

GENERATING NARRATIVE VARIATION IN INTERACTIVE FICTION

Nick Montfort

A DISSERTATION

in

Computer and Information Science

Presented to the Faculties of the University of Pennsylvania
in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2007

Mitchell P. Marcus, Supervisor

Gerald Prince, Supervisor

Rajeev Alur, Graduate Group Chair

Acknowledgments

I am extremely fortunate to have had the guidance, encouragement, and close attention of my advisors, Mitch Marcus and Gerry Prince, as I worked on this project. I also am grateful for the help I have received from members of my committee: Aravind Joshi, Mark Liberman, Fernando Pereira, and Marie-Laure Ryan.

I appreciate many discussions and interactions that have been of specific help in developing this dissertation work. Thanks to Sean Barrett, for several conversations about various aspects of IF and my work on this system; John Blitzer, for looking carefully at the system's output; Jim Carpenter, for discussions about object-oriented creative text generation; Nihkil Dinesh, for discussion of the modalities of narrating and simulating and for his careful consideration of the system's output; Jeremy Douglass, for his work on and discussion of IF generally, but particularly for discussing with me the distinction between different functions of the player character; Jason Dyer, for comments on my IF system as I was developing it; David Herman, for his work editing some of my other writing about IF and narratology, which also helped me to improve this document; Stuart Moulthrop, for collaborating with me on an article on IF and for his continued encouragement and IF-related discussions since the days of *Winchester's Nightmare*; Daniel Ravidpinto, for discussion of the system and the method of specifying variable text; Dan Shiovitz, for numerous discussions about IF and the system and for collaborating with me on *Mystery House Taken Over*, where we explored the content/expression distinction through creative IF development; Emily Short, for numerous discussions about IF criticism, theory, and system-building, and for collaborating with me on *Mystery House Taken Over*; Marc Simpson, for discussing my IF system with me from the standpoint of an IF author and system designer; and Rob Wheeler, for discussions of the nature of the player character.

I am also grateful for my collaborators on *Grand Text Auto* — Mary Flanagan, Michael Mateas, Scott Rettberg, Andrew Stern, and Noah Wardrip-Fruin — who have helped me in the blogosphere and beyond. I appreciate ifMUD, which is a generally good group of people as well as a particularly apropos community for someone undertaking interactive fiction research.

On my path to Penn, I had plenty of chances to learn about interactive fiction outside of the academic framework, but I also had the benefit of formal education that involved writing three related theses for previous degrees. I thank my previous advisors and committee members, who helped me to better understand language, literature, and IF: John Slatin and Michael Benedikt at the University of Texas at Austin; Justine Cassell, Glorianna Davenport, and Janet Murray at MIT; and Robert Pinsky and David Ferry at Boston University. At Penn, I am very glad to have had the opportunity to work with and learn from Michael Kearns. Thanks to the Institute for Research in Cognitive Science and Laurel Sweeney for the pleasant work environment.

Many other people have provided perspectives, support, encouragement, and collaborative effort on related creative and research projects. Some offered me valuable contexts to present and discuss my work related to interactive fiction, which I greatly appreciate. At the risk of omitting many names, I will thank Espen Aaeseth, Norm Badler, Charles Bernstein, Ian Bogost, Robert Coover, Roderick Coover, Peter Decherney, Al Filreis, William Gillespie, Kenneth Goldsmith, Terry Harpold, Fox Harrell, Matt Kirschenbaum, Stephen Lane, Bob Perelman, Peter Stallybrass, Brian Kim Stefans, Stephanie Strickland, Partha Pratim Talukdar, and Jill Walker.

Last but not first, thanks to Hanna and to my family.

