Narrative In The Information City

Mark Bernstein[1,*,†]

[1]Eastgate Systemms, Inc. 134 Main St, Watertown MA 02472 USA

Abstract

In the 1990s, hypertext pioneer and artist Andreas Dieberger suggested that a cityscape might inspire a useful hypertext visualization. A fresh implementation of this idea is presented in this year's conference [1], motivated in part by a desire to reduce our visualizations' dependence on textual labels. Here, I speculate on some unexpected narrative elements that arise in using an information city. Varied building types turn out to be important for practical as well as aesthetic reasons, but those variations cry out for etiological explanation. Though LLMs are notoriously facile and unreliable, it seems likely that they can invent useful explanations for why our imaginary city looks as it does.

Keywords

hypertext, visualization, narrative,

1. On the Information City

In this year's Proceedings I have a paper with Silas Hooper and Mark Anderson about returning "Back To The Information City" [1]. The Information City was a hypertext visualization proposed by the late hypertext researcher and artist Andreas Dieberger in his dissertation [2] and elaborated in subsequent papers [3]. Today we know more about computer graphics, we possess GPUs, we have better screens and much faster processors. The time seemed ripe to try a fresh implementation.

Dieberger's chief concern was navigation. Investigators of that era were concerned that readers would become lost in hypertexts [4], and that readers traversed hypertexts inefficiently, revisiting nodes they had already seen [5]. Dieberger believed that people already knew how to recognize and move around in cities, and could apply that knowledge to the visualization. (In practice, as Dieberger knew, people are not in fact very adept at understanding cities [6]). Some form of 3D navigation is likely necessary for getting around in a large information city [7], but 3D navigation is actually quite difficult for humans accustomed to moving over the ground [8].

Fortunately, worries about navigation proved largely ill-founded [9, 10]. Contemporary concerns have changed. Our best hypertext visualizations remain variations on a diagram of boxes and arrows, one that would be familiar to users of many of the earliest hypertext systems. What else might be done [11]?

One problem with boxes and arrows is that the screen rapidly fills with the names of notes. A large map view (Figure 1) may have 100 notes; that is a lot of text labels to read. Moreover, we have several layers of calligraphic lines placed together and, often, overlapping: text labels, the outlines of the boxes, and the arcs that represent links all lie in some way on the picture plane, and the result is not always aesthetically pleasing. In actual towns and cities, few buildings are prominently labeled. Those rare architectural sites with label densities characteristic of hypertext visualizations, such as Times Square, Piccadilly, or Shibuya are considered cacophonous. One goal of our visualization was to remove almost all text from the screen at rest, presenting transient labels only when the mouse moves over an object.

None of the authors of "Back To The Information City" were entirely satisfied that this new visualization will prove useful. Despite these reservations, I would like to suggest some tentative ideas about narrative in the Information City. Narrative is not an obvious concern of cityscapes; narrative

NHT 2025: Narrative and Hypertext September 15-18, 2025, Chicago, IL

bernstein@eastgate.com (M. Bernstein)

https://www.eastgate.com/ (M. Bernstein)

0000-0001-8861-5664 (M. Bernstein)

