

Curveship: An Interactive Fiction System for Interactive Narrating

Nick Montfort

Program in Writing and Humanistic Studies
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
nickm@nickm.com

Abstract

Interactive fiction (often called “IF”) is a venerable thread of creative computing that includes *Adventure*, *Zork*, and the computer game *The Hitchhiker’s Guide to the Galaxy* as well as innovative recent work. These programs are usually known as “games,” appropriately, but they can also be rich forms of text-based computer simulation, dialog systems, and examples of computational literary art. Theorists of narrative have long distinguished between the level of underlying content or story (which can usefully be seen as corresponding to the simulated world in interactive fiction) and that of expression or discourse (corresponding to the textual exchange between computer and user). While IF development systems have offered a great deal of power and flexibility to author/programmers by providing a computational model of the fictional world, previous systems have not systematically distinguished between the telling and what is told. Developers were not able to control the content and expression levels independently so that they could, for instance, have a program relate events out of chronological order or have it relate events from the perspective of different characters. Curveship is an interactive fiction system which draws on narrative theory and computational linguistics to allow the transformation of the narrating in these ways. This talk will briefly describe interactive fiction, narrative variation, and how Curveship provides new capabilities for interactive fiction authors.

1 Curveship and Its Contexts

This paper addresses those interested in aesthetic and computational, work with language, whether or not they are familiar with interactive fiction or narrative theory. I describe the twofold motivation behind Curveship, explaining why I find interactive fiction compelling and why I find narrative variation a worthwhile capability for a literary computer system. I then sketch the way that Curveship works, pointing to aspects of the system that will, I hope, interest interactive fiction authors and also have relevance beyond interactive fiction.

Several histories of interactive fiction are available, including book-length (Montfort 2003) and briefer ones (Nelson 2001, Montfort 2007a). This paper focuses on how interactive fiction works, and on explaining its conventions, rather than on detailing the history of the first interactive fiction, *Adventure* (written in 1976 by Will Crowther and Don Woods), the “mainframe” games that followed, interactive fiction in the commercial marketplace (including the many notable contributions of Infocom), and the surge in development by individual, non-commercial interactive fiction authors. This paper also doesn’t provide any information about how to download or run interactive fiction, and very little about how to play it, although several FAQs and other resources are available online with that information (IFWiki 2009, Aiken 2008, Cadre n.d.).

After offering some of the motivation for this work, this paper provides a high-level introduction to Curveship and its capabilities. The details on narrative variation and on how it has been implemented computationally are available in the

author's dissertation (Montfort 2007b). While the interface to the system has been streamlined and other changes are being made to facilitate the use of Curveship by interactive fiction authors, the approach to narrative variation that is described there remains current.

2 The Interactive Fiction Tradition

Since the release of the first work of interactive fiction in 1976, those working in the form have developed conventions and traditions and have undertaken many different sorts of exploration and innovation.

A snippet of a transcript from this author's *Book and Volume* (2005) gives a basic idea of what interaction is like:

Your Apartment Building's Lobby

Clean, if not homey. The doorman waits in a Plexiglas cube that looks curiously bullet-proof.

>leave

Onward, into the intentional city. nWare crafted it well -- built from the ground up with security and scalability in mind, it's fully pedestrianized but hums, almost full-size, like a city center.

The night sky is completely cloudless. The dry warmth of the air hints at the desert that surrounds nTopia.

Abalone & Fifth

Your apartment building is on the southwest corner. Opposite it is a windowless, blocky building. On the southeast corner Pharmicopia sprawls. Otto's Automat is to the northwest.

A window pops up in the lower right corner of your nLap. In it, the uncanny red stapler wiggles its hinged head to the accompaniment of a cheery robot voice:

I see you're outside! Hey, the nice thing about this city is that it uses the grid system, so you can just go NORTH, EAST, SOUTH, or WEST, and you'll pretty much always end up somewhere else!

The window winks away.

