Back To The Information City

Mark Bernstein*
Eastgate Systems, Inc.
Watertown, MA, USA

Silas Hooper
Three Trees Labs, Inc.
Vancouver, BC, Canada

Mark W. R. Anderson
Web and Internet Science Group
University of Southampton
Southampton, Hants, United Kingdom

Abstract

In his 1994 doctoral dissertation, Andreas Dieberger proposed to visualize hypertexts as cityscapes. Revisiting this concept with the aid of modern processors and displays, we have found that the Information City provides a new perspective on hypertext visualization itself, a practice inspired both by structuralist and by existentialist thought. Where conventional spatial hypertext has tended to focus on proximity, the Information City often foregrounds the implicit semantics of the spaces between buildings.

CCS Concepts

• **Human-centered computing** → **Visualization theory, concepts and paradigms**; **Interaction design theory, concepts and paradigms**; • **Applied computing** → Architecture (buildings).

Keywords

spatial hypertext, information city, knowledge discovery, sensemaking, memory palaces, visualisation, information recall, VR

ACM Reference Format:
Mark Bernstein, Silas Hooper, and Mark W. R. Anderson. 2025. Back To The Information City. In *Proceedings of the 36th ACM Conference on Hypertext and Social Media (HT 2025), September 15–19, 2025, Chicago, IL, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3720553.3746664

1 Dieberger's City

In this paper, we revisit the Information City, an early proposal to visualize hypertexts as urban environments. When the idea was first advanced, neither processors nor computer screens were sufficient for more than a suggestive mockup. Today, a fresh effort may be of interest.

The late Austrian hypertext researcher Andreas Dieberger proposed the Information City in his doctoral dissertation [30] and elaborated that concept in several subsequent reports [29, 31, 32, 34, 35], culminating in 'A City Metaphor to Support Navigation in Complex Information Spaces' [33]. The core idea is that hypertext visualizations need not be limited to boxes and arrows, but that a hypertext might usefully be presented as a cityscape.

*Corresponding author.

HT 2025, Chicago, IL, USA

ACM ISBN 979-8-4007-1534-1/2025/09
https://doi.org/10.1145/3720553.3746664

> "The Information City is not a metaphor meant to be implemented straight from our specifications. Instead, it defines an ontology of spaces and connections that we think is useful for talking about systems of spatial metaphors and how they interrelate."

Dieberger's Information City makes an interesting contrast to SemNet [39, 40], a contemporary abstract information space examined by Boeing. Both projects were concerned with navigation and the challenge of finding a desired resource amidst many other resources. Three-dimensional spaces offer attractive affordances; in particular, all graphs can be embedded in three dimensions without edge intersections whilst comparatively few *planar* graphs can be thus embedded in two dimensions. The intimidating tangle of dense link networks was clearly a concern of both research groups. Both groups offered interesting proposals for the best way to get around their visualizations; SemNet thought helicopter controls might be most easily mastered, while Dieberger contemplated buses and subways. The speculative KNOESPHERE proposal envisioned a spaceship equipped with a pilot and tour guide [66].

While Dieberger's Information City [31] has the form of a city, it need not be a plausible or real urban environment. In fact, cityscapes in games often omit things we would miss were the cities actually inhabited. There is no need for sewers or ventilation [94, ch. 1]. A fifteen minute walk from our house to the grocer might be acceptable in real life, but in a game it would grow boring [22]. Many factors make a built environment sympathetic to the people who live in it [1], but in the Information City we do not need to worry whether the roof is easy to repair or if the teenagers have sufficient space.

At the second European Hypertext Conference (ECHT'92), Jay David Bolter suggested that the *memory palace* of the rhetoricians might make a suitable model for embedding a hypertext in virtual reality [19]. The memory palace is an aid to remembering the topics to be covered in a long speech; one associates each topic with a familiar place, and can reconstruct the desired sequence by imagining oneself walking through that place. The place need not be a literal palace: the traditional seder plate serves the same function, and we can as easily envision an art gallery or a streetscape.

A key, in either case, is management of visibility: if we stand *here*, what do we see as we look around? [105] This would typically include the adjacent buildings, but might also include distant landmarks [69]. Gateways can call attention to a district while blocking the view of its interior, while gaps between buildings can reveal things that might otherwise be hidden. It is desirable to break up long vistas both for utility and for performance [22].

2 Visualizing Hypertext

Hypertext systems have always featured visualizations of their content. Mockup visualizations are prominent in Nelson's early visions [87, 88], and the first review of hypertext systems was inclined to consider visualization a necessary component of a hypertext [24]. Among the earliest systems, visualizations were a prominent feature of NoteCards [52], Intermedia [119], Storyspace [20], SEPIA[107], gIBIS [25], and Document Examiner [117]. Though somewhat abated, interest in novel visualizations has never vanished [14, 53, 61, 92].

Hypertext visualizations might be seen as merely convenient, a facet of human-computer interaction that could conceivably lead to modest productivity gains [89, ch. 9]. They may be better understood, however, as the contested ground between two distinct intellectual reactions to the Second World War.

2.1 Hypertext and Structuralism

The widespread intuition that hypertexts can and should be visualized is entangled with *structuralism*, an intellectual movement that stems from the initial wartime encounter (in a seminar on *cybernétique* at the *École Libre des Hautes Études*, a New School annex for refugee francophone scholars) between the *sciences humaines* and the notion of Hilbert spaces implicit in nascent cybernetics [81, pp. 204-209] [48]. The lecturer, Roman Jakobsen, was chiefly interested in *phonemes*, the myriad atomic sounds that make up human languages. These differ endlessly among languages and dialects, but Jakobsen saw that they could all be characterized by placing them in a multi-dimensional space with axes like voiced/unvoiced, clipped/continuous, and so forth. This framework offered a promising route to systematize phonetics; there are lots of sounds but only so many degrees of freedom in the vocal tract. A sociologist, Claude Lévi-Strauss, proposed a corresponding approach to systematize kinship relations: every culture has its unique approach to kinship, but perhaps these too could be indexed systematically through a similarly finite number of axes. Lévi-Strauss would proceed to build the Human Relations Area Files, a huge storehouse of systematized microfiche data on human cultures that recalls both Bush's Memex [110]) and Luhmann's *zettelkasten* [68].

Almost all hypertext researchers have been favourably inclined to *structuralism*. The core intuition of structuralism is that information — *interesting* information at any rate — has informative structure. This was no longer given (or ordained) in post-Enlightenment thought, and both Modernism and Expressionism were open to the idea that no structure could be found beneath the empty sky, save for whatever conventional structure civilized man could wrest from

> The secrets of the hoarie deep, a dark
> Illimitable Ocean without bound,
> Without dimension, where length, breadth, and highth,
> And time and place are lost; where eldest Night
> And Chaos, Ancestors of Nature, hold
> Eternal Anarchie, amidst the noise
> Of endless Warrs, and by confusion stand.
> (Paradise Lost II 891-7)

The development of the natural sciences in the 18th and 19th centuries had suggested that things really did have order and structure. Atoms combined in molecules, and those molecules formed crystals

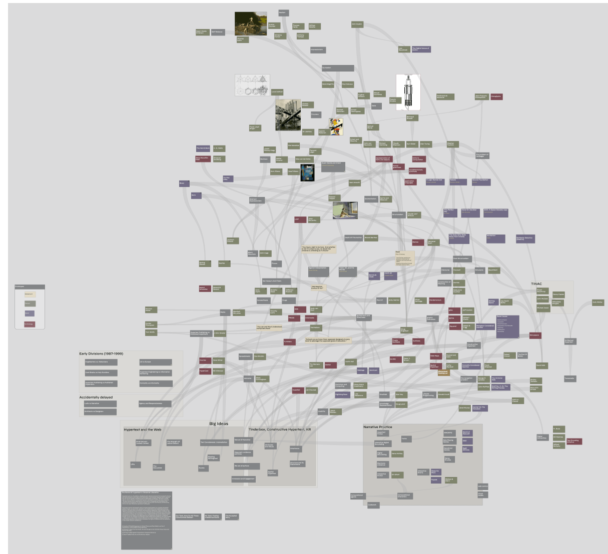

Figure 1: A conventional map view in Tinderbox [10], visualizing notes on art and hypertext theory.

through orderly processes susceptible to discovery and description. Mathematics itself could be shown to be consistent and correct (although not complete) without reference to revelation [21, 71].

