Literatronic: Use of Hamiltonian Cycles to Produce Adaptivity in Literary Hypertext

Juan B. Gutierrez
Department of Mathematics
Florida State University
208 Love Building Tallahassee
Tallahassee, FL, 32308, USA
E-mail: jgutierr@math.fsu.edu

Abstract

Literatronic is an adaptive hypermedia system for hypertext fiction. Its adaptive features are based on an algorithm that simulates a Hamiltonian cycle on a weighted graph. The algorithm maximizes narrative continuity and minimizes the probability of loosing a reader’s attention. The metric for this optimization is defined as the minimization of hypertextual friction and hypertextual attraction. We consider the challenges involved with modeling such hypertext, and we offer specific examples of this type of adaptivity.

1. Introduction

Literary hypertext is a genre of electronic literature. Its fundamental attributes have been mostly faithful to the origin of electronic text: a set of linked episodes that contain hypermedia elements. The reader chooses links to move from one page to the next, building in this fashion a story from a large set of possible stories.

Whether some features of literary hypertext could be reproduced in printed media or not has been subject of debate by opponents and proponents of digital narratives. However, as the electronic media evolves, some traits truly unique to literary hypertext have appeared. Significant effort has been invested in creating hypertexts responsive to reader’s actions by using conditional links; additionally there have been efforts to create systems capable of producing fiction, with varying degrees of success. Both of these approaches have in common that they grant greater autonomy to the computer, thus making of it an active part of the literary exchange. These changes have not only redefined the concepts of reader and author, but also they have created a new element in the reading and authoring processes: the “digital author”, which is a media element that acts on the message. The “digital author” positions itself between the reader and the author to perform a task that did not exist in printed media: the adaptive optimization of the text according to rules set by the author and the interaction of the reader.

In order to analyze a work of fiction, specially when it is a literary hypertext, it is convenient to differentiate between the story that is told as a linear sequence of temporally episode-based moments and the way that story is told, not necessarily as a linear sequence; i.e. *fabula* and *plot*. Sometimes the distinction is difficult because of phenomena such as the explanation of events that have happened in the past but have not yet been identified (*analepsis*) and the anticipation of something still to occur (*prolepsis*) [9, 5]. Given a set of pages in a hyperfiction, the plot corresponds to the sequence of pages the reader selects among many options. The fabula remains unchanged; however, the perception of it might vary depending upon the selected plot.

In this paper, we explore a mathematical model that transforms the computer in some sort of "digital author" who actively participates in the construction of the plot. The assumptions at the base of this work are: (i) the interface should focus on rendering the fabula in a way that it makes the most sense to readers according to some logic defined by authors; we can think of it as narrative continuity, regardless of the temporal or spatial continuity of events, i.e. a linear plot with variation in analepsis and prolepsis at author's discretion, (ii) readers want to explore the entire literary space, and (iii) the computer should assume the burden of assembling the pieces rather than giving them to the reader, i.e. it should become a digital author.

This paper is organized as follows: Section 2 gives references and establishes the background for this work. Section 3 describes the mathematical model and how it was implemented. Section 4 describes some results obtained with this model. Section 5 discusses some ideas presented in this paper and describes future directions.

2. Background

Literary hypertext, primarily an aesthetic form, differs from other forms of hypertext such as scholarly communication, learning systems and other types of hypertext systems in many aspects. For instance, it is not necessary to include an abstract in a work of fiction, require the existence of hyperlinks in keywords, or expect the user to complete a successful training cycle. It has been proposed and widely accepted that: (i) reader's choice, intervention and empowerment are the key elements of literary hypertext[1] [15], (ii) hypertext reading fosters both passive and active reading, where links provide decision points [16] and (iii) the suggestive power of literary hypertext lies in the lyric quality of links; lyric because of its particular intensity when searching for meaning, similar to the way we read poetry [17].