>go to the automat

Otto's Automat

Someone, perhaps nWare's monarchically wealthy founder, shellf, really got into *Dark City*. There couldn't be any other excuse for this combination vending machine and cafeteria, even though it may strain to look modern and appetizing. Pixelated smiley faces in an urgent red decorate everything, even the edges of the food-dispensing wall.

The opportunity to avoid social contact during food selection is appealing.

>buy some food

Food does seem like a good idea. You select something palatable from the food-dispensing wall and begin to consume it, face flickering with mammalian happiness.

A guy you met during new employee reorientation -- Knut, you think his name is -- furtively joins you as you're dining. He looks even more feverish and high-strung than he did a few days ago.

"Hi. Hi again. Gotta tell. Tell you something."

Here, the player character — the one controlled by the player via textual commands — begins in the lobby of his apartment building. The player then types a command for that character to "leave," which he does, ending up outside at an intersection. Because the player character happens to be carrying his laptop computer, which is turned on, an annoying animated agent appears and dispenses some advice. The next command, "go to the automat," directs the character to enter a different building. ("Go northwest," "walk northwest," "northwest," or simply "nw" would have all done the same thing.) Finally, the character is commanded to "buy some food," which the character does — something that is only possible because the character has his electronic cash device, isn't sated, and is in place where food is sold. Although these are not dramatically interesting actions, they, in

this game and in others, can be part of figuring out the nature of an unusual world.

Of course, not every input works as smoothly as these did in this context. Actions often work if they are warranted by the situation in some way; so “buy some food” is appropriate in an automat while “donate my organs to science” isn’t. A character can be commanded to walk in a particular direction, or to enter a building that is nearby, but not, usually, to “go to Tangiers.” And there is usually no need for fine-grained positioning or describing the manner in which an action is done, so instead of issuing the command “hop spryly over to the coffee table” to prepare for setting down one’s mug, it’s fine to just go directly to typing “put my mug on the coffee table.”

Moving a character around using compass directions is a very notable convention originating with *Adventure*, although there were other ways to get around in that game. However it’s done, traversing a virtual space is very important to interactive fiction.

There are four important characteristics of interactive fiction that make it interesting from a research standpoint as well as from the standpoint of poetics. A work of interactive fiction is:

- A limited domain that serves as a simulated “microworld.” It has a complete model of the things that can be manipulated in the simulation and can be usefully talked about.
- A dialog system. Natural language is provided as output, and the system accepts commands that, although simple and short, are instances of English text.
- A computer game, providing enjoyment and fun. Although not the preeminent form of computer entertainment today, as it was around 1980, interactive fiction is something that many people find enjoyable and interact with for its own sake.
- A form of aesthetic expression and literary art. As with any form or medium, only a few use a significant amount of this potential. But the computational, literary nature of interactive fiction gives it the capability to do aesthetic work that could not otherwise be done.

Since many people don’t realize that interactive fiction extends beyond the cave setting and fantasy

genre, it’s worth mentioning a few examples of work from the last few years, work that gives an idea of the range of interactive fiction today — all of which is available for free download and easily found online:

Anchorhead, by Michael Gentry, 1998: An expansive interactive fiction with deep secrets and action that runs over several days, inspired in tone and style by H. P. Lovecraft.

Bad Machine, by Dan Shovitz, 1998: Manifesting itself as confusing a mix of status reports, error messages, this interactive fiction takes place in a strange robot-run factory.

Narcolepsy, by Adam Cadre, 2003: A seemingly contemporary, ordinary interactive fiction that branches hilariously into strange genre scenarios.

Slouching toward Bedlam, by Star C. Foster and Daniel Ravipinto, 2003: A steampunk science fiction piece set in an asylum and involving technological and kabbalistic themes.

Savoir-Faire, by Emily Short, 2002: The return to a childhood home provides opportunities to remember the past and requires that the player figure out an intricate system of sympathetic magic.

Spider and Web, by Andrew Plotkin, 1998: A science-fiction spy thriller that has the player reenact past events to the satisfaction of an interrogator.