One key to effective elucidation of such structures had been known since Leibniz: devise an effective notation after defining a suitable frame of reference from which to view the question [55, 62]. Several early hypertext projects were explicitly structuralist in ambition, notably gIBIS [25] and SEPIA [107], systems for exploring diagrammatic representation of argumentation [122]. Even researchers who had been convinced by the postmodern problematization of Lévi-Strauss's orthogonal axes (or "binary oppositions") believed that there existed *some* structure to be teased out and understood [20, 65]. Tensions in the literature reflect differing interpretations of existing textual practices and their place in hypertext. Intermedia viewed links as fundamentally annotative and hand-crafted approximately-hierarchical overviews of topical locales as a reaction against the perceived chaos of automatically-generated maps (Figure 1) [119], while the TINAC collective preferred a profusion of cycles and contours mediated by dynamic link behaviour and insisted that recurrence — the revisiting of an already-seen writing space that had been dismissed as an inefficient impediment to hypertext adoption [28] — was not a defect, but rather a way to comprehend structure [15].

2.2 Hypertext and Existentialism

Structuralism sought to find order beneath an empty sky amidst the noise of endless war, and found a partial answer in the binary logic of n-dimensional vector spaces. Its visualizations aspire to be proofs, objective demonstrations that an argument is sound or specious [25, 122].

Existentialism, on the other hand, is always suspicious of certainty. The central problem that existentialism confronts is personal choice, such as deciding whether one should flee Vichy to join the

Free French or instead remain and fight with the Resistance [81, pp. 63-69]. Existentialism holds that this choice must be free and authentic, that it should express individual character. Analysis cannot always help because there may be no right answer: if there were, the choice would not be free. The core problem is to choose *for oneself* rather than conforming in *mauvaise foi* to what your parents, preacher, or nation expect. There is no exit from this demand, and therefore we must raise our consciousness and liberate ourselves from convention.

Ted Nelson used the language of existentialism to surprise, shock, and proselytize in terms of *Computer Lib* and *Literary Machines*. Again and again, Nelson returns to the plight of the student forced to study what bores her because her teacher insists she ought [88, pp. DM16-DM19]. Though his rhetoric was less exuberant, Engelbart's full-throated embrace of existentialist thought is clear in the dynamics of his group and its convolution with EST [8, 115][7, pp.201–208].

Existentialism in hypertext is perhaps most evident in *spatial hypertext*, an approach that tries to free hypertext from the constraints of formal systems [5]. The crucial jump, in retrospect, was Aquanet [73], "a system to hold your knowledge in place." Aquanet attempted to look at 'knowledge structuring tasks' using spatial layout to explore and impart semantic meaning to emergent knowledge, and permitted users to define structured type such as multi-headed links. But study of AQUANET in use [74] confirmed earlier observations that choosing types could become an impediment to getting work done. Moreover, types are hardest to choose in precisely the situations in which they could do the most good, when the problem domain is not well understood [102]. Spatial hypertext tries to express hypertextual connections without explicit links, and often uses a spatial parser to help identify possible structure in analysed documents [6, 76, 101]. Meaning emerges from proximity and geometry [11], but crucially, for spatial hypertext meaning resides in the mind of the user. We are not proving theorems or conducting jovial debates of the curious ways people resist innovation; the spatial hypertext user is trying to figure things out for themselves.

3 Diagramming and Finding Texts

Recent years have seen great interest in the study of mechanisms for representing textual relations. Marginal annotations to call attention to (or cast doubt upon) a passage are an ancient practice, standardized and institutionalized by Zenodotus, the first Librarian of Alexandria [97]. The scholarly footnote[1] itself dates to Gibbon [49] and to the practice of Leopold von Ranke, the founder of archival history [51]. Notably, the footnote does not merely decorate the main historical narrative: it creates a second narrative that explains how we know that the main narrative is correct. This creates a complexly hypertextual rhythm, even on paper. Studies of other kinds of marginalia have yielded fascinating insights into the way books are used [27, 38, 56, 72, 93].

Grafton also studied the timeline, which originated as a graphic approach to representing events that occurred simultaneously in different places [98]. The use of maps and trees as interpretive aids for texts was pioneered by Moretti [85], and by Turchi [114]. Lima provides a lavish compendium of tree visualizations [67]. Prescriptive approaches to information graphics have been elucidated by Tufte [112, 113] and comprehensive descriptive approaches, in art-book format, by Rendgen et al. [96].

An abundant literature studies way in which people have sought to discover and express structural relations within and among documents. It was clear by the 13th century, and likely much earlier, that there was simply too much to know [18] and that scholarship would in the future depend on extrinsic or artificial mechanisms for finding and understanding text.

> "We have reason to fear that the multitude of books which grows every day in a prodigious fashion will make the following centuries fall into a state as barbarous as that of the centuries that followed the fall of the Roman Empire. Unless we try to prevent this danger by separating those books which we must throw out or leave in oblivion from those which one should save and within the latter between what is useful and what is not. (Adrien Baillet, 1685) (from [17, p.59])

Judith Flanders examined the development of alphabetical order [42] and Kevin Jackson studied dedications, blurbs and other paratexts [57], while Wilson-Lee explored the origins of library catalogs in the 16th century [118].

Recent years have seen fascinating growth in scholarship and understanding of various forms of notebook. The origins of writing in the West (it originated separately in China and in the Americas) derives from the desire to keep notes about transactions [95, 100]. People in antiquity made notes on wax tablets, on folded vellum, or in papyrus codices; these were not entirely satisfactory because wax and vellum were susceptible to forgery while papyrus was fragile. Only in the mid 13th century did paper notebooks become available in the Mediterranean; because ink soaks into paper it resists erasure, while the sewn signatures of a bound book make replacement of entire pages visible [2]. Notebooks were to become tremendously important and powerful tools for the laboratory and the field scientist, for the artist and for the Romantic poet abroad in Nature [36, 54].

We are particularly interested in *constructive hypertext* [59]), which is to say in the use of hypertext in the process of composition, whether the end-product is a hypertext, a paper like this one, or a book. Indeed, there may be no end-product at all; we may abandon the book or we may take research notes to satisfy our own curiosity or to improve our own understanding. It may frequently be the case that we ourselves do not know the ultimate uses of our notes, yet still find note-taking rewarding.

4 A Use Case

The specific use case that concerns us here is collecting and using research notes in connection with crafting a long piece of non-fiction. Recent observations on this task often foreground the discovery or utilization of implicit structure [79] [23].

The most familiar family of hypertext visualizations displays each hypertext node as a labelled geometric shape (Figure 1), representing links as labelled arcs among them. Labels are a problem. In English (and many other scripts), labels are predominantly horizontal, which constrains the shape and spacing of the nodes. The

[1] The print footnote, as such, dates earlier to London, c.1658 [121, Ch.2].