Abstract

GENERATING NARRATIVE VARIATION IN INTERACTIVE FICTION

Nick Montfort

Mitchell P. Marcus and Gerald Prince

A general method for the generation of natural language narrative is described. It allows the expression, or narrative discourse, to vary independently of the underlying events and existents that are the narrative's content. Specifically, this variation is accomplished in an interactive fiction (IF) system which replies to typed input by narrating what has happened in a simulated world. IF works have existed for about 30 years as forms of text-based computer simulation, instances of dialog systems, and examples of literary art. Theorists of narrative have carefully distinguished between the level of underlying content (corresponding to the simulated world in interactive fiction) and that of expression (corresponding to the textual exchange between computer and user) since the mid-1960s, when the field of narratology began to develop, but IF systems have not yet made use of this distinction. The current project contributes new techniques for automatic narration by building on work done in computational linguistics, specifically natural language generation, and in narratology. First, types of narrative variation that are possible in IF are identified and formalized in a way that is suitable for a natural language generation system. An architecture for an IF system is then described and implemented; the result allows multiple works of interactive fiction to be realized and, using a general plan for narrating, allows them to be narrated in different ways during interaction. The system's ability to generate text is considered in a pilot evaluation. Plans for future work are also discussed. They include publicly released systems for IF development and narratology education, adding a planning capability that uses actors' individual perspectives, and adapting automatic narration to different sorts of interactive systems.

Table of Contents

1 Overview.....	1
1.1 Contributions to Narratology.....	2
1.2 Contributions to Computational Linguistics.....	2
1.3 Potential for Teaching about Narrative and for IF Development.....	4
1.4 Outline of the Dissertation.....	4
2 Insight from Narratology.....	7
2.1 The Uses of Interactive Fiction.....	7
2.2 A Perspective from Narratology.....	9
2.2.1 The Content/Expression Distinction.....	9
2.2.2 Events and Existents.....	10
2.2.3 An Extended Example of Narrative Variation.....	10
2.2.4 A Mapping to Interactive Fiction.....	13
2.3 Other Benefits for Research and Creative Practice.....	14
3 Review of Related Work.....	16
3.1 Existing Interactive Fiction Systems.....	16
3.2 Story Generation.....	17
3.3 Interactive Drama and an IF Research System.....	19
3.4 Commercial Games.....	22
3.5 Other Interactive Natural Language Generation and Dialog Systems.....	24
4 Steps toward a Potential Narratology.....	26
4.1 Theorizing Interactive Fiction.....	26
4.2 Interactive Fiction and the Interactor.....	29
4.3 Sessions, Interactions, Traversals.....	31
4.4 Cycles, Exchanges, and the IF World.....	32

4.5	Initial Situation to Final Situation, Prologue to Final Reply.....	36
4.6	Player Characters, Non-Player Characters, and Other Persons.....	39
4.7	World, Rooms.....	41
4.8	Diegesis, Hypodiegesis, and Extradiegesis.....	42
4.9	Winning and Losing.....	45
4.10	Puzzles and Their Solution.....	47
4.11	Typology of Interactive Fiction Outputs.....	51
4.11.1	Reports.....	52
4.11.2	Replies.....	53
4.11.3	Generating Diegetic, Interaction, and Narration Texts.....	57
4.12	Further Steps.....	58
5	A New IF Architecture.....	60
5.1	Elements of the World Models.....	60
5.1.1	Modeling Existents.....	62
5.1.2	The Tree of Existents and Relations.....	64
5.1.3	Modeling Events.....	67
5.1.4	Modeling Actions.....	70
5.2	IF Actual World versus Focalizer Worlds.....	72
5.2.1	Updating Focalizer Worlds.....	74
5.2.2	Use of Focalizer Worlds for Narrating.....	75
5.2.3	Potential Use of Focalizer Worlds for Planning.....	79
5.2.4	Determining What is Perceived.....	81
5.2.5	Zero Focalization.....	82
5.3	Discourse Model.....	84
5.4	The Modules.....	84
5.4.1	Preparer.....	84
5.4.2	Recognizer.....	85
5.4.3	Joker.....	85
5.4.4	Clarifier.....	85
5.4.5	Simulator.....	86