Interactive fiction as it exists now is a type of virtual reality, a simulation of not only a space and the characters and things in that space but also of physical and metaphysical laws that obtain in a world. Furthermore, it’s a virtual reality that works well, one in which conventions have evolved about the level of abstraction and the types of commands that will work. An effective way of interacting has been negotiated.

Although more could be done to better simulate a world and to better understand language in interactive fiction, the Curveship project has a different goal. Curveship is being developed to add to interactive fiction’s well-established capability for simulation a new capability for narration, one that will allow the telling to be parametrically varied.

3 Narrative Variation

For more than three decades, interactive fiction programs have simulated fictional worlds. By allowing control over settings, characters, and the

incidents that happen, they have provided very useful facilities. However, literary works are not powerful and compelling merely because of what happens in them. They also rely on these events being told in an interesting way, on the different types of narrating that can be done. The interactive fiction system I am describing, *Curveship*, uses natural language generation to allow the narrating to be varied parametrically. To understand why this is a significant capability, it is worth turning to non-digital novels, stories, and narrative poems to see how they accomplish their effects.

We may consider different novels, stories, and poems to be “great” — powerful, affecting, transforming, deeply pleasing to read — but whichever ones we prefer, it is unlikely that we appreciate them simply because of what happens in them. The way these events are narrated is also important. A paraphrase or summary is generally not considered to be as interesting as is a great work of literature, even an ancient one. A timeline of events would hardly compare to *The Odyssey*, in which Odysseus tells some of the events himself, in which he weeps as he hears a bard, who does not know Odysseus’s identity, relating the events of the Trojan War and his own exploits to him. This is not to say that there can be no interesting retellings of *The Odyssey*, only that any telling will be interesting or not based on how the narrating is done.

The study of narrating, of how the same underlying events can be told in different ways, has been undertaken systematically in the field of narrative theory or *narratology*, in which the distinction between story/content and discourse, between that

which is *narrated* the *narrative* itself, has been central. Specifically, the model that Gérard Genette presents in *Figures III*, translated into English as *Narrative Discourse* (Genette 1980) and later revised in *Narrative Discourse Revisited* (Genette 1988), has provided the basis for narrative variation in *Curveship*.

A variant of a simple story given as an example by E. M. Foster is represented in figure 1. There are five underlying events: The death of the king, the grieving of the queen, the death of the queen, the usurping of the throne by a clown, and the laughing of the jester. These can be told one after another in their chronological order, as the top part of the diagram shows. But it is also possible to narrate the same underlying contest by saying “The king and queen died. The jester laughed — after the clown usurped the throne.” This telling represented in the bottom part of the diagram, and corresponds to changes in three of Genette’s categories: *frequency* (whether there is one telling per event, one for several events, or several for one event), *speed* (how rapidly or slowly events are related), and *order* (the sequence in which events are represented as compared to their chronological order in the story world). In this case, the king and queen’s death are both narrated with a single statement, a change in frequency; the queen’s grief is skipped over as rapidly as is possible and thus omitted entirely, a change in speed; and the clown’s usurping of the throne is mentioned last, after the jester’s laughter, which it apparently occasioned — a change in order.

Genette describes several other categories of variation, two of which are important for this paper. The *time of narrating* describes the temporal relationship between the narrating and the events of the story. For instance, in “I was driving down the road and it started raining frogs” the narrating is happening after the events, but a different (and still perfectly plausible) telling of this story, “So I’m driving down the road and all of the sudden it starts raining frogs,” the narrating and the events take place at the same time, giving a more immediate feel to the narrative. We could gloss this different as one of “past tense” and “present tense,” but this simple reference to grammar breaks down as a story gets more complex. If the narrator-character were to continue by noting “I had just had the wiper blades replaced” in the first case and “I just had the wiper blades replaced,” the story would no