Since literary hypertext seems to be qualitatively different to other forms of hypertext, the natural question we must ask is: What are the characteristics of literary hypertext? There are two aspects of this question:

First, speaking of hypertext in general, there are several well known problems with interfaces [10]: (i) Users often do not know how to get to desired pages or how to return to previously-visited pages, (ii) they can become frustrated when they keep "rediscovering" the same page over and over, (iii) at some point the number of links may overwhelm the reader, and (iv) an author's reasons for including specific links may not always be clear to readers. In addition to the intrinsic difficulty of literary hypertext, readers are expected to be [5]: (i) tolerant enough with the up to this point experimental nature of literary hypertext, (ii) skilled enough to cope with different reading devices, and (iii) open enough to accept the non-linear result.

Secondly, there is experimental evidence reported by Gee [8] that refutes the premises described above regarding literary hypertext. According to the experiment: (i) multilinearity causes disorientation and results in readers skimming rather than reading; (ii) readers want a single starting point, (iii) they prefer narrative structures more or less linear with moderate branching, and (iv) they do not seem to be clamoring to be co-authors or to be empowered. The work by Gee has special significance. It seems to be the only factual study about readers of literary hypertext. Since most of the systems used to compose literary hypertext work with client-side technology, and those that work with server-side technology does not seem to track user interaction, there is no way to consolidate user's behavior beyond anecdotal and subjective information. However, these results might be confounded with the fact that only one tool was used, and generalization to other systems of reading might not be valid.

---

[1] That is an early 90s utopian assumption that has been challenged many times. It basically points out that links can give readers more control than printed text does.

Significant work has been done by proponents of adaptive literary hypertext. Mostly, it can be divided into two types of systems [6]: (i) Conditionality-based systems: certain rules are triggered when specific conditions are met. Examples of conditional systems are Storyspace [1] and Connection Muse [14]; some variations around this theme are what Bernstein [2] calls calligraphic vs. sculptural hypertext, that is, hypertext built by addition of links between episodes, or by the removal of them. (ii) Adaptivity -based systems: the choice of what conditions to meet depends on reader's interaction. An example of an adaptive system, though not literary-oriented, is AHA! [7].

Another way to achieve a dynamical response of text, though not necessarily adaptive, is algorithmic literature. Algorithmic generation of text driven by rules and adaptive generation has been explored by Bootz in poetry [4]. A similar goal is pursued by StoryEngine, a client constructed on top of Auld Linky structural link server [2]. The most classic attempts to use mathematics as a tool to model narrative with mathematics is Queneaus's *Cent Mille Milliards de Poèmes*.

Little, however, has been said about the computer as an author. We must make a careful distinction about exactly what part the computer authors. It is possible, in principle, to build a device that creates narratives automatically (dialogs, actions, etc.) in such a way that responds to user interaction [3]. However, in practice, this type of algorithmic literature has been more focused in interactive fiction[2]. One of the most recent examples is Façade[3].

What we call the digital author is the role of the computer when it builds links between episodes in such a way that corresponds to user interaction. While the concept of adaptivity is not new, the novelty in the approach described in this paper is that adaptivity is defined for the literary case, instead of using existing frameworks for adaptivity that were not created for the literary problem. As we will show in this paper, we develop concepts such as hypertextual friction, hypertextual attraction, and we use the well-known concept of Hamiltonian cycles, which do not have a counterpart in existing adaptive hypermedia systems.

3. Mathematical Model

In this section we deal with the formal representation of literary hypertexts. A *graph* consists of two finite sets: a set V of points, called *vertices*, and a set E of connecting lines, called *edges*, such that each edge connects two vertices, called the endpoints of the edge. In a graph $G = (V,E)$ we can walk from a vertex v_l along some edges to some other vertex v_k. The vertices belonging to an edge are called the ends, endpoints, or end vertices of the edge. An edge $e = (v_i, v_j)$ is considered to be directed from v_i to v_j ; v_j is called the head and v_i is called the tail of the edge. A graph with directed edges is called a *directed graph*; otherwise, it is called undirected. Any sequence of adjacent edges is known as a *walk*. A *trail* is a walk in which a vertex may occur at most once. A *path* is a specific kind of trail that does not visit any vertex more than once, except that its first and last vertices may be the same, in which case we call this a *cycle*. If we choose a path that passes exactly once through each vertex of an undirected graph G, we call this a *Hamiltonian path*. If the Hamiltonian path finishes in the initial vertex of an undirected graph G, thus allowing repetition of only one vertex, we call this a *Hamiltonian cycle*.