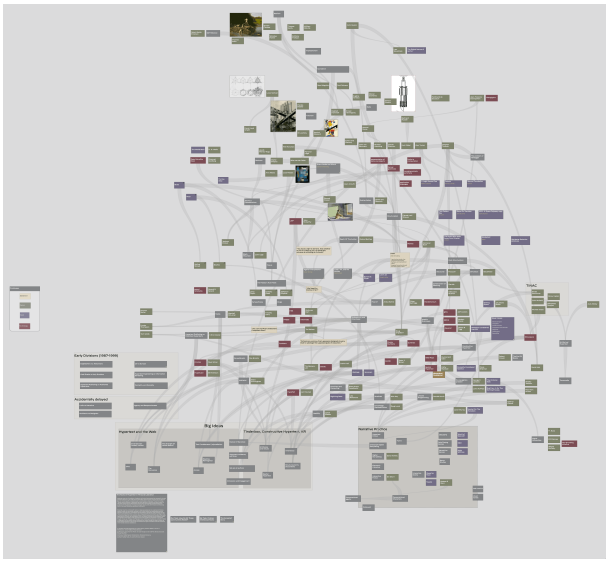

Figure 1: A conventional map view

happens within the city [12, 13] or (sometimes) to the city [14, 15] (though see the architectural theory of Bernard Tschumi and of Peter Eisenman for contrary views). Still, personal experience using our Information City to write some papers suggests a few stories worth a brief telling.

2. Footnotes

My interdisciplinary papers often have abundant references and footnotes [1][16][17]. Documentation reassures reviewers, helping them check assertions that surprise them and calming their discomfort at reviewing interdisciplinary work. But references also play a narrative role: as Grafton observes, they can explain how we know things in counterpoint to the main exposition that explains what we know [18].

One of the very first implementation tests for our Information City was a simple grid of 100 stylized Cape Cod houses (Figure 2), which I imagined to be a sort of subdivision of references. If performance was acceptable with 100 references, I thought we would be safe: surely any real paper would have fewer. (In the end, "Back To The Information City" wound up with about 120.) Yet it became clear at once that remembering where a specific reference was would prove difficult. At rest, we have a hundred little white houses, all alike. Pointing at one house lets its label float up from the house toward the sky. Brushing along several houses gives us several labels. Still, the writer looking for a specific paper has an O(n) journey in store.

If, on the other hand, we could divide the references into several visually-distinct classes or categories, those categories could limit the amount of searching. If we already know, say, that blue houses are foundational studies, red houses are art history and aesthetic theory, purple houses are computer graphics, then hunting a specific reference requires **O(n/m)** time with m categories. Even if we do not know what a blue house is intended to signify, experienced readers may be able to intuit the likely categories from an example or two. This gives us **O(m+n/m)**. Moreover, if other properties of the reference (e.g. its date and length) are reflected in other visual facets (e.g. trim color, window style), many of the references will acquire unique visual characteristics. This is, in fact, how we navigate in cities; we know what the intersection of State and Madison looks like even though we don't actually know very much about the buildings.

Interestingly, this is how ancient societies organized large cities for message delivery. Ancient Rome had a million people, so using linear search to find a named stranger would be intractable. Even specific clues — the little red house just off the street of the fish-sellers — were of limited use when a city might

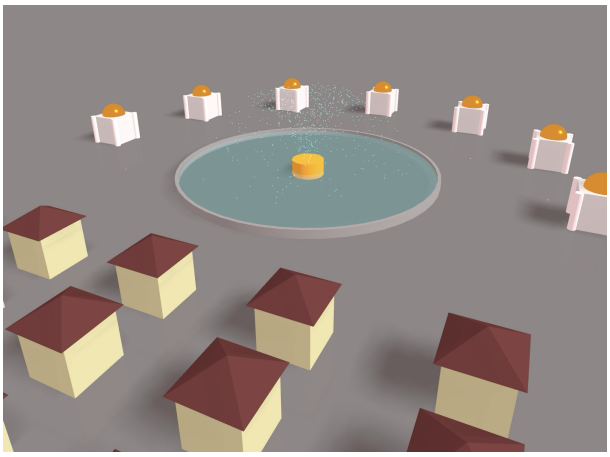

Figure 2: A neighborhood of reference abuts a group of discursive notes around an ornamental fountain

have many distinct fish markets. Rome was divided into 14 *regiones*, each of which was divided into a variable number of *vici* or "wards". In all, Rome had 256 *vici*. You could address a letter to your friend Marcus in that red house near the fish market in *regio* VI *vicus* I, and expect your messenger to find the correct place. Even with a million residents, the messenger has narrowed his search to a neighborhood of 4,000. Add a simple locale (near the fishmarket, not far from the baths) and it's likely that asking a few people on the street would pin the recipient down. [1]

On reflection, actual towns and cities — especially towns and cities that people like — tend to have many buildings, each different. Some unities help make a place look coherent; commonalities of style, cornice line, color palette, or architectural vocabulary make Paris look Parisian and not like Santa Fe [20]. Yet even in tightly-controlled places and even when buildings all began from a standard plan, over time each structure tends to acquire its own characteristics [21].