5.4.6	Narrator.....	86
5.4.7	Presenter.....	87
5.5	What an Author Writes for a Particular Work.....	87
6	Reply Planning and Variations in Order.....	90
6.1	Narrative Tense.....	93
6.2	Pure Order of Events in the Narrative.....	94
6.3	Pure Ordering of Events in the Narrating.....	97
6.4	Time and Grammatical Tense in Ordering Events.....	100
6.4.1	Chronological, Simultaneous, with Time Words.....	104
6.4.2	Chronological, Subsequent, no Time Words.....	105
6.4.3	Achrony, Simultaneous.....	105
7	Reply Planning and Variations in Speed, Frequency, Mood, and Voice.....	106
7.1	Speed.....	106
7.2	Frequency.....	108
7.2.1	A Connection between Order and Frequency.....	109
7.3	Narrative Mood.....	111
7.4	Narrative Voice.....	111
7.5	Interactive Mood and Voice?.....	113
7.6	Other Parameters of the Plan for Narrating.....	114
8	Microplanning and Surface Realization.....	116
8.1	Microplanner.....	116
8.1.1	Mapping the Narrative to the Grammatical.....	117
8.1.2	A String-with-Slots Representation for Building Sentences.....	118
8.1.3	Section, Paragraph, and Sentence Representations.....	121
8.2	Realizer.....	122
9	Creative Work Using Narrative Variation.....	124

9.1	Distance in <i>Lost One</i>	125
9.2	Baseball Narration in <i>Oddly Angled</i>	130
10	Pilot Evaluation.....	134
10.1.1	Naturalness.....	135
10.1.2	Agreement about Events.....	139
10.1.3	Ability to Communicate Events and Order.....	139
10.1.4	Comments from Annotators.....	139
10.2	Lessons for IF Output Evaluation.....	140
11	Discussion.....	142
11.1	Consideration of the Project and Early nn Development.....	142
11.2	Summary of Advances.....	144
11.2.1	A Theory of IF Based on Narratology.....	144
11.2.2	A New IF Architecture Abstracting Content from Expression.....	144
11.2.3	A Working IF System, nn.....	144
11.2.4	Formalization of Narrative Variation.....	145
11.2.5	An Automatic Narrator Based on Genette’s Theories.....	146
11.2.6	Implementation of Standard IF.....	146
11.2.7	Implementation of New Pieces with Complex Plans for Narrating.....	146
11.2.8	Pilot Evaluation.....	146
11.3	Future Work.....	147
11.3.1	Digital Media / Digital Writing.....	147
11.3.2	Narratology.....	148
11.3.3	Dialog Systems / Text Generation / HCI.....	149
11.3.4	Artificial Intelligence.....	150
12	Glossary.....	152
13	Appendix A	
	Example Interactive Fiction Transcript	

From Adam Cadre' <i>Varicella</i> (1999).....	159
14 Appendix B	
Example Interactive Fiction Transcript	
From Emily Short's <i>Bronze</i> (2006).....	161
15 Appendix C	
Example nn Code	
<i>Cloak of Darkness</i>	163
16 Appendix D	
Example nn Code	
Excerpted from <i>Adventure</i>	168
17 Appendix E	
Example Plans for Narrating.....	183
18 Appendix F	
Evaluation Texts and Instructions.....	191
19 Bibliography.....	207
19.1 Works of Interactive Fiction.....	207
19.2 Other References.....	209

List of Figures

Figure 1: Interior structure of the world models component. The simulation runs based only on the IF Actual World, but the Simulator also updates all of the Focalizer Worlds. The Narrator then bases its narration on one of these Focalizer Worlds — in this example, that of focalizer 2. An external, “omniscient” Narrator would base its narration on the COSMOS’s Focalizer World, which models everything in the IF Actual World. The Joker’s ability to read from and reset the world models is not represented here.....72

Figure 2: The system architecture, with input on the lower left and output on the lower right. The smaller circles are simpler modules less relevant to the proposed research. Of the major modules, only the Simulator can write to the world models (the IF Actual World and Focalizer Worlds), while only the Recognizer (which parses and understands input) and the Narrator (which does all the preparation of output except for a final formatting step) access the Discourse Model.....83