The graph representation of a hypertext places pages on the vertices. The edges are the hyperlinks. The goal is to find a reading path that visits every page exactly once, i.e. a Hamiltonian path. Tradition in literary hypertext does not require visiting every page that belongs to the hyperfiction; on the contrary,

---

[2] Interactive fiction refers to software containing simulated environments in which players use text commands to control characters.

[3] Façade, a one-act interactive drama. http://www.interactivestory.net/

dead-ends and conditionally-accessible pages became the norm in earlier hypertexts. However, this was influenced by the tools used to produce hypertext rather than by a conscious effort to produce an aesthetic effect. My position as a fiction author is that every word should be carefully crafted and deserves a place within the text. If a sentence is not needed, it should be eliminated. Likewise, if a sentence is missing, a vital part of the text is missing.

After reading all available pages, the last page of the hypertext could display the laconic sentence “The End” and dismiss the reader. But the reading of a hypertext does not necessarily ends in the last page, as printed text does. As we will show below, it is possible to create edges dynamically that simulate a Hamiltonian path, effectively creating many possible readings. Therefore, after the Hamiltonian path has been completed, the reader can be taken back to the initial page, completing in this way a Hamiltonian cycle. This serves two purposes: (i) it becomes an indication to the reader that the entire text was read, and (ii) it emphasizes the fact that there are many readings of the same text. Closing the path is not necessary, only convenient.

The reading along a path makes sense if and only if the author creates directed edges between pages that have narrative continuity. It is obvious that while the sequence of text v_j, v_k defined by the author makes sense, the sequence v_k, v_j does not necessarily have to make sense. Thus, we can enounce more precisely the problem we are trying to solve: we want to find a directed Hamiltonian cycle in the hypertext. We must emphasize the fact that “narrative continuity” is a subjective measure set by the author; it means that if an edge exists between the vertices v_j and v_k it should be possible to aggregate sequentially the text and have a seamless reading according to the author’s plan.

There are two technical complications with the model we have described: (i) given a page v_i, and some other pages to which it is connected, it will occur most of the times that some connections make more sense than others. We need some type of measure that describes the narrative continuity between pages, and (ii) a Hamiltonian cycle does not always exist for a given graph. Even if the Hamiltonian cycle existed, and the problem could be solved, nothing would guarantee that the mathematically optimal solution would make sense from the point of view of the narrative. A solution to each one of these problems will be explained in the following two subsections.

3.1. A Measure of Narrative Distance. Let us assume for the moment that page repetition is allowed. We will refer from now on to vertices as pages. The problem regarding narrative continuity is easily solved if we introduce an additional element in the graph: the narrative cost between pages. We will call it *hypertextual cost*. It is a positive real number proportional to how disruptive the transition from page v_i to page v_j is. It is necessary to define only the hypertextual cost between pages that are directly connected; the hypertextual cost between pages that are not connected is assumed to be infinite. The graph $G = (V,E)$ along with its associated costs is called a weighted graph. Since small costs maximize narrative continuity, the shortest paths between every pair of pages v_i and v_j in the graph will render the “best” plot for the fabula that starts in v_i and finishes in v_j. It is possible, in principle, to calculate this distance for each pair of pages. This is the well-known problem of finding multiterminal network flows.

Minimum paths do not necessarily solve the Hamiltonian cycle problem. In fact, left unchecked, reading a hypertext following minimum paths would result in repeated pages most of the cases. Some authors of classic hypertext used precisely repetition as narrative devices in their works. In fact, many critics and practitioners see in it an inescapable consequence of the essence of literary hypertext. However, arguably repetition has a negative impact in readers’ attention thus encouraging readers’ desertion.