To facilitate variation, our Information City supports an open set of simple model buildings that correspond to Tinderbox notes. Multistory buildings present a container of notes, which can be useful for presenting collections like Release Notes Tower. A note may have an optional dictionary of modifications, moreover, which affect things like door treatment, window style, interior lighting, wall color and trim color.

Exceptional, uncharacteristic, or conspicuous buildings can serve as landmarks or navigational guides [6]; Release Notes Tower may not be important, but knowing it's there can help keep us anchored.

The tension between wondering "what is this?" and discovering "that's how it works!" is essentially narrative, a journey toward understanding. It is, I think, a central theme in open world games like Elder Scrolls or Cyberpunk 2077 [22].

3. Looking At Pictures

Most buildings in our Information City are unlike the majority of their neighbors. This looks nice and improves findability. Not every building has to be unique; a short run of identical, or nearly identical, row houses can be attractive and it may be easy enough to distinguish "the one of the left", "the middle house", and "the one on the right". Over time, working in these synthetic places becomes fairly seamless; when I felt it was too hard to figure out which house was which, it was easy enough to put out a flag or add some street furniture to distinguish it. Metadata might also be reflected automatically in appearance; older notes might have denser ivy on their walls, and often-visited notes might have more wear on the path to their door. Finally, it is easy to write actions to randomly mutate some characteristics, simple to differentiate a note from its neighbors.

[1] Modern computer scientists might have proposed 16 *regiones* of 16 *vici*, but the Romans came close to an optimal configuration [19].

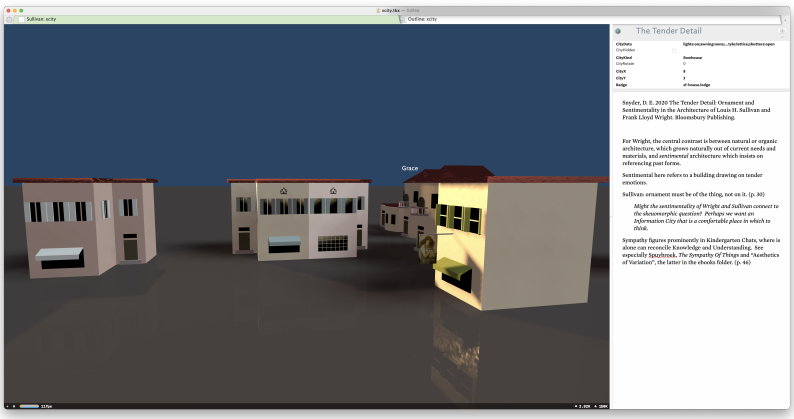

Figure 3: A view of an information city looks at Aesthetics Square, with the Ruskin Pub at right and notes on Ruskin's Gothic virtues—savageness, changefulness, naturalism etc. — along the lane at right rear.

As in real cities, and in paintings of townscapes, familiarity with the (artificial) built environment is an invitation to narrative speculation. We see a house with red trim on a street where the other houses have white trim, and we imagine the personality who chose that color, and perhaps the reaction of the neighbors. We see a house with an unusual veranda and imagine the old lady who used to sit there with her mint julep and her knitting, and wonder why that old lady had so few visitors. Implicit *etiological* narratives can help explain or rationalize appearance, and over time might help a user recall where she put things.

4. Biographies Of Buildings

Because we don't often tell stories of buildings, it might be worth taking a minute to recall some precedents. I exclude for this purpose archaeological studies, which are often diachronic but do not accord with our present purpose.