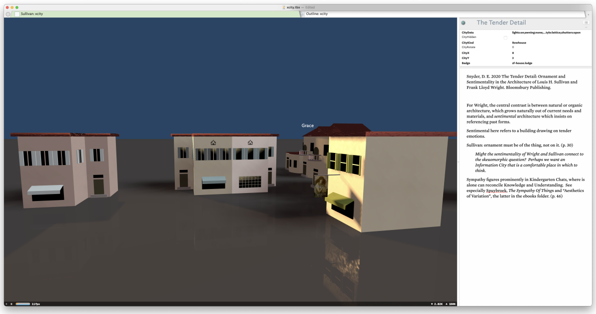

Figure 2: A view of some notes about Ruskin's theory of the Gothic virtues in architecture and urban design.

texts add visual noise, and rendering them presents a performance bottleneck. In practice, it seems difficult to fit more than a hundred or so notes onto a screen at any one time. Formerly, we could allude to the eventual production of better screens, but today's screens are already approaching practical limits of human vision. Scrolling, segmentation, and elision [12, 20] all help, but only take us so far. Non-Euclidean geometries are intriguing [77], but these, too, have practical limits.

In plane geometry, each note tends to have about 6 ± 2 proximate neighbours [14]. This, too, imposes representational constraints. In the pioneering era, these limitations were not experienced as shortcomings of visualization because hypertexts were small. *afternoon, a story* [58], had 562 lexia, *Victory Garden* [86], considered to be enormous when it appeared, had 972. Yet a thousand research notes seem quite modest for even a short book. A thousand notes in a Tinderbox document was prodigious as late as Tinderbox 6 (2014 [4]); today, note collections five times as large are common.

Might we find new visualizations that do not depend on omnipresent textual labels yet prove useful for writing?

5 Architectonic Spaces and Visual Composition

In an Information City, hypertext nodes are represented as buildings, or perhaps parts of buildings. We walk (or fly) among them. Selecting a building lets us examine and edit the node's title, text, and attributes. Whilst it is easy to move vertically in the Information City, we generally prefer to remain at ground level and use the vertical dimension chiefly for fast movement across long distances. A bird's eye view, just above the rooftops, is often useful, and resonates with model railroads (which also hold hypertextual interest[2] [64]).

5.1 What Is A Building?

The Information City represents hypertext nodes as structures. The choice of atomic units is not uninteresting. For example, the bibliography of a technical paper is a list of some dozens of references; might we represent them as an office building — a structure designed to replicate a generic "office" as many times as required [108]?

[2] https://victorianweb.org/cv/models/apg/history.html

We adopted the principle that a node ought to be visible from any direction, though of course other buildings might obstruct your view. Thus, our multi-story structures are like Tokyo's *zakkyo* buildings, tall but narrow commercial buildings built on very small lots in which each tenant occupies an entire floor [3, pp. 60-98]. This provides a convenient set of affordances; for example, the roof slab for each story is black if that unit has no title or text, and is otherwise red. Alternatively, we might show interior lights in "occupied units" and leave vacant units dark.

Creation of empty hypertext nodes has not been much (if ever) discussed in the literature, but it makes sense both in terms of the city metaphor and in terms of constructive hypertext [59]. As we sit down in a library to take some notes, we often have a sense of how many notes we expect. When we created a container for References for this paper, we knew roughly how many references a paper like this typically has. Why not make them right away, and use them as needed? Joyce's description of constructive hypertexts as "versions of what they are becoming, a structure for what does not yet exist" can seem gnomic, but what is more direct than making space for the notes, sections, or references that you are bound to need?

In retrospect, we might relax the requirement that every node be visible from every direction. While the visibility principle is a useful guide, it may unnecessarily constrain the representation of subordinate nodes or components. These might be represented as accessory units — shops on the ground floor of residences, street vendors, basement flats. Though not as visible as we would like, these remain findable to the extent that the structure in (or on) which they reside is findable. The contribution of such mixed uses has become a central concern of contemporary urban design [3, 105]. We would also like to explore the potential of markings on walls that add meaning, such as graffiti or tramp signs, or even markers of age, read-wear or modification such as the growth of ivy. The appearance of worn pavement would be a natural signal for trails [44].

5.2 Montage

In the Information City, the spaces between buildings take on an intriguing valence. We seldom see buildings in isolation. Even when standing so close to a structure that it fills our field of view, we recall what we saw as we approached. If I am standing in a plaza near buildings labeled as "Beowulf", "Caedmon", and "The Ruins," this empty space comes to mean "Old English verse" without formal designation. This effect has recently become an important consideration in city planning [46, 105], though we did not know this at first and identified it with the implicit semantics of spatial hypertext. Significantly, though, most of the spatial hypertext literature considers the relations among nodes, where here we experience semantic effects of the spaces *between* nodes. The effect is related to cinematic montage, in which a cut between two shots gives rise to a third idea, one not contained in either shot [37].

The Information City can seem like the weird twin of locative hypertext. Where locative hypertext starts out from meaningful places and adds layers of location-based narrative, buildings in the Information City acquire meaning from their contents and their neighbours. Route planning matters a lot in locative hypertext

[82, 83], but distance matters less in the Information City since we need not climb hills or worry that the parking lot will close.

We assemble our Information City, as we usually assemble conventional hypertext maps, node by node, placing each new item in whatever locations seems convenient or evocative. If our initial choice turns out to be wrong, buildings can be dragged to a new place. Yet we don't do this as often as we had expected; neighbourhoods tended to retain much of their layout and character. This echoes the experience of real cities, in which the forces of lot boundaries and real-estate prices provide additional conservative force.

The effect of *montage* raises the question of whether a *Planner* algorithm might be devised to accept an existing collection of notes and to lay out a useful Information City. A Planner might also examine newly-added or revised notes and suggest alternative placements or propose relevant colour schemes or architectural details [75]. This incremental planner could benefit not only from examining the contents and metadata of new notes but also from reviewing the evolution of neighbourhoods over time, much as the semantics of spatial hypertexts may sometimes be induced by observing the objects' histories [99].

6 Symbolic or Skeuomorphic?

Buildings in the Information City differ in many respects from real buildings. In the Information City, if a building is in the wrong place, we can drag it elsewhere. We seldom move buildings in real life. We can add details to differentiate one building from another—door treatments, shutters, paint, pavement fixtures—with a simple choice from a contextual menu. This raises the question of whether the components of the Information City are useful as symbols, or merely skeuomorphic illusions. Might we be as well off with SemNet's labeled boxes floating in space?[40]

Our initial experience suggests that the utility of the Information City resides significantly in management of visibility: what can be easily seen, what can be seen in the distance, or even things that might not be visible themselves but that we know are near a prominent landmark that we can see [69]. The spatial reasoning involved here is intuitive, and though few people design cityscapes, the same issues arise when planning a kitchen or organizing a closet. One advantage of the urban metaphor in this case is that we can anticipate places that are likely to be viewpoints, such as gateways, plazas, and intersections, and pay particular attention to the view from these locations.

Though it might be tempting to organize one's Information City in a regular grid of broad avenues, narrow and irregular streets and small plazas offer better opportunities to control sight lines. Not every note ought to be equally visible at all times; when we stand in the plaza of Old English Verse, we hope to see the notes we expect and also some other notes that might be pertinent. By avoiding long linear streets, we can also avoid the overhead of rendering distant structures that are less likely to inspire interesting ideas.

Our Information City is not populated. Depopulated virtual spaces are common, but always uncanny [84].

When we query hypertexts, we typically ask the system to list isolated notes that contain a text pattern, possess a tag, or have metadata values in a range that interests us. Query in the Information City could foreground results by rendering structures that do not satisfy the query as translucent. Stored queries might work like Lenat's goggles, allowing the user to see things normally hidden [66]. A different sort of query problem also arises from the city itself: we might want to relocate a particular neighbourhood from its visualization whilst not recalling the identify of any of the constituent building. A user might tell us, "I remember there was an intersection with a fire station on one corner and a restaurant with pink shutters on the other." This was once a problem familiar to hypertext research [50], and though we seldom talk about link-network queries today, perhaps we ought.