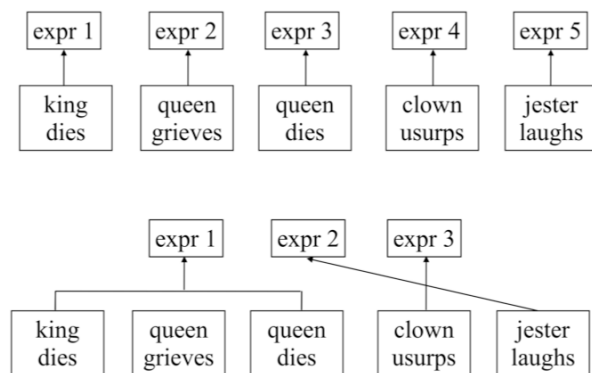


Figure 1. The same underlying events can be represented in a straightforward chronological way (above) or with different frequency, speed, and order (below).

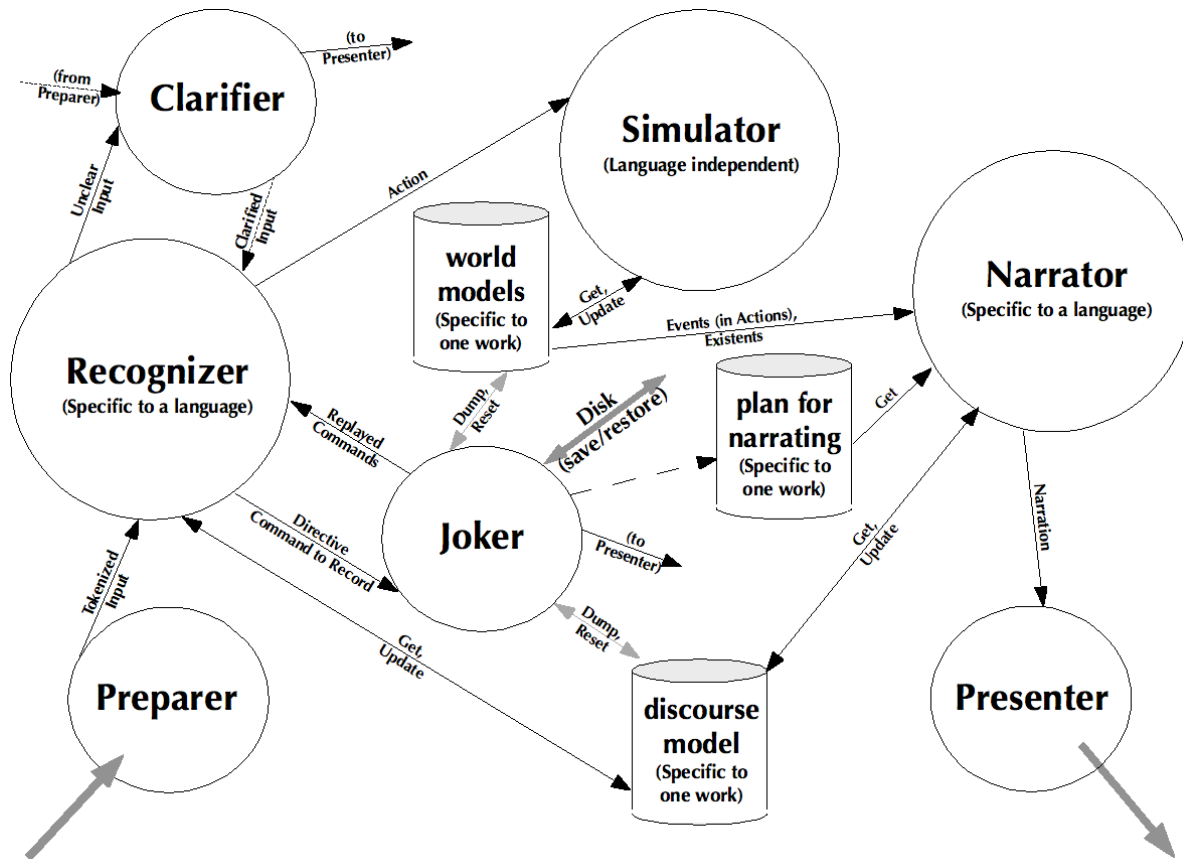


Figure 2. The architecture of Curveship. Each module is responsible for one more or less complex function; for instance, the Joker allows for *save*, *restore*, *restart*, and similar manipulation of the game state. The Simulator determines what events transpire in the IF world, while the Narrator deals with how to represent those events.