The most biographical treatment of a large, modern building that I know is *Reconstructing the Garrick : Adler & Sullivan's lost masterpiece* by John Vinci et al. [23]. The Garrick Theater (1890-1960) was a Chicago skyscraper designed by Adler and Sullivan. It was, briefly, one of the tallest buildings in the world. It was built as part of an effort by German-American immigrants to emphasize that Germans were modern and, often, highly educated; its office tower was intended to subsidize the theater in a program of German performances. Over the years, inevitably, design features that once read as advanced came to seem out-dated and passé. The meaning of *Bildung* to the American Midwest changed with two wars, growing English-language competence, and with the gradual acceptance of German Jews in the US. The salience (and profitability) of live theater was damaged by film and damaged again by television, and though the Garrick was converted in turn to a movie house and a television studio, it was never ideal for either task. Ultimately, the owners decided to replace it with a parking garage. Vinci documents the building's changing uses and office lessees in detail, as well as documenting demolition to better understand its construction.

A different approach was adopted by Paulina Bren in *The Barbizon: the hotel that set women free* [24]. The Barbizon was originally developed (in 1927) as a residence for unmarried women who wish to reside for a time in Manhattan. It offered a convenient address, chaperonage, very small rooms, and respectability. Over time, the gradual liberation of women reduced the demand for womens' hotels, but the Barbizon was kept afloat by a secretarial school, a modeling agency, and by Mademoiselle's Guest Editor Contest, which for decades brought winners from all over the country for an extended immersion in editing, fashion journalism, and New York. (The author's mother, Patsy Starrels, was one of those girls at Mademoiselle; Sylvia Plath was a winner the following year). There is little about the building *per se* in this biography; what matters are the people who passed through those doors.

Figure 4: A sentimental image of Provincetown, Massachusetts.

5. Generated Narrative

As we take notes on our scientific results, we may not want to expend much effort to imagine, store, or recall implicit etiological narratives. LLMs are well adapted to generating these, and seem remarkably successful at doing so. In this role, their propensity to lie is inconsequential.

For example, suppose we have a small house with red accents on a street old larger white houses. Its front walk is paved with the bottoms of beer and soda bottles. Claude Sonnet 4 gives us an explanation:

> Old Ezra Dunham was the town's most notorious contrarian. When his neighbors painted their houses pristine white for the summer people in 1924, Ezra slapped barn-red trim on his shutters and door. "If they want white," he'd grumble, "they can look elsewhere."
>
> The bottle walkway came later, after his wife Mehitabel passed in '31. Ezra had been the town's unofficial bottle collector for decades — folks would drop their empties at his place rather than cart them to the dump. After the funeral, something broke in him. He couldn't bear to throw away all those bottles that had once held the simple pleasures of his neighbors' lives.

Generating explanations could easily be automated, providing backstories on request or presenting bits of them in idle moments.

Focus on related characters and incidents might help deepen these explanations while making them easier to remember and to care about. The mechanics of tabletop role playing games are quite good at this [25], and might be used either to focus the attention of the LLM or to incrementally refines its prompt. Having established old Ezra and Mehitabel, it might be no great stretch to imagine another house, perhaps a bit down the road, where their daughter lived. Ezra cannot have been an easy neighbor, and that, too provides plenty of narrative springboard.

These narratives are sheer excess, almost entirely beside the point. Our goal is not to think about Ezra and Mehitabel; it is to remember where we put those notes we wrote last semester about the physics of the compound bow and the curious ways people resist innovation [26]. But if we remember that cantankerous neighbor, we may more easily recall the tiny house with its garish, if weather beaten, shutters.

Figure 5: A sentimental image of Provincetown, Massachusetts.

Acknowledgments

I emphasise that these speculations are my opinions only, and that my co-authors are not at fault for blunders herein. Regardless, I owe a great debt to much discussion and correspondence on The Information City with Mark W.R. Anderson, Silas Hooper, and Kimera Royale. I also thank Lars Spuybroek and John Vinci for helpful discussions on the buildings of Louis H. Sullivan, and also thank the Frances Loeb Library of the Graduate School of Design at Harvard University.