6.1 Beyond Realism

The elements of our Information City look roughly like actual buildings with doors, windows, lintels, and corbels. As such, they can acquire some emotional resonance with actual buildings that happen to be familiar to us, or with our imaged experience of inhabiting an actual space like the space that the Information City depicts. This cannot be entirely avoided: a space defined by Platonic solids would exert *some* emotional response.

Both Dieberger's proposal and our implementation assume a fairly literal rendering of building models, but this is by no means necessary. We might easily, for example, use vertex shaders to replace the photographic rendering with a pointillist rendition in the manner of Signac's *Châteaux des Papes*, though this raises interesting questions of cinematic continuity as we move through the view. Similarly, softening and smudging of edges (as in Pierre-Auguste Renoir's views of Cagnes), or atomospheric effects like fog or rain (c.f. Caillebotte's *Rue de Paris, temps de pluie*), are no more arbitrary than the photographic rendering we happened to select [13].

6.2 Software aesthetics

The elements of the Information City represent buildings, and so invite an emotional or aesthetic response. There is, of course, a long history of representing abstract or imaginary but emotionally-intense structures in painting: Piranesi, Palladio, Hogarth, Sheeler, and Hopper are some obvious examples. The craft of software has tended to relegate aesthetics to the realm of marketing or to customer whim (though see [41]). Contemporary architects, on the other hand, have embraced the complexity of our aesthetic response to structures [106, 116].

Users are influenced by the aesthetics (or visual appeal) of a user interface, even when expressly trying to evaluate its usability or productivity [63]. This has been corroborated in different cultural contexts [111]. Norman explored the vital role of aesthetics and emotional responses to design in how users experience and evaluate products [91], while Nielsen's heuristics for UI design include a call for aesthetic and minimalist design [90].

Traditional evaluation models, based on predefined linear paths or usability metrics, can be challenged in the context of spatial hypertext. Evaluation typically focuses instead on how users make sense of information, the ways in which spatial relationships support or hinder sensemaking, and how interfaces scaffold emergent organisation. For example, see Shipman, Marshall and Moran's treatment of implicit structure in spatial information layouts [104]. Blackwell argues that metaphors in UI design let us make abstract concepts and metaphors more concrete (reification) [16]. In spatial hypertext, metaphor often underpins the aesthetic logic of the space: space as a metaphor, leading to cognitive immediacy.

Reified metaphors are not merely functional, but have aesthetic properties: for example, a tool reifiying the ‘desktop’ has neatness, icons and containment (order), while a spatial hypertext reifying ‘sketching’ might emphasise fluidity, incompleteness, and messiness (emergence, play).

Users may judge spatial hypertexts not only on performance but also on how well the metaphorical aesthetic aligns with their approach (organising vs exploring, brainstorming vs structuring). The Information City might wear different clothes and so offer distinct experiences; our course syllabus might perhaps be a cheerful village green or the midway of a World’s Fair, while our overdue book project might be a Piranesian gloom.

7 Implementation Notes

We implemented a prototype information city in SceneKit as a Tinderbox view [14]. This required roughly 3500 lines of C++ and Objective C, written by one developer over a span of about four months.

7.1 Performance

The implementor expected that performance would be a constant impediment to this project, as it had been to Dieberger. This was mistaken, and though our GPU reports it is working fairly hard, at no point in the project did an actual performance issue arise. We seem able to accommodate hundreds of simple stuctures without much trouble, and so are at least roughly comparable in scope to conventional hypertext maps.

7.2 Labels

One reason we decided to revisit the Information City was to explore hypertext visualizations that are not constrained by their need to accommodate hundreds or thousands of textual labels. Few of our buildings have textual signage. Instead, following Zellweger’s FLUID annotations [120], we display a transient label on mouseover that drifts upward into the sky.

Despite the lack of labels, the most-used buildings soon became familiar. This process was assisted by supporting numerous small variations in the colour and appearance of the buildings. By changing window shape and trim, door transoms, shutters, accent colors, and lighting, we can distinguish noteworthy structures while retaining the same underlying repertoire of models.

7.3 Links

Though unimplemented at this time, our intent is to represent links as (abstract) birds that fly in a looping trajectory from source to destination. This keeps the links in the sky, which otherwise lacks information density [112, 113]. This scheme has the useful effect of identifying nodes with many links and should make it easy to distinguish those with high in-degree from those with high out-degree.

7.4 Navigation

Dieberger paid substantial attention to navigational facilities such as tour buses, subways or portals. These can be useful, but being able to move vertically is efficient [43]. As SemNet discovered, we do not need plausible flight mechanics [40]. Simply getting up above the rooftops facilitates rapid travel over long distances, though for the modest city involved in writing this paper, walking was usually preferable.

It is less obvious, at least to us, what navigational controls are most useful. For example, if pressing the W key moves you “forward”, should you advance toward toward the selected object or toward the direction in which you are looking? Each approach can be useful.

Faulty navigation makes things terrible: the system won’t let you do what you want to do, and the effect of better imagery is to make this more, not less, frustrating. For several weeks during development, navigation was possible but confusing, because getting the correct frames of reference for camera rotation was unexpectedly tricky. Navigation is hard to implement using test-driven development [9] because you can easily find yourself with a suite of unit tests that confirm that you are, indeed, doing what you intended, when you intended to do the wrong thing.

8 Toward Evaluation

The success metric for our proposed use case is, clearly, the successful and timely completion of a good book. To measure this in terms of milliMasterpieces per fortnight, though, is likely infeasible.

Tenen [109]) considers this question in one of the better recent considerations of literary machines and their role in writing. He concludes that computer assistance cannot play a role in works of genius, but that they might be very useful for more modest “hack” work, such as pulp fiction generated by *Plotto* [26]. This Romantic division of literary labor strikes us as ill-advised, especially since a number of works and authors keep sliding across the proposed gulf. Attempts to enforce a similar dichotomy on cinema collapsed half a century ago in the face of postmodern criticism [60, 80], and we have no reason to expect that rebuilding the wall between high and low brow would be more successful in computation.

Any effort like the Information City must take into account the Hawthorne Effect: changes in the work environment frequently improve productivity simply because the environment changes and the subjects know they are subjects. But this reservation, too, cuts both ways: if a new view is not in fact objectively superior to our customary visualization but, by engaging us through its novelty, leads to improved transient performance, that performance is not an illusion: the book we wrote is real. As Emerson wrote in his notebook (Journal. 19 April 1835), “It is a happy talent to know how to play.”

8.1 Our Doubts

Our experience of the Information City is mixed. As we planned this paper, we sometimes found ourselves reaching for other, more familiar tools.

It seems clear that the Information City is better suited to constructive than to exploratory hypertext. It is now almost five months from our starting point, and that is sufficient time to learn that the

To Do List is the green bow-fronted house in the centre of town, and that the latest Release Note is the top story of a tall (and growing) tower on the outskirts. This awareness grows naturally in a personal space, but the label-free design makes it hard for collaborators to grasp.

On the other hand, the opaqueness of the label-free design does offer some interesting affordances for artifactual hypertexts [45, 78], in which the reader moves freely within a narrative-rich environment. Here, the story is shaped not by conventional narratology [47] but through directing the viewer's attention through architectonic space. The analogy to immersive theater [70, pp. 41-43,170-172] is intriguing.

8.2 And Yet...

Despite our doubts, in working on this paper we have often stood in front of the hot dog stand that holds our reading list and looked across the way to the Ruskin tavern which holds notes on architectural aesthetics. Further down the narrow, winding street are notes on "changefulness," "sympathy," and "grace." A tall tower of Louis Sullivan buildings is just offscreen to our right, but we know that it is right there if we want it.