Figure 3: The architecture of the Narrator module. A standard three-stage pipeline is used. The Reply Planner first converts events, existents and a plan for narrating into a reply structure, an ordered tree of proposed expressions that represents what is to be narrated and in what order. Next, the Microplanner accepts this and uses discourse information to do aggregation and referring expression generation, converting the RS into a list a paragraph proposals. The Realizer renders each sentence in these as text and outputs them, updating the Discourse Model as it does this.....87

Figure 4: Three simple reply structures: (a) represents a present-tense chronicle with time words used; (b) represents a similarly-ordered chronicle, but with no time words and told in the past tense; (c) represents a retrograde narration in the past tense.....99

Figure 5: Two reply structures: (a) represents achrony, with events represented in the present tense; (b) generates a present-tense narration of events 4, 5, and 6 with a past-tense analepsis (including 1, 2 and 3) embedded within it. Note that the sequence of events is the same in both cases; the ordered tree representation allows for the same sequence to be generated coherently in the latter case and confusingly in the former, with both corresponding to different types of order identified by Genette.....101

List of Tables

Table 1. Recognized inputs and the outputs that correspond to them may be diegetic or not.....35

Table 2: Ratings of naturalness from the two annotators, each normalized over the interval that annotator used in assigning scores. Order is not the only thing that varies, but it is shown here to characterize the texts.....135

Table 3: Whether events identified by annotators matched up and what chronological ordering was assigned to events. “_” indicates that one annotator did not mark something as an event when the other did. Numerals represent the chronological rank assigned, with ranks 10, 11, 12 ... indicated with A, B, C... The essential identification by each annotator of the prevailing order in each passage is indicated in the last column.....138

1 Overview

This project develops an abstract, general method for generating different sorts of narrative discourse — in the plane of expression — based upon a particular fixed representation of existents and events in the underlying content plane. This method has been implemented in a complete interactive fiction development system. The research contributions of the project are intended to be mainly at the interface between narratology and natural language generation. The project offers to narratology a precise formalization of certain concepts of narrative and an extension of narrative theory to account for one sort of interactive system that narrates. To computational linguistics, the project offers a way of mapping a higher-level narrative specification to a particular syntactical representation for generation, along with some insight into how narrative theory can inform the generation of narrative texts. The research system that has been developed, called *nn*¹, has the potential to be further developed into a tool for teaching about narrative, supplementing the study that students now undertake by writing about narrative and by reading or viewing narratives in different media. It has the potential to be released to the public and to be used to support new sorts of expressive, creative work in this form as well.

1.1 *Contributions to Narratology*

First, this document describes how existing concepts from narratology and existing ideas of narrative variation have been applied to interactive computer programs that generate narratives. Specifically, these programs are works of interactive fiction. They are best considered potential narratives — spaces of possible narrative — rather than narratives of the usual sort, so the application of narrative theory requires that certain additional distinctions be made, such as the distinction between narrative that is simulated and that which is not simulated and that between characters who are commanded via user input or not. Still, existing narrative theory is used as the basis for understanding narrative variation in IF.

¹ The name *nn* is meant to suggest the distinction between narrated events (content) and the narrating (expression) of them.

Second, different types of narrative variation have been specified precisely here so that they can be implemented computationally. Narratological concepts have been expressed with rigor and sometimes in a way that computer scientists would call formal, for instance, in a formal grammar to distinguish narratives from non-narratives (Prince 1982); in a grammar of American conversational stories (Polanyi 1985); and using the approach of computational narratology for the analysis of texts (Meister 2003). But to implement models of narrative variation for text generation, it is also necessary to have truly formal specifications of how different narrative discourses can be generated from the same underlying content. This project offers such formulations for each of the main categories of narrative discourse as discussed by Genette (1980, 1988) as they have been further developed and refined (Chatman 1978, Prince 1982, Bal 1997). The concepts of these theorists of narrative need some alteration, refinement, re-specification, and further formalization as they are applied to interactive fiction, since interactive computer programs that narrate are not the same as narratives. This work is undertaken here as well.