longer be entirely in the simple present or simple past. The important difference here, although it is reflected in the grammar, is actually a narrative one.

Focalization, briefly, describes the way that the information available to the narrator is regulated. If the narrative follows a character and tells us only what that character knows, it is focalized by that character. Whether the character is referred to in the main level of the narrative as “I,” in the third person (as in a typical Hemingway story), or even as “you” (the standard case in interactive fiction) is a separate matter. Specifically, that has to do with who the narrator and naratee are and if there are characters within the story who have this role.

4 The Architecture of Curveship

State-of-the-art IF systems (including TADS 3 and Inform 7) have innovated in many ways, but

they are similar in offering two main modules, the “parser,” which deals with recognizing intended actions based on typed user input, and the rest of the program, which handles both the simulation of the IF world and the narrating of events and description of existents in that world.

Curveship has a parser as well (the Recognizer) but, as shown in figure 2, it is further separated into modules that deal with different functions the interactive fiction system and program have to carry out. Significantly, it has separate Simulator and Narrator modules. The Simulator is potentially independent of the human language of a particular interactive fiction, although Curveship has only been implemented in English as yet. It updates the world models to reflect the new state of the underlying simulated world and the new theories that characters have about this world. Then, the Narrator module, which is quite specific to a particular human language, builds a narrative reply using a

world model and a plan for narrating. The Simulator is the only module that updates the world models. Similarly, the discourse model is written only by the Recognizer (which updates this model to reflect the user's contributions to the discourse) and the Narrator (which produces the system's contributions to the discourse and updates the model to reflect these).

Curveship's somewhat unusual name is meant to call attention to how the system models the essential qualities of variation — the curve of a story through its telling — just as friendship and authorship represent the essence of being a friend and author.¹ The word “curveship” was coined by Hart Crane (1899-1932) in the last line of his poem “To Brooklyn Bridge,” in which he addresses the bridge: “And of the curveship lend a myth to God.”

5 Order and Time of Narrating

The order of events as narrated does not have to correspond to the order of events in a fictional, simulated, or historical world. Genette represents the order of events in the narrating as a sequence, of the form “3451267,” but he also notes that events can be reordered in many different ways, for different purposes and to different effects. For instance, in “3451267,” the earliest two events, “12,” may have been narrated as what is commonly called flashback (which Genette calls an analepsis). But perhaps not: perhaps “345,” “12,” and “67” all fell into different categories, and the narration was done according to these categories — using syllepsis, in Genette's system. Or, perhaps the events have been jumbled at random to confuse the reader about their temporal relationship; this is called achrony. Cue words and tense will be used differently in these three cases, so “3451267” is not an adequate representation when text is to be generated, rather than just analyzed.

Instead of representing the order of events in the narrative as a sequence, Curveship uses an ordered tree representation called a reply structure. It describes not only the sequence of events but also which level each event is at and what its relationship is to the level above. To determine the tense, the system uses a theory that relates how three

¹ This may seem like an obscure name, but at least it's better than the name the system previously had, during the main phase of my research work: “nn.” People often couldn't even identify this as a word, whether it was spoken or written.