Declaration on Generative AI

During the preparation of this work, the author(s) used Claude Sonnet 4 in order to demonstrate its use, and to locate some forgotten references.

References

[1] M. Bernstein, M. W. R. Anderson, S. Hooper, Back to the information city, in: C. Hargood, D. Lee (Eds.), Proceedings of the 36th ACM Conference on Hypertext and Social Media, ACM, New York, 2025.

[2] A. Dieberger, Navigation in Textual Virtual Environments using a City Metaphor, Ph.D. thesis, Vienna University of Technology, Vienna, 1994.

[3] A. Dieberger, A. U. Frank, A city metaphor to support navigation in complex information spaces, Journal of Visual Languages and Computing 9 (1998) 597–622.

[4] J. Raskin, The hype in hypertext: a critique, in: Proceedings of the ACM conference on Hypertext, 1987, pp. 325–330.

[5] L. De Young, Linking considered harmful, in: N. Streitz, A. Rizk, J. André (Eds.), Hypertext: Concepts, Systems and Applications, Cambridge University Press, Cambridge, 1990, pp. 238–249.

[6] K. Lynch, The Image of the City, MIT Press, 1982.

[7] G. Furnas, Effective view navigation, CHI'97 (1997) 367–374.

[8] K. M. Fairchild, S. E. Poltrock, G. W. Furnas, Semnet: three-dimensional graphic representations of large knowledge bases, in: Readings in information visualization: using vision to think, Morgan Kaufmann Publishers Inc., 1999, pp. 190–206.

[9] G. P. Landow, The rhetoric of hypermedia: Some rules for authors, in: P. Delany, G. P. Landow (Eds.), Hypermedia and Literary Studies, The MIT Press, Cambridge, 1990.

[10] M. Bernstein, Deeply intertwingled hypertext: The navigation problem reconsidered, Technical Communication 1991 (1991) 41–47.

[11] M. Bernstein, A new view, in: Proceedings of the 35th ACM Conference on Hypertext and Social Media, 2024, pp. 65–69.
[12] C. D. Simak, City., Gnome Press, New York, 1952.
[13] S. Anderson, Winesburg, Ohio : a group of tales of Ohio small town life, B.W. Huebsch, New York, 1919.
[14] R. Meiggs, Roman Ostia, Clarendon Press, Oxford [Eng.], 1973.
[15] C. Mieville, The City & The City, Del Rey Ballantine Books, New York, 2009.
[16] C. Hargood, D. Lee (Eds.), What is it like to be Augmented?, ACM, 2025.
[17] Amiability and the Web: Lessons from Red Vienna and the Origin of Computing, ACM, New York, 2025.
[18] A. Grafton, The Footnote: A Curious History, Harvard University Press, Cambridge, 1997.
[19] C. Atzenbeck, J. Rubart (Eds.), Information Interfaces In The Antique World, ACM, 2025.
[20] C. Alexander, S. Ishikawa, M. Silverstein, e. al., A Pattern Language, Oxford University Press, 1977.
[21] H. J. Gans, The Levittowners : ways of life and politics in a new suburban community, Columbia University Press, New York, 1982.
[22] A. Bulavina, Urban planning in games, in: GDC, 2024. URL: https://www.youtube.com/watch?v=Q3Scw3dzxWE.
[23] J. Vinci, T. Samuelson, E. Nordstrom, C. Ware, Reconstructing the Garrick : Adler & Sullivan's lost masterpiece, Alphawood Foundation, Chicago, 2021.
[24] Bren, Paulina, The Barbizon : the hotel that set women free, Simon & Schuster, New York, NY, 2021.
[25] M. Miller, Traveller, Game Design Workshop, Normal, Illinois, 1977.
[26] V. Bush, As we may think, Atlantic Monthly July 1945 (1945) 101–108.