The narratological questions the project has sought to answer are: What is a *potential narratology* of interactive fiction, particularly with regard to narrative variation? How can specifications for narrative variation be expressed in formal, general ways?

1.2 Contributions to Computational Linguistics

Just as this project does not seek to document previously unseen or unanalyzed types of narrative variation, the project is also not directed toward the development of, for instance, new surface realization techniques. Rather, the contributions are intended to be in determining how reasonably high-level specifications of narrative variation can be used to constrain and assist in natural language generation, to help produce texts that are more fluent and aesthetically pleasing and which are narrated in a wider variety of ways.

Because of this focus, a standard three-stage pipelined architecture is the basis for the proposed system's Narrator² component. This natural language generation architecture and the basic functions of text generation are not, themselves, the core research questions of this project. The basics of text generation are given, as is, for instance, the ability to change tense and aspect

2 To avoid any confusion between the IF system's Narrator module as designed and implemented and, more generally, "the one who narrates, as inscribed in a text" (Prince 2003:66), the names of this module, and all modules and models of the system, are capitalized.

appropriately during generation to reflect different temporal relationships. The research to be done is in the gap between the thorough understanding of narrative variation in unilinear texts — an understanding that has been supplied by narratology — and the existing techniques and architectures for natural language generation — approaches which have been supplied for many other text-generating situations by researchers in computational linguistics.

Accomplishing some aspects of narrative variation is fairly simple. Other parts of this gap have taken significantly more work to close, as is described in the text that follows. But a systematic connection between narrative variation and text generation offers tremendous benefits for interactive fiction and for narrating systems in general. This project aims to make important research contributions in computing and narrative, and to begin a significant exploration into how computer systems of all sort (beyond interactive fiction) can narrate. By being geared toward a cross-platform, widely downloadable and installable system rather than an elaborate contraption that will only ever run on one system, this project has also attempted to provide the foundation for an important new aesthetic system and offer an example of how an academically-developed natural language generation system can, eventually, be widely deployed and used.

The essential computational linguistics question for this project is: Given a model of events and existents and a specification of how these are to be narrated, what are the new ways to accomplish content selection, content structuring, aggregation, and surface realization in terms of narrating specifically, ways that would not apply to non-narrative texts (for example, to the representation of sequences of events that are not linked by causality) or to the generation of narratives that are always told as uniform, unvarying chronicles?

1.3 Potential for Teaching about Narrative and for IF Development

This project has not yet resulted in a system that is simple to use for teaching about narrative or one that is easy for IF developers to use, and certainly has not yielded a system that can be carefully evaluated in either of these capacities. However, the project has built the foundation for the further development of educational and expressive systems. The basic implementation allows for a range of narrative variation that will be useful for students and teachers as well as IF authors. With refinement — including the addition of a capability for quickly generating non-interactive narratives for educational use and the incorporation of better interactive and

simulative capabilities for IF authors — the system could see a reasonable amount of adoption by users in two different domains. A released version could, of course, also be used as a platform for further research by others or as a stage in a story generation or other sort of pipeline.

1.4 *Outline of the Dissertation*

This document provides background; describes the theoretical work that has been done (both in formulating narratological ideas for use with IF in this system and in developing the architecture for an IF system); covers the existing implementation of the Narrator; and relates the results from a pilot evaluation of the system's language-generating capability.

Specifically, the next chapter further motivates interactive fiction as a domain for computational linguistics research. It continues to describe how one fundamental idea in narratology — that the content plane (or story) can be separated from the expression plane (or discourse) — can inform the development of interactive fiction and allow for narrative variation to be accomplished generally and flexibly.

Chapter 3 reviews related work in interactive fiction itself and in related fields: story generation, interactive drama, commercial games, and other NLG and dialogue systems.

In chapter 4 a narratological theory of interactive fiction is provided. Since an IF program is an interactive system capable of producing many narratives, that is, a potential narrative, a preliminary *potential narratology* of IF is described here. The theory encompasses, among other things, the elements of IF, the way that narrative levels map to levels of simulation and narration in IF, and a theory of puzzles. The chapter closes with a detailed analysis of IF outputs, undertaken to further inform the development of the Narrator. This section considers what specific functions are accomplished by reports and replies, the two types of textual output from the system.