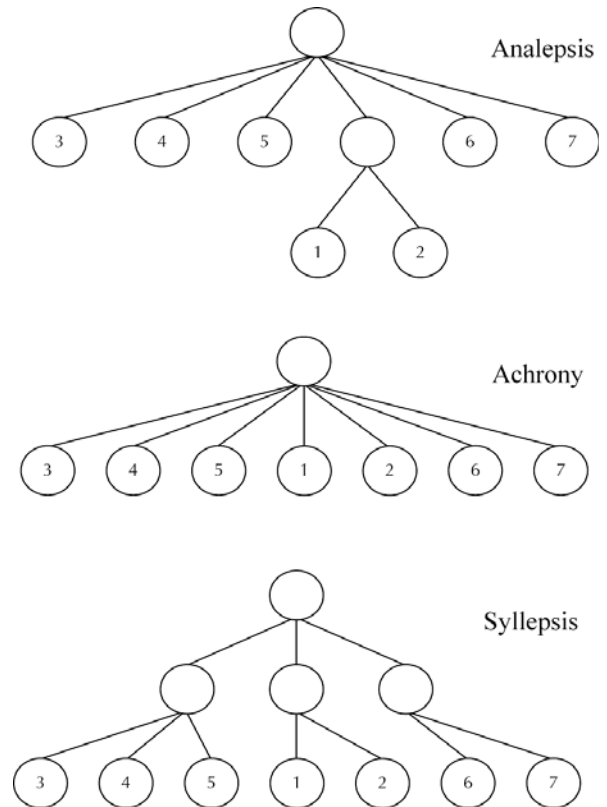


Figure 3. The reply structures corresponding to three different orderings, all of which would look the same if a simple sequence were used as a representation.

points in time — speech time (S), reference time (R), and event time (E) — correspond to a particular grammatical tense (Reichenbach 1947). Event time is supplied by the simulator; the other two times are determined based on the plan for narrating and the reply structure as text generation are done. The reply structure representation allows for different orderings to be composed, so, for instance, within a flashforward, the events can be jumbled achronously, and within each sylleptic category the narration can be done in a different temporal way.

6 Focalization

Curveship implements a system for changing focalization based on Marie Laure-Ryan's concept of a Fictional Actual World which the reader re-centers upon (Ryan 2001). In the formulation of this concept for interactive fiction, it is useful to consider an Interactive Fiction Actual World that

represents what is actual, or real, to the characters in the game. Each character, then — each potential focalizer — has his or her own world model, a theory of this world which may be mistaken and almost certainly is partial. The Narrator, then, never even sees the underlying simulation, but instead relates events based on the focalizer’s current theory of the world.

Because the Narrator may tell about things that happened before the current state of the world, each focalizer maintains not only a current theory of the world but also a history of how the world appeared in the past.

7 Text Generation in Curveship

The Narrator, which does text generation in Curveship, is organized into a standard three-stage pipeline. First comes the highest-level operation of content selection and ordering, which is done by the Reply Planner (essentially a document planner, but here part of a discourse is being planned). Then, the Microplanner determines the grammatical specifics of the output based on the plan for narrating. Finally, the Realizer accepts the paragraph proposals from the Microplanner and produces a string.

The problem of authoring for generation is a difficult one. Interactive fiction authors would like to be able to write as they do now, simply associating strings with objects and events. This representation is not suitable for the generation task, however. Something more general is needed to allow narrative variation to be automatically produced.

Advanced research and commercial text generation systems use highly abstract representations of sentences (different ones for each system) to allow text to be flexibly transformed, aggregated, and changed in tense, aspect, and person. While the power of this approach is unquestionable, taking this direction is also unsuitable, because it would require a tremendous investment on the part of authors, who would spend perhaps a hundred times the time and effort to create the same textual output that they could jot off in the typical interactive fiction system. It is unlikely that anyone would undertake this voluntarily, and, if people did, it would almost certainly disrupt the authorship process.