The Information City can hold roughly as much as the conventional map view; the sky is at present not used efficiently but provides a way to separate links from nodes that has eluded familiar visualizations. Reading someone else's Information City is difficult, but this is also the experience of all complex hypertext maps.

Map views in spatial hypertext systems tend to be concerned with composites and clusters. An interesting aspect of the Information City is that the City is more interested in "What can I see from here?" than "What is this part of?" The Information City lends itself to exploring what else might be useful to think about, when you are thinking about a note.

Though we are sympathetic to the structuralist conviction that information has meaningful structure, we do not see a clear prospect that Dieberger's Information City can be proven to yield more insights per hour than other existing or conceivable visualizations. We are also sympathetic to the existentialist claim that what really matters in a hypertext visualization takes place in an individual mind, not on the screen or the page. The Information City may be superb for some and intolerable for those who might prefer to work in a Piranesi dungeon.

Acknowledgments

We are grateful to the Frances Loeb and Widener Libraries of Harvard University. Kimera Royal provided valuable implementation advice. Gina Mosca, Dan Kearney, Lars Spuybroek, and John Vinci all provided helpful advice on aspects of architecture and urban design.

References

[1] Christopher Alexander. 1978. *A Pattern Language: Towns, Buildings, Construction.* OUP USA, New York, NY, USA. 1216 pages. https://global.oup.com/academic/product/a-pattern-language-9780195019193
[2] Roland Allen. 2023. *The Notebook: A History of Thinking on Paper.* Profile Books, London, UK. 342 pages. https://profilebooks.com/work/the-notebook/
[3] Jorge Almazán. 2022. *Emergent Tokyo. Patterns of Spontaneous Micro-urbanism.* Oro Editions, Toronto, OT, CA. 250 pages.
[4] Mark W. R. Anderson. 2025. *Previous Versions To Current Baseline.* atrbref.org. Retrieved April 11, 2025 from https://www.acrobatfaq.com/atbref10/index/Previous_Versions_To_Current_Baseline.html
[5] Mark W. R. Anderson. 2025. W(h)ither Spatial Hypertext?. In *Proceedings of the 36th ACM Conference on Hypertext and Social Media (HT '25)* (Chicago, IL, USA). ACM, New York, NY, USA. doi:10.1145/3720553.3746683
[6] Claus Atzenbeck and Peter J. Nürnberg. 2005. Constraints in Spatial Structures. In *Proceedings of the Sixteenth ACM Conference on Hypertext and Hypermedia (HTYPERTEXT '05)* (Salzburg, Austria). ACM, New York, NY, USA, 63–65. doi:10.1145/1083356.1083368
[7] Thierry Bardini. 2000. *Bootstrapping: Douglas Engelbart, Coevolution, and the Origins of Personal Computing.* Stanford University Press, Stanford, CA. 320 pages. https://www.sup.org/books/title/?id=308
[8] Thierry Bardini and Friedwald. 2001. Chronicle of the Death of a Laboratory: Douglas Engelbart and the Failure of the Knowledge Workshop History of Technology. In *23*. Bloombury Press, London, UK, 191–213. doi:10.5040/9781350018983.0014
[9] Kent Beck. 2003. *Test-Driven Development: By Example.* Addison-Wesley Professional, Boston, MA, USA. 241 pages. https://archive.org/details/est-driven-development-by-example/test-driven-development-by-example/
[10] Mark Bernstein. 2002. *Tinderbox.* Eastgate Systems. Retrieved 2019-08-01 from http://www.eastgate.com/Tinderbox/
[11] Mark Bernstein. 2011. Can We Talk About Spatial Hypertext?. In *Proceedings of the 22nd ACM Conference on Hypertext and Hypermedia* (Eindhoven, NL). ACM, New York, NY, USA, 103–112. doi:10.1145/1995966.1995983
[12] Mark Bernstein. 2016. Storyspace 3. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media* (Halifax, NS). ACM, New York, NY, USA, 201–206. doi:10.1145/2914586.2914624
[13] Mark Bernstein. 2023. The Indefinite Idea Plane Artistically Considered. *International Journal of Design Sciences and Technology* 25, 2 (2023), 1–10.
[14] Mark Bernstein. 2024. A New View. In *Proceedings of the 35th ACM Conference on Hypertext and Social Media* (Poznań, Poland). ACM, New York, NY, USA, 65–69. doi:10.1145/3648188.3678214
[15] Mark Bernstein, Michael Joyce, and David Levine. 1992. Contours of Constructive Hypertexts. In *Proceedings of the ACM Conference on Hypertext* (Milan, IT). ACM, New York, NY, USA, 161–170. doi:10.1145/168466.168517
[16] Alan F. Blackwell. 2006. The reification of metaphor as a design tool. *ACM Trans. Comput.-Hum. Interact.* 13, 4 (2006), 490–530. doi:10.1145/1188816.1188820
[17] Ann Blair. 2003. Reading Strategies for Coping with Information Overload, ca.1550-1700. *Journal of the History of Ideas* 64, 1 (2003), 11–28. doi:10.1109/MC.1987.1663693
[18] Ann M. Blair. 2010. *Too Much to Know.* Yale University Press, New Haven, CT, US. 416 pages. https://yalebooks.yale.edu/book/9780300165395/too-much-to-know/
[19] Jay David Bolter. 1992. Virtual Reality and the Future of Hypertext (Keynote Abstract). In *Proceedings of the ACM Conference on Hypertext* (Milan, IT). Association for Computing Machinery, New York, NY, USA, 2–2. doi:10.1145/168466.168473
[20] Jay David Bolter and Michael Joyce. 1987. Hypertext and Creative Writing. In *Proceedings of the ACM Conference on Hypertext* (Chapel Hill, NC, USA). ACM, New York, NY, USA, 41–50. doi:10.1145/317426.317431
[21] Stephen Budiansky. 2021. *Journey to the Edge of Reason: The Life of Kurt Gödel.* Oxford University Press, Oxford, UK.
[22] Anna Bulavina. 2024. *Urban Planning in Games.* Games Developers Conference. Retrieved April 12, 2025 from https://www.youtube.com/watch?v=Q3Scw3dzxW
[23] Robert A. Caro. 2019. *Working.* Knopf Doubleday, New York, NY, USA. 241 pages. https://www.penguinrandomhouse.com/books/605159/working-by-robert-a-caro/
[24] Jeff Conklin. 1987. Hypertext: An Introduction and Survey. *IEEE Computer* 20, 9 (1987), 17–41. doi:10.1109/MC.1987.1663693
[25] Jeff Conklin and Michael L. Begeman. 1988. gIBIS: A Hypertext Tool for Exploratory Policy Discussion. *ACM Transactions on Information Systems (TOIS)* 6, 4 (1988), 303–331. doi:10.1145/58566.59297
[26] William Wallace Cook. 1928. *Plotto.* Ellis Publishing Company, Battle Creek, MI, USA. 300 pages.
[27] Bradin Cormack and Carla Mazzio. 2005. *Book Use, Book Theory: 1500-1700 (Exhibition pamphlet).* The University of Chicago Library, Chicago, IL, USA. https://storage.lib.uchicago.edu/pres/2011/pres2011-0098.pdf
[28] Laura De Young. 1990. Linking Considered Harmful. In *Proceedings of the First European Conference on Hypertext, Hypertext: Concepts, Systems and Applications* (Proceedings of the First European Conference on Hypertext). Cambridge University Press, Cambridge, UK, 238–249.
[29] Andreas Dieberger. 1993. The Information City–A Metaphor for Navigating Hypertexts. In *BCS-HCI (Vol.93)*. BCS Learning & Development Ltd, Swindon, UK.
[30] Andreas Dieberger. 1994. *Navigation in Textual Virtual Environments Using A City Metaphor.* Ph. D. Dissertation. Vienna University of Technology, Vienna,