Chapter 5 describes the IF architecture that is implemented in this project. The chapter begins with how the basic world model — used in both the IF Actual World and in the Focalizer Worlds, which represent actors' perceived worlds — is defined. This discussion covers the representation of the entities that exist in the world (existents) and that which happens in the world (events and actions). The system's modules are then discussed. While the Narrator is the important component from a text generation and narrating standpoint, the architecture itself is an

important innovation that divides content from expression and modularizes the functions of system so that it is possible to deal with the simulation and narration separately.

Chapters 6 and 7 cover the work at the heart of this project's computational contributions. These chapters explain how narrative variation can be specified at a high level — that is, how a compact and general definition can be used instead of an exhaustive and inflexible description of exactly what and how is to be narrated. This includes a strict, computationally implementable formalization of elements from every major branch of Gérard Genette's theory of narrative discourse. The workings of the first, highest-level component of the Narrator, the Reply Planner, are described in detail along the way. Chapter 6 focuses on a single particularly interesting category of narrative variation, variation in order. Chapter 7 covers the way that speed, frequency, and some qualities that fall within mood and voice are varied.

Chapter 8 deals with the remaining two components of the three-stage pipeline of the Narrator: the Microplanner and the Realizer. The string-with-slots formalism is described. This representation allows sentences to be very quickly specified at a level of abstraction that is fairly low, but which still allows for several sorts of variation, including changes in tense, aspect, and number.

Chapter 9 comments on transcripts from two short creative demos that have been written in the system. These show how more complex sorts of narrative variation can be assembled from the simpler dimensions of variation provided. These two more complex narrative styles involve increasing distance and generating a sort of baseball-commentator narration. These demos also show different ways that variation can be made to change based on user input in an interactive system.

Chapter 10 describes a pilot evaluation that was done and offers a discussion of the results and of the annotators' comments.

Chapter 11 concludes by considering the work that has been done, the challenges that were encountered, and what has been learned about the nature of the problem of narrative generation. Then, the work that has been completed and the contributions that have been made are summarized. The chapter concludes with plans for future work of several different sorts. Possible projects are categorized as more or less extensive.

A glossary with key terms in narratology and interactive fiction is included after chapter 11.

Although many of the terms there are also defined when they are first used in the dissertation, readers less familiar with narratology and IF may wish to look over this glossary first.

2 Insight from Narratology

2.1 *The Uses of Interactive Fiction*

Interactive fiction (often abbreviated IF) can be compelling from a literary and aesthetic standpoint. At its best, interactive fiction can provide transformative experiences that can help readers to understand the world from new perspectives. It is difficult to explain the literary and reflective qualities of IF without showing how particular pieces of IF actually function, just as it would have been difficult to explain a motion picture a hundred years ago without showing one. Just as cinema was initially dismissed as a novelty, as entertainment, and as unsuitable for artistic expression, interactive fiction is — when it is considered at all — often quickly dismissed both by those interested in commercial video games and by those interested in literary work. From the standpoint of a video game developer, player, or critic, interactive fiction is frequently considered a historical curiosity with no commercial potential, uninteresting because of its textual rather than graphical nature and the way it engages language rather than 3D representations. To many (but fortunately not all) in contemporary poetics and literature, IF is, ironically, the opposite — “just a game.” For these reasons, even within the narrow intersection of computing with literary and artistic practice, IF is underappreciated. Brilliant new works continue to be written, programmed, and published by IF authors; interactive fiction is written internationally in several languages and provides challenges and provocation to those who interact with it; and many people have undertaken useful studies of the form’s nature and its gaming, literary, and educational potential (see Jerz 2001 for a broad survey of the writing about IF up to that point). Despite all of this, IF has had hardly any support or recognition from institutions that have traditionally promoted literature.