We will call hypertextual attractors to those pages that disrupt the narrative continuity by repetition. We will call hypertextual attraction to the metric used to define a hypertextual attractor. Attractors pose a serious problem to authors because normally they wear down a reader's interest. Let c_{ij} represent the hypertextual cost between the origin page v_i and the destination page v_j. Let e_i represent the number of edges connected to page v_i Let V represent the total number of pages. Let A_i be the hypertextual attraction of page v_i. Then,

$$A_i = \frac{1}{\sum_{m=1}^{V} \sum_{n=1}^{V} c_{mn}} \sum_{j=1}^{e_i} \frac{c_{ij}}{2}. \qquad (1)$$

The numerator in equation 1 is the sum of costs of all connections to the page i, while the denominator is the total sum of hypertextual costs. Small values of hypertextual attraction mean that the cost of including this page in a reading path is low, therefore the page v_i is very likely to appear in many reading paths. The measure of attraction is problematic if, for example, one of the costs c_{ij} is much higher than the others; in this case, the total measure of attraction could seem appropriate, but in fact the page could be very attractive by the action of many small costs. Therefore, attraction must be used in conjunction with page stiffness, which is the ratio of the highest distance to the smallest one among those edges connected to a single page.

We will call hypertextual friction between two pages to the probability of loosing readers' attention during a navigation event that goes from the origin page v_i to the destination page v_j; we assume it is directly proportional to the hypertextual cost between the two pages, directly proportional to the hypertextual attraction of page v_j, and inversely proportional to the total number of pages read R. The latter means that narrative costs are compensated by reader's increased knowledge. High measures of friction indicate a high probability of reader's desertion. Let D_{ij} represent the hypertextual friction. Let c_{ij} represent the hypertextual cost between the origin page v_i and the destination page v_j. Let A_i be the hypertextual attraction of page v_i as defined in equation 1. Then,

$$D_{ij} = \frac{c_{ij} A_j}{R}. \qquad (2)$$

For authors, the problem of maintaining a manageable volume of narrative, and at the same time minimizing the effect of hypertextual attractors and hypertextual friction as defined in equations 1 and 2, poses an enormous practical challenge.

3.2. The Simulation of a Hamiltonian Cycle. Finding a Hamiltonian cycle in an undirected graph is a problem NP-complete. Finding it in a directed graph is even more difficult if not impossible in most cases. Therefore, we will use the special characteristics of the problem at hand to produce an approximate solution. The most important feature of literary hypertext is that if there is no direct connection between an origin page v_i and a destination page v_j, but the reader has read the intermediate pages needed to traverse the path from v_i to v_j, then a virtual edge between v_i and v_j exists. This is possible because the reader already has the intermediate information needed to fill the narrative gap between those two pages.

Here is where the computer comes into play as a digital author. Let us assume that the reader observes in each page a number of hyperlinks that allows visiting other pages. For simplicity, let us limit the number of links to three. Since it is possible to find the solution of the multiterminal network flow, the narrative cost and the shortest path between each pair of pages is known. We can assume without loss of generality that links shown to the reader are sorted by narrative cost: the first option corresponds to the page connected by the lowest narrative distance; the second option corresponds to the next bigger cost, and so on. After reading a page, the reader will know that: (i) reading is possible through several paths, and (ii) some paths offer greater narrative continuity than others. When the user selects a destination page, there are shown in the new page three links pointing to the closest pages, which do not contain pages already visited. It is in this sense that the system exhibits adaptation. The effect is that the hypertext is read through a Hamiltonian cycle, and the reading is sorted according to shortest narratives distances. A reading process that follows the previous algorithm is guaranteed to be optimum from the point of view of the author. Note that even though in practice the reader reads a hypertext that corresponds to a Hamiltonian cycle, such mathematical solution has not been found.

4. Methods and Results

The system Literatronic was implemented on the web following the algorithms described above. It is available at http://www.literatronic.com. It was developed with server-side technology, meaning that all readers' interactions are processed and stored in a central database. It has been possible for authors to mine user interaction data in order to find readers' patterns, points of desertion, attractive pages, etc.