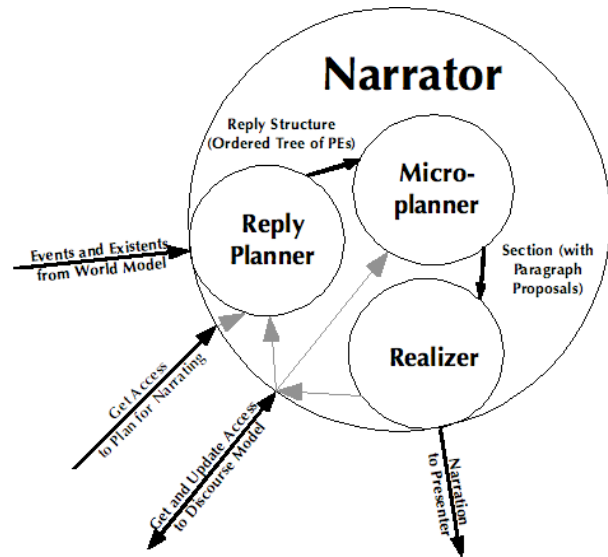


Figure 4. The Narrator module uses a standard three-stage pipeline for text generation.

As a compromise, Curveship uses a string-with-slots representation that offers significant flexibility in generation without the extreme complexity of most sentence representations. It allows authors to “cheat” and indicate that something should be treated as an entity in the discourse even if there is no model of it in the simulation. For instance, the text at the beginning of *Adventure* can be generated from the following strings:

```
'S_FC V_stand_PROG at the_end of
a_road before
a_small_brick_building'
'a_small_stream V_flow_S out of
the_building and down a_gully'
```

The first slot, S_FC, indicates that the focalizing character is to be named there (pronominalized if appropriate) and will be the subject of the sentence. The next, V_stand_PROG, says that the verb “stand” is to appear in the progressive. It is not necessary to specify the number; without such a specification, the verb will agree in number with the subject. The rest of the first string looks ordinary, except that noun phrases have been connected with underscores. This indicates that they should be treated as entities in the discourse even though they are not simulated: The system will, for instance, output “a road” the first time around and, since the road is then given in the discourse, it will

output “the road” afterwards. Finally, in the second string, there is the slot V_flow_S. The subject of the sentence is not indicated, but it is not necessary, since the “_S” indicates that the verb “flow” should be output in the singular.

Depending on the plan for narrating and the state of the discourse, this can produce:

You are standing at the end of a road before a small brick building. A small stream flows out of the building and down a gully.

As well as:

You were standing at the end of the road before the small brick building. The small stream flowed out of the building and down the gully.

Along with more exotic strings that result from unusual narrative settings and the use of text filters.

References

- Aiken, Jim. “Getting Started with Interactive Fiction,” 2008. http://www.musicwords.net/if/if_getting_started.htm
- Cadre, Adam. “Interactive Fiction — Getting Started.” <http://adamcadre.ac/content/if.txt>
- Genette, Gérard. *Narrative Discourse: An Essay in Method*. Trans. J. E. Lewin. Ithaca, NY: Cornell University Press. 1980.
- Genette, Gérard. *Narrative Discourse Revisited*. Trans. J. E. Lewin. Ithaca, NY: Cornell University Press. 1988.
- IFWiki. “Interactive Fiction FAQ,” 2009. <http://ifwiki.org/index.php/FAQ>
- Montfort, Nick. *Twisty Little Passages: An Approach to Interactive Fiction*. The MIT Press, 2003.
- Montfort, Nick. “Generating Narrative Variation in Interactive Fiction.” Dissertation, University of Pennsylvania, 2007a.
- Montfort, Nick. “Riddle Machines: The History and Nature of Interactive Fiction.” In *A Companion to Digital Literary Studies*, pp. 267–282. Editors, Ray Siemens and Susan Schreibman. Basil Blackwell, 2007b.
- Nelson, Graham. “A short history of interactive fiction.” *The Inform Designer’s Manual* section 46, 4th edition, 2001. <http://www.inform-fiction.org/manual/html/s46.html>
- Reichenbach, Hans. *Elements of Symbolic Logic*. New York: Macmillan. 1947.
- Ryan, Marie-Laure. *Narrative as Virtual Reality*. Baltimore: Johns Hopkins University Press. 2001.