Austria.
[31] Andreas Dieberger. 1997. A City Metaphor to Support Navigation in Complex Information Spaces. In *A City Metaphor to Support Navigation in Complex Information Spaces* (Laurel Highlands, PA, USA), Stephen C. Hirtle and Andrew W. Frank (Eds.), Vol. International Conference on Spatial Information Theory. Springer, Berlin/Heidelberg, 53–67. doi:10.1007/3-540-63623-4_42
[32] Andreas Dieberger and Jay D. Bolter. 1995. On the design of hyper “spaces”. *Communications of The ACM* 38, 8 (1995), 98. doi:10.1145/208344.208366
[33] Andreas Dieberger and Andrew W. Frank. 1998. A City Metaphor to Support Navigation in Complex Information Spaces. *Journal of Visual Languages & Computing* 9, 6 (1998), 597–622. doi:10.1006/jvlc.1998.0100
[34] Andreas Dieberger, M. Pohl, and P. Purgathofer. 1995. Einige Fragen zur Entwicklung von Hypertext-Autorenwerkzeugen. *Hypermedia in der Aus-und Weiterbildung* 17 (1995).
[35] Andreas Dieberger and Jolanda G. Tromp. 1993. The Information City project-a virtual reality user interface for navigation in information spaces Was to appear in. In *Proceedings of the Symposium Virtual Reality Vienna* (Vienna, Autstria). IDG GMBH, Vienna, Austria.
[36] Matthew Daniel Eddy. 2022. *Media and the Mind: Art, Science, and Notebooks as Paper Machines, 1700-1830.* University of Chicago Press, Chicago, IL, USA. 531 pages. https://press.uchicago.edu/ucp/books/book/chicago/M/bo173084820.html
[37] Sergei Eisenstein. 1943. *The Film Sense.* Harcourt, Brace and Co., New York. 288 pages.
[38] Ayelet Even-Ezra. 2021. *Lines of Thought.* University of Chicago Press, Chicago, IL, USA. 272 pages. https://press.uchicago.edu/ucp/books/book/chicago/L/bo63098990.html
[39] Kim Michael Fairchild and Steven E. Poltrock. 1987. Soaring through Knowledge Space: SemNet 2.1 (video). In *Human Factors in Computing Systems* (Toronto, ON, Canada). ACM, New York, NY, USA. https://www.youtube.com/watch?v=yecrkVY4kFo Videotape presented at CHI'87 Human Factors in Computing Systems.
[40] Kim Michael Fairchild, Steven E. Poltrock, and George W. Furnas. 1988. SemNet: Three-Dimensional Graphic Representations of Large Knowledge Bases. In *Cognitive Science and Its Applications for Human-computer Interaction*, Raymonde Guindon (Ed.). Lawrence Erlbaum Associates, Hillsdale, NJ, 201–233. doi:10.4324/9780203782040
[41] Paul A. Fishwick. 2006. *Aesthetic Computing.* MIT Press, Cambridge, MA, USA. 476 pages. https://mitpress.mit.edu/9780262562379/aesthetic-computing/
[42] Judith Flanders. 2020. *A Place For Everything.* Pan Macmillan, London, UK. 272 pages. https://www.panmacmillan.com/authors/judith-flanders/a-place-for-everything/9781509881581
[43] George W. Furnas. 1997. Effective View Navigation. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, GA, USA). Association for Computing Machinery, New York, NY, USA, 367–374. doi:10.1145/258549.258800
[44] Richard Furuta, Frank M. Shipman, III, Catherine C. Marshall, Donald Brenner, and Haowei Hsieh. 1997. Hypertext Paths and the World-Wide Web: Experiences with Walden's Paths. In *Proceedings of the Eighth ACM Conference on Hypertext* (Southampton, UK). ACM, New York, NY, USA, 167–176. doi:10.1145/267437.267455
[45] Steve Gaynor. 2003. *Gone Home.* The Fulbright Company. Retrieved April 12, 2025 from https://gonehome.com
[46] Jan Gehl. 2011. *Life Between Buildings: Using Public Space.* Island Press, Washington, DC, USA. 216 pages. https://islandpress.org/books/life-between-buildings
[47] Gérard Genette. 1980. *Narrative Discourse: An Essay in Method.* Cornell University Press, Ithaca, NY, USA. 285 pages. https://www.cornellpress.cornell.edu/book/9780801417580/narrative-discourse-revisited
[48] Bernard Dionysius Geoghegan. 2012. La cybernétique « américaine » au sein du structuralisme « français ». Jakobson, Lévi-Strauss et la Fondation Rockefeller. *Revue d'Anthropologie des Connaissances* Vol. 6, n° 3, 3 (2012), 585–601. doi:10.3917/rac.017.0117
[49] Edward Gibbon. 1776. *The History of the Decline and Fall of the Roman Empire.* Vol. I–VI. Strahan & Cadell, London, UK.
[50] Gene Golovchinsky. 1997. What the query told the link: the integration of hypertext and information retrieval. In *Proceedings of the Eighth ACM Conference on Hypertext* (Southampton, United Kingdom). Association for Computing Machinery, New York, NY, USA, 67–74. doi:10.1145/267437.267445
[51] Anthony Grafton. 1997. *The Footnote - A Curious History.* Harvard University Press, Cambridge, MA, USA. 255 pages. https://www.hup.harvard.edu/books/9780674307605
[52] Frank G. Halasz, Thomas P. Moran, and Randall H. Trigg. 1986. NoteCards in a Nutshell. In *Proceedings of the SIGCHI/GI Conference on Human Factors in Computing Systems and Graphics Interface* (Toronto, Ontario, Canada). ACM, New York, NY, USA, 45–52. doi:10.1145/29933.30859
[53] Heiko Haller and Andreas Abecker. 2010. iMapping: a zooming user interface approach for personal and semantic knowledge management. In *Proceedings of the 21st ACM Conference on Hypertext and Hypermedia* (Toronto, OT, Canada). Association for Computing Machinery, New York, NY, USA, 119–128. doi:10.1145/1810617.1810638
[54] Jillian M. Hess. 2022. *How Romantics and Victorians Organized Information: Commonplace Books, Scrapbooks, and Albums.* Oxford University Press, Oxford, UK. 320 pages. doi:10.1093/oso/9780192895318.001.0001
[55] Kenneth E. Iverson. 1980. Notation As A Tool of Thought (ACM 1979 Turing Award lecture). *Commun. ACM* 23, 8 (1980), 444–465. doi:10.1145/358896.358899
[56] H. J. Jackson. 2001. *Marginalia: Readers Writing in Books.* Yale University Press, New Haven, CT, USA. 324 pages. https://yalebooks.co.uk/book/9780300097207/marginalia/
[57] Kevin Jackson. 2000. *Invisible Forms: A Guide to Literary Curiosities.* St Martin's Press, New York, NY, USA. 310 pages.
[58] Michael Joyce. 1987. *afternoon: a story.* Eastgate Systems, Inc., Watertown, MA.
[59] Michael Joyce. 1991. Storyspace As a Hypertext System for Writers and Readers of Varying Ability. In *Proceedings of the Third Annual ACM Conference on Hypertext* (San Antonio, TX, USA). ACM, New York, NY, 381–387. doi:10.1145/122974.125110
[60] Pauling Kael. 1969. *Trash, Art, And The Movies.* Harper' Magazine. Retrieved April 13, 2025 from https://harpers.org/archive/1969/02/trash-art-and-the-movies/
[61] Paul Kahn and Krzysztof Lenk. 2001. *Mapping Web Sites.* Rotovision SA, Hove, UK. 144 pages.
[62] Michael Kempe. 2024. *The Best of All Possible Worlds: A Life of Leibniz in Seven Pivotal Days.* Puskin Press, London, UK. 179 pages. https://pushkinpress.com/book/the-best-of-all-possible-worlds
[63] Masaaki Kurosu and Kaori Kashimura. 1995. Apparent Usability vs. Inherent Usability: Experimental Analysis on The Determinants of The Apparent Usability. In *Conference Companion on Human Factors in Computing Systems* (Denver, Colorado, USA). Association for Computing Machinery, New York, NY, USA, 292–293. doi:10.1145/223355.223680
[64] George P. Landow. 1979. Make Mine an Old Timer: the joys of railroading in the 1880s as seen in one day in the life of the AP&G. *Railroad Modeller* 9 (1979), 48–55. https://victorianweb.org/cv/models/articles/oldtimer2.html
[65] George P. Landow. 1992. *Hypertext: The Convergence of Contemporary Critical Theory and Technology.* Johns Hopkins University Press, Balimore, MD, USA. 326 pages.
[66] Douglas B. Lenat, Alan Borning, David W. McDonald, Craig Taylor, and Steven Weyer. 1983. Knoesphere: Building Expert Systems With Encyclopedic Knowledge. In *International Joint Conference on Artificial Intelligence* (Karlsruhe, DE). William Kaufmann, San Francisco, CA, USA, 167–169. doi:10.5555/1623373.1623410
[67] Manuel Lima. 2014. *The Book of Trees.* Princeton Architectural Press, New York, NY, USA. 192 pages. https://papress.com/products/the-book-of-trees-visualizing-branches-of-knowledge
[68] Niklas Luhmann. 1992. What is Communication? *Communication Theory* 2, 3 (1992), 251–259. doi:10.1111/j.1468-2885.1992.tb00042.x
[69] Kevin Lynch. 1964. *The Image of the City.* MIT Press, Cambidge, MA, USA. 206 pages. https://mitpress.mit.edu/9780262620017/the-image-of-the-city/
[70] Josephine Machon. 2019. *The Punchdrunk Encyclopaedia.* Routledge, London, UK. 322 pages.
[71] Norman Macrae. 1999. *John von Neumann: the scientific genius who pioneered the modern computer, game theory, nuclear deterrence, and much more.* American Mathematical Society, Providence, RI. https://www.routledge.com/The-Punchdrunk-Encyclopaedia/Machon/p/book/9781138556799
[72] Catherine C. Marshall. 1998. Toward an Ecology of Hypertext Annotation. In *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia : Links, Objects, Time and Space—structure in Hypermedia Systems* (Pittsburgh, PA, USA). ACM, New York, NY, USA, 40–49. doi:10.1145/276627.276632
[73] Catherine C. Marshall, Frank G. Halasz, Russell A. Rogers, and William C. Janssen, Jr. 1991. Aquanet: A Hypertext Tool to Hold Your Knowledge in Place. In *Proceedings of the Third Annual ACM Conference on Hypertext* (San Antonio, TX, USA). ACM, New York, NY, USA, 261–275. doi:10.1145/122974.123000
[74] Catherine C. Marshall and Russell A. Rogers. 1992. Two Years before the Mist: Experiences with Aquanet. In *Proceedings of the ACM Conference on Hypertext* (Milan, Italy). ACM, New York, NY, USA, 53–62. doi:10.1145/168466.168490
[75] Catherine C. Marshall and Frank M. Shipman, III. 1993. Searching for the Missing Link: Discovering Implicit Structure in Spatial Hypertext. In *Proceedings of the Fifth ACM Conference on Hypertext* (Seattle, WA, USA). ACM, New York, NY, USA, 217–230. doi:10.1145/168750.168826
[76] Catherine C. Marshall, Frank M. Shipman, III, and James H. Coombs. 1994. VIKI: Spatial Hypertext Supporting Emergent Structure. In *Proceedings of the 1994 ACM European Conference on Hypermedia Technology* (Edinburgh, UK). ACM, New York, NY, USA, 13–23. doi:10.1145/192757.192759
[77] Catherine C. Marshall, Frank M. Shipman, III, and Mark LeMere. 1999. Beyond location: hypertext workspaces and non-linear views. In *Proceedings of the tenth ACM Conference on Hypertext and hypermedia: returning to our diverse roots: returning to our diverse roots* (Darmstadt, Germany). ACM, New York, NY, USA, 121–130. doi:10.1145/294469.294498