The costs c_{ij} between the origin page v_i and the destination page v_j were arbitrarily chosen. The books in Literatronic were initially designed with five values of hypertextual cost: 5 for immediate connection, 10 for near connections, 15 for medium connections, 20 for far connections and 30 for improbable connections. Statistical analysis of hypertextual friction, hypertextual attractors and user interaction of approximately 4000 anonymous users and approximately 200 registered users, showed that a distance of 30 has little effect in the reading, and a distance of 5 behaves almost exactly as a distance of 10. Currently only three hypertextual costs are being used: $c_{ij} = \{10, 15, 20\}$.

The multiterminal network flow was solved using the Gomory-Hu algorithm [11] because of its simplicity at development time, and because it not only finds the minimum costs between every pair of pages, but also the minimal paths. It is worth noticing that calculations performed by the Gomory-Hu algorithm have to be done when the author decides to test the hypertext. It is analogous to "compiling" the hypertext before reading it. Gomory and Hu created an elegant and simple N^3 algorithm; calculation time can be noticeable for medium-size and large hypertexts (more than 100 pages).

5. Discussion

Following the algorithms described in this paper, it was possible to build an information system for literary hypertext that positions itself between the reader and the author to perform a task that did not exist in printed media: the adaptive optimization of the text according to rules set by the author and the interaction of the reader.

A high value of hypertextual attraction it indicates that a single page could potentially be linked from many other pages. The question authors must ask themselves in this case is: Will the reading make sense in all linked cases? A high value of hypertextual friction could indicate that the lyrical quality of links is being stretched. The question authors must ask themselves in this case is: Does this connection make sense? The system takes the burden of performing these checks during the writing process and advises the author accordingly. As mentioned earlier, the author can use precisely hypertextual friction and hypertextual attractors as narrative devices. The system would just make sure the author is aware of it.

Adaptivity in this context has a precise meaning: the system adapts the connections in the graph in order to produce a Hamiltonian cycle in which hypertextual friction and hypertextual attraction have been minimized. As we have seen, the challenge for the author is to set proper values to the hypertextual cost of every link. Even though it is possible to use an automated optimization method, author's discretion is always the main criteria to assign costs between pages.

The Hamiltonian cycle is not actually found but simulated by creating edges on the fly in the graph. From the system's perspective, the Hamiltonian cycle is not even requested; from the user perspective, the Hamiltonian cycle is exactly what is offered. Pages already visited are removed from the list of possibilities to continue reading. However, if too many pages are removed from the path, the effective cost between the current page and the destination page could be too big, leading to an increase of hypertextual friction, or loss of narrative continuity.

One of the main arguments opponents of hypertext have shown in the past against it is the fragmented story that is offered to the reader. In this case, the reader receives a plot that is optimized, from the narrative perspective. That is to say, the reader receives a linear text most of the times. This raises a question: Are multilinearity and fragmentation the goal of hyperfiction, or are they the product of the state of the art when the first literary hypertexts were produced? Is fragmentation a paradigm that we want to preserve? Our intention is not to answer this complex question here. However, we want to indicate that we have the ability to produce a text that exploits the essence of digital media, and that at the same time preserves the essence of narrative in a classical sense: immersion.

The experiment reported by Gee [8] must be regarded carefully. An approach is to consider that the subjects of the experiment were still too close to the paper universe, and therefore could not appreciate the full extent of hypertextual possibilities. Another approach would be to consider that in fact readers enjoy linear texts more, regardless of the media that supports them. Although it is not possible to generalize from one single experiment, these results are consistent with the author's own experience with literary hypertexts. We consider these results a strong signal of what readers expect, and what can be expected from them.

The future direction of Literatronic will be focused on extracting information from the growing database of user interaction. As of May, 2006, there are 185 members. Between November 2003 and January 2005, Literatronic.com has received about 4,000 unique visitors. This user interaction data will be analyzed anonymously to calibrate processes of supervised learning (though neural networks) and non-supervised learning (through clustering techniques). The goal will be to increase the probability of occurrence of paths favored by readers. At that point, there will be a fuzzy line between a human author and a digital author.

