Simon 2001]

## Potential Uses

**Cognitive science laboratory**

**Creativity support**

**Digital libraries**

**Education**

**Prototyping tool for information visualization and knowledge model development**

**Related work analysis**

**Research topic seeking**

**Therapy (Rorschach and TAT meet the computer...)**

## Issues not addressed

### Collaboration

While my focus has been on an individual's use of the environment, certainly there is nothing that precludes, in principle, this environment from being used collaboratively. Certain issues, such as versioning and namespace control, that are important for the single-user environment become critical in a collaborative environment. Other issues arise that are specific to the needs of collaboration. As I described above, Frank Shipman in his VKB paper [Shipman et al. 2001] comments on the problems that collaborative VIKI users had which influenced the design of VKB. I would like to find a way to keep the design richness that is a fundamental component of ConceptLab design while addressing those issues and I hope that software engineering may suggest strategies that will allow that. However, I don't yet have any alternative solutions to offer.

# Bibliography

[Buzan 1996] Tony Buzan. The Mind Map Book : How to Use Radiant Thinking to Maximize Your Brain's Untapped Potential. Plume, ISBN: 0452273226, 1996.

[Davis et al. 1993] Randall Davis, Howard Shrobe, and Peter Szolovits. "What is a Knowledge Representation?" in AI Magazine, 14(1):17-33, 1993.

[Fairchild et al. 1988] Kim M. Fairchild, Steven E. Poltrock, and George W. Furnas. SemNet: Three-Dimensional Representations of Large Knowledge Bases" in R. Guindon, editor, Cognitive Science and its Applications for Human-Computer Interaction, pages 201-233 Lawrence Erlbaum, 1988. URL: http://panda.iss.nus.sg:8000/kids/fair/webdocs/semnet/semnet-1.html.

[Gendlin 1981] Eugene Gendlin. Focusing. Bantam Books, 1981.

[Hoskin 1999] Noon Quilt. http://trace.ntu.ac.uk/quilt/info.htm, 1999.

[Ingalls et al. 1997] Dan Ingalls, Ted Kaehler, John Maloney, Scott Wallace, and Alan C. Kay. "Back to the Future: The Story of Squeak - A Usable Smalltalk Written in Itself" in Proceedings of ACM OOPSLA '97, 1997.

[Kirby et al. 1999] R. Michael Kirby, H. Marmanis, and David Laidlaw. "Visualizing Multivalued Data from 2D Incompressible Flows Using Concepts from Painting" in Proceedings of IEEE Visualization '99, 1999.

[Larsen 1998] Deena Larsen. Samplers. Eastgate Systems, URL: http://www.eastgate.com/catalog/Samplers.html, 1998.

[Lenat & Guha 1990] Douglas B. Lenat and R. V. Guha. Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project. Addison-Wesley Publishing Company, Inc. 1990. URL: http://www.cyc.com/

[Lieberman 1994] Henry Lieberman. Powers of Ten Thousand: Navigating In Large Information Spaces, Conference on User Interface Software Technology, Marina del Rey, California, November 1994. URL: http://lieber.www.media.mit.edu/people/lieber/Lieberary/Macroscope/Powers/Powers.html

[Marshall et al. 1994] Catherine C. Marshall, Frank M. Shipmann III, and James H. Coombs. "VIKI: Spatial Hypertext Supporting Emergent Structure", in Proceedings of ACM ECHT '94, 1994.

[Mayfield & Nicholas 1994] James Mayfield and Charles Nicholas. "SNITCH: Augmenting hypertext documents with a semantic net" in the International Journal of Intelligent and Cooperative Information Systems, 2(3), 1994.

[Metcalf & Simon 2001] Linda Trichter Metcalf & Tobin Simon. Proprioceptive Writing: Clearing the Mind workshops. URL: http://www.eomega.org/omega/workshops/8531ce50792aec241b4b5819f5e27d59/?GCS_Session=a77521c1ba95858859455045b622ac2e

[Nelson 2001] Theodor Holm Nelson. ZigZag. http://www.xanadu.net/zigzag/, 2001.

[Reinert et al. 1999] Olav Reinert, Dirk Bucka-Lassen, Claus Aagaard Pedersen and Peter J. Nürnberg. "CAOS: a collaborative and open spatial structure service component with incremental spatial parsing" in Proceedings of ACM Hypertext '99, 1999.

[Rico 2000] Gabriele Rico. Writing the Natural Way. J P Tarcher, ISBN: 0874779618, 2000.

[Shipman et al. 2001] Frank M. Shipman III, Haowei Hsieh, Preetam Maloor, J. Michael Moore. "The Visual Knowledge Builder: A Second Generation Spatial Hypertext" in Proceedings of ACM Hypertext 2001.

[Simpson et. al., 1986] Rosemary Simpson, John Mann, and JoAnn Brooks. "LMI-Gateway: A Syntactic Tool for Organizing Hetergeneous Data Sets." Submitted to AAAI '86., 1986.

[Simpson 1997] "Beyond the Plane: Spatial Hypertext in a Virtual Reality World", http://www.cs.brown.edu/stc/resea/telecollaboration/beyond_the_plane.html, 1997.

[Simpson 2001] Rosemary Michelle Simpson. "Experiences with Web Squirrel: My Life on the Information Farm" in Proceedings of ACM Hypertext 2001., Aarhus, Denmark, 2001.

[Tufte 1997] Edward R. Tufte. Visual Explanations: Images and Quantites, Evidence and Narrative. Graphics Press, 1997.