[78] John McDaid. 1992. *Uncle Buddy's Phantom Funhouse.* Eastgate Systems, Inc., Watertown, MA, USA. https://dtc-wsuv.org/projects/uncle-buddy/archival/index.html

[79] John McPhee. 2017. *Draft no. 4: On The Writing Process.* Farrar, Straus and Giroux, New York.

[80] Louis Menand. 1995. *Finding It at the Movies.* The New York Review. Retrieved April 13, 2025 from https://www.nybooks.com/articles/1995/03/23/finding-it-at-the-movies/

[81] Louis Menand. 2021. *The Free World: Art and Thought in the Cold War.* Farrar, Straus and Giroux, New York, NY, USA. 880 pages. https://us.macmillan.com/books/9781250829641/thefreeworld/

[82] David E. Millard and Charlie Hargood. 2017. Tiree Tales: A Co-operative Inquiry into the Poetics of Location-Based Narrative. In *Proceedings of the 28th ACM Conference on Hypertext and Social Media* (Prague, Czech Republic). ACM, New York, NY, USA, 15–24. doi:10.1145/3078714.3078716

[83] David E. Millard, Charlie Hargood, Michael O. Jewell, and Mark J. Weal. 2013. Canyons, Deltas and Plains: Towards a Unified Sculptural Model of Location-Based Hypertext. In *Proceedings of the 24th ACM Conference on Hypertext and Social Media* (Paris, France). ACM, New York, NY, USA, 109–118. doi:10.1145/2481492.2481504

[84] Rand Miller and Robin Miller. 1993. *Myst.* Cyan, Spokane, WA, USA.

[85] Franco Moretti. 2007. *Graphs, Maps, Trees: Abstract Models for a Literary History.* Verso, London, UK. 124 pages. https://www.versobooks.com/en-gb/products/1939-graphs-maps-trees

[86] Stuart Moulthrop. 1991. *Victory Garden.* Eastgate Systems, Inc., Watertown, MA, USA. https://www.eastgate.com/catalog/VictoryGarden.html

[87] Theodor Holm Nelson. 1972. As We Will Think. In *Proceedings of the International Conference on Online Interactive Computing, Volume 1* (Uxbridge, UK). Online Computer Systems Ltd, Brunel University, Uxbridge, UK, 439–454.

[88] Theodor Holm Nelson. 1974. *Computer Lib/Dream Machines.* Vol. 1978 Reprint of 1974 Edition. Self-published, n/a. 336 pages. https://computerlibbook.com

[89] Jakob Nielsen. 1990. *Hypertext And Hypermedia.* Academic Press Professional, Inc., San Diego, CA, USA. https://dl.acm.org/doi/abs/10.5555/79094

[90] Jakob Nielsen. 1994. Enhancing the explanatory power of usability heuristics. In *Conference Companion on Human Factors in Computing Systems* (Boston, MA, USA). Association for Computing Machinery, New York, NY, USA, 210. doi:10.1145/259963.260333

[91] Donald A. Norman. 2004. *Emotional Design: Why We Love (or Hate) Everyday Things.* Basic Books, New York, NY, USA. 276 pages. https://www.hachettebookgroup.com/titles/don-norman/emotional-design/9780465051366/

[92] Alexandra Okada, Simon J. Buckingham Shum, and Tony. Sherborne (Eds.). 2008. *Knowledge Cartography: Software Tools and Mapping Techniques.* Springer, New York, NY, USA. 540 pages. doi:10.1007/978-1-4471-6470-8