An adaptive literary hypertext system is a fascinating possibility. The idea is very suggestive and promises advances in this area. Even if it did not offer any answers, it would help us to improve our perspective about hyperfiction. Quoting Stephen Jacobson: "Before you do it right, you have to do it all" (Wearable Robots. Technology Review, July/August 2004. Cambridge, MA).

6. Acknowledgements

The platform Literatronic was built as part of Juan B. Gutierrez's graduate research project at the Department of Mathematics, Florida State University, between August 2002 and August 2005. The hyperfiction used in this project, Extreme Conditions was financed by the Ministry of Culture of Colombia (grant COLCULTURA/SECAB 014/1996) and the Repository of Artistic Projects of the Institute of Culture of Bogota (grant 514/1997, grant 410/1998). The third version of Extreme Conditions, the version used in this experiment, was finished in November, 2003.

References

[1] M. Bernstein. Patterns of hypertext. In Proceedings of the Nineth ACM Conference on Hypertext and Hypermedia, pages 21–29. ACM Press, 1998.

[2] M. Bernstein, D. E. Millard, and M. J. Weal. On writing sculptural hypertext. In HYPERTEXT '02: Proceedings of the thirteenth ACM conference on Hypertext and hypermedia, pages 65–66, New York, NY, USA, 2002. ACM Press.

[3] J. D. Bolter and M. Joyce. Hypertext and creative writing. In HYPERTEXT ’87: Proceeding of the ACM conference on Hypertext, pages 41–50, New York, NY, USA, 1987. ACM Press.

[4] P. Bootz. Vers un multim´edia contraint et a-m´edia. http://motsvoir.free.fr/, 176(1):101– 108, June 2002.

[5] L. Calvi. ’lector in rebus’: The role of the reader and the characteristics of hyperreading. In Proceedings of the Tenth ACM Conference on Hypertext and Hypermedia, pages 101–109. ACM Press, 1999.

[6] L. Calvi. Adaptivity in hyperfiction. In Proceedings of the Fifteenth ACM Conference on Hypertext and Hypermedia, pages 101–109. ACM Press, 2004.

[7] P. de Bra, A. Aerts, B. Berden, B. de Lange, B. Rousseau, T. Santic, D. Smits, and N. Stash. Aha! the adaptive hypermedia architecture. In Proceedings of the fourteenth ACM conference on Hypertext and hypermedia, pages 81–84. ACM Press, 2003.

[8] K. Gee. The ergonomics of hypertext narrative: Usability testing as a tool for evaluating and redesign. ACM Journal of Computer Documentation (JCD), 25(1):3–16, 2001.

[9] G. Genette. Paratexts Thresholds of Interpretations. Cambridge University Press, New York, NY, 1997.

[10] G. Golovchinsky. Queries? links? is there a difference? In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pages 407–414. ACM Press, 1997.

[11] R. E. Gomory and T. C. Hu. Multiterminal network flows. Journal of the Society of Industrial and Applied Mathematics, 9:551–570, 1961.

[12] J. B. Gutiérrez. Hipertexto en contexto. Signo y Pensamiento, XIX(36):36–40, 2000.

[13] J. B. Gutiérrez. Literatrónica: Hipertexto literario adaptativo. In 2o Congreso ONLINE del Observatorio para la CiberSociedad, Barcelona, Spain, 2002. Observatorio Para La Cibersociedad. URL: http://www.cibersociedad.net/congres2004/grups/fitxacom publica2.php?grup=2&id=242&idioma=es.

[14] R. Kendall. Toward an organic hypertext. In Proceedings of the Eleventh ACM Conference on Hypertext and Hypermedia, pages 161–170. ACM Press, 2000.

[15] G. P. Landow. Hypertext 2.0. Johns Hopkins University Press, Baltimore, MD, 1997.

[16] I. Snyder. Hypertext: The Electronic Labyrinth. New York University Press, New York, NY, 1999.

[17] S. P. Tosca. A pragmatic of links. In Proceedings of the Eleventh ACM Conference on Hypertext and Hypermedia, pages 77–84. ACM Press, 2000.