[93] Stephen Orgel. 2015. *The Reader in the Book: A Study of Spaces and Traces.* Oxford University Press, Oxford, UK. 186 pages. https://global.oup.com/academic/product/the-reader-in-the-book-9780198737568

[94] Luke Caspar Pearson and Sandra Youkhana. 2022. *Videogame Atlas: Mapping Interactive Worlds.* Thames & Hudson, London, UK. 288 pages. https://thamesandhudson.com/videogame-atlas-mapping-interactive-worlds-9780500024232

[95] Amanda H. Podany. 2022. *Weavers, Scribes, and Kings: A New History of the Ancient Near East.* Oxford University Press, Oxford, UK. 673 pages. https://global.oup.com/academic/product/weavers-scribes-and-kings-9780190059040

[96] Sandra Rendgen. 2019. *History of Information Graphics.* Taschen, Köln, DE. 462 pages. https://www.taschen.com/en/books/graphic-design/03435/history-of-information-graphics/

[97] L. D. Reynolds and N. G. Wilson. 2013. *Scribes and Scholars: A Guide to the Transmission of Greek and Latin Literature.* Oxford University Press, USA, new York, NY, USA. 353 pages. https://global.oup.com/ukhe/product/scribes-and-scholars-9780199686339

[98] Daniel Rosenberg and Anthony Grafton. 2010. *Cartographies of Time.* Princeton Architectural Press, New York, NY, USA. 272 pages.

[99] Thomas Schedel and Claus Atzenbeck. 2016. Spatio-Temporal Parsing in Spatial Hypermedia. In *Proceedings of the 27th ACM Conference on Hypertext and Social Media* (Halifax, NS, Canada). ACM, New York, NY, USA, 149–157. doi:10.1145/2914586.2914596

[100] Denise Schmandt-Besserat. 1996. *How Writing Came About.* University of Texas Press, Austin, TX, USA. 216 pages. https://utpress.utexas.edu/9780292777040/

[101] Frank M. Shipman, III, Haowei Hsieh, Preetam Maloor, and J. Michael Moore. 2001. The Visual Knowledge Builder: A Second Generation Spatial Hypertext. In *Proceedings of the 12th ACM Conference on Hypertext and Hypermedia* (Århus, Denmark). ACM, New York, NY, USA, 113–122. doi:10.1145/504216.504245

[102] Frank M. Shipman, III and Catherine C. Marshall. 1993. *Formality Considered Harmful: Experiences, Emerging Themes, and Directions (Technical Report ISTL-CSA-94-08-02).* Technical Report. Xerox Palo Alto Research Center. https://people.engr.tamu.edu/shipman/formality-paper/harmful.html Revised/republished in 1999 as [103].

[103] Frank M. Shipman, III and Catherine C. Marshall. 1999. Formality Considered Harmful: Experiences, Emerging Themes, and Directions on the Use of Formal Representations in Interactive Systems. *Computer Supported Cooperative Work (CSCW)* 8, 4 (1999), 333–352. doi:10.1023/A:1008716330212

[104] Frank M. Shipman, III, Catherine C. Marshall, and Thomas P. Moran. 1995. Finding and Using Implicit Structure in Human-organized Spatial Layouts of Information. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Denver, CO, USA). ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 346–353. doi:10.1145/223904.223949

[105] David Sim. 2019. *Soft City: Building Density for Everyday Life.* Island Press, Washington, DC, USA. 256 pages. https://islandpress.org/books/soft-city

[106] Daniel E. Snyder. 2021. *The Tender Detail: Ornament and Sentimentality in the Architecture of Louis H. Sullivan and Frank Lloyd Wright.* Bloomsbury Publishing, London, UK. 305 pages. https://www.bloomsbury.com/uk/tender-detail-9781350236714/

[107] Norbert A. Streitz, Jörg Haake, Jörg Hannemann, Andreas Lemke, Wolfgang Schuler, Helge Schütt, and Manfred Thüring. 1992. SEPIA: A Cooperative Hypermedia Authoring Environment. In *Proceedings of the ACM Conference on Hypertext* (Milan, Italy). ACM, New York, NY, USA, 11–22. doi:10.1145/168466.168479

[108] Louis H. Sullivan. 1896. The Tall Office Building Artistically Considered. *Lippincott's Magazine* 57 (1896), 403–408. https://archive.org/details/sim_mcbrides-magazine_1896-03_57/page/402/mode/2up

[109] Dennis Yi Tenen. 2024. *Literary Theory for Robots: How Computers Learned to Write.* WW Norton & Company, New York, NY, USA. 176 pages. https://wwnorton.com/books/9781324105053

[110] The Atlantic 1945. As We May Think. *The Atlantic Monthly* 176, 1 (1945), 101–108 pages.

[111] Noam Tractinsky. 1997. Aesthetics and apparent usability: empirically assessing cultural and methodological issues. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems* (Atlanta, GA, USA). Association for Computing Machinery, New York, NY, USA, 115–122. doi:10.1145/258549.258626

[112] Edward R. Tufte. 1982. *The Visual Display of Quantitative Information.* Graphics Press, Cheshire, CT, USA. 197 pages. https://www.edwardtufte.com/book/the-visual-display-of-quantitative-information/

[113] Edward R. Tufte. 1990. *Envisioning Information.* Graphics Press, Cheshire, CT, USA. 126 pages. https://www.edwardtufte.com/book/envisioning-information/

[114] Peter Turchi. 2007. *Maps of the Imagination: The Writer as Cartographer.* Trinity University Press, San Antonio, TX, USA. 259 pages. https://tupress.org/9781595340412/maps-of-the-imagination/

[115] Jacquers Vallee. 1982. *The network revolution : confessions of a computer scientist.* And/Or Press, Berkeley, CA, USA. 228 pages. https://archive.org/details/networkrevolutio00vall

[116] John Vinci. 2021. *Reconstructing the Garrick: Adler & Sullivan's Lost Masterpiece.* Alphawood Exhibitions, Chicago, IL, USA. 256 pages.

[117] Janet H. Walker. 1987. Document Examiner: Delivery Interface for Hypertext Documents. In *Proceedings of the ACM Conference on Hypertext* (Chapel Hill, NC, USA). ACM, New York, NY, USA, 307–323. doi:10.1145/317426.317448

[118] Edward Wilson-Lee. 2018. *The Catalogue of Shipwrecked Books.* William Collins, London, Uk. 416 pages. https://www.williamcollinsbooks.co.uk/products/the-catalogue-of-shipwrecked-books-young-columbus-and-the-quest-for-a-universal-library-edward-wilson-lee-9780008146245/

[119] Nicole Yankelovich, Bernard J. Haan, Norman K. Meyrowitz, and Steven M. Drucker. 1988. Intermedia: The Concept and the Construction of a Seamless Information Environment. *Computer* 21, 1 (1988), 81–96. doi:10.1109/2.222120

[120] Polle T. Zellweger, Bay-Wei Chang, and Jock D. Mackinlay. 1998. Fluid Links for Informed and Incremental Link Transitions. In *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia : Links, Objects, Time and Space—Structure in Hypermedia Systems* (Pittsburgh, PA, USA). ACM, New York, NY, USA, 50–57. doi:10.1145/276627.276633

[121] Chuck Zerby. 2003. *The Devil's Details: A History of Footnotes.* Touchstone, New York, NY, USA. 162 pages. https://www.simonandschuster.com/books/The-Devils-Details/Chuck-Zerby/9780743241755

[122] Siyi Zhu, Robert Haisfield, Brendan Langen, and Joel Chan. 2024. Patterns of Hypertext-Augmented Sensemaking. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology* (Pittsburgh, PA, USA) *(UIST '24).* Association for Computing Machinery, New York, NY, USA, Article 143, 17 pages. doi:10.1145/3654777.3676338