A transcript of interaction is not adequate to show how the reader must figure out the simulated interactive fiction world, typing back to the program in a dialogue. To illustrate something about the variety and quality of recent interactive fiction works, however, the first two appendices provide short excerpts from transcripts of interaction with two IF pieces: Adam Cadre’s 1999 *Varicella* and Emily Short’s 2006 *Bronze*.

Among its other virtues, interactive fiction can serve as a useful context for computational linguistics research. As others who have undertaken text-adventure-based research recently have explained,

From the perspective of computational linguistics, the computer game setting is interesting because it naturally restricts what utterances the user will produce. For example, players will typically only refer to objects they can “see” in the simulated world. This simplifies the language processing tasks, constrains the way in which ambiguous inputs are to be interpreted, and allows the inference problems to scale rather well to larger game worlds. There is a natural notion of a context (both with respect to what has been said before and in the sense of being situated in the game world), and the world can be freely specified and tailored to different levels of complexity. (Koller *et al.* 2004:188-189)

These researchers go on to suggest the uses of interactive fiction systems in the lab as testbeds. But there has been little discussion in computational linguistics about how interactive fiction development systems are used and IF games are played *in vivo* and not just *in vitro* — by numerous people, worldwide, using many different platforms, with many different styles of interaction and initial expectations. Interactive fiction is a demanding aesthetic application. It provides a way to disseminate research widely, bring ideas about generation into the popular consciousness, and have one’s work tried and tested by a huge range of users. An innovation in the use of natural language generation in interactive fiction would be well-poised to reach other real-world systems beyond IF and to attain widespread use in a variety of contexts.

Interactive fiction can be beneficial in several other ways, although these are not the main focus of this dissertation. Works of interactive fiction and constructive approaches to interactive fiction (developing games, not just playing them) have been used in various sorts of educational contexts, to teach reading comprehension and other literacy skills; to assist students in language learning; and even to model scientific, historical, and cultural systems for learners to interact with. In the past, interactive fiction has been commercially viable as entertainment software, and it is possible — although far from certain — that a substantial market for interactive fiction might arise again. Interactive fiction is also useful for exploring underlying issues in digital media of many different sorts, whether the explorers are student game designers or digital media scholars. The design of graphical adventure games and other simulated worlds can be prototyped in all-text interactive fiction. Some issues are specific to language as it is used in interactive fiction, and some are specific to the visual presentation or real-time action that is important to many other digital media systems. Nevertheless, there are many issues that arise in designing systematic

worlds, creating virtual spaces, setting up interactions among characters, and narrating what has happened that are seen in interactive fiction and that also manifest themselves in other digital media systems.

2.2 *A Perspective from Narratology*

2.2.1 *The Content/Expression Distinction*

The basic distinction between *story* and *discourse* has long been noted in discussions of narrative. This distinction was the starting point of the project described here. The “content plane” of *story* has been discussed, since Aristotle, as *mythos*, *fabula*, *histoire*, and *narrated*; it is, essentially, what is told about. The “expression plane” of *discourse* has been framed as *logos*, *sjuzet*, *récit*, and *narrating*; that is, the telling itself. Because *story* has a commonplace meaning and *discourse* a computational linguistics meaning, both of which differ from the way the terms are used in narratology to make this distinction, the terms *content* and *expression* are used for the remainder of this dissertation. Ideally, terms that are less general than these would be used, since *content* and *expression* are not specific to narrative. They suffice to make this important distinction, though.

There are certainly some differences between the classical, Russian Formalist, and contemporary narratological ideas about the story/discourse distinction. Furthermore, calling attention to the content/expression distinction is not meant to suggest that these are the only two planes of interest: Genette identified three, *histoire* or story, *récit* or narrative, and *narration* or narrating (Genette 1980:27) and his discussion was neither the first nor the last to point to more than two planes of narrative (Barthes 1975:243, Bal 1997:9, Rimmon-Kenan 2002:3). Certainly, though, the idea that what is told about can be considered as distinct from the telling itself is not a particularly controversial one. In fact, the idea has been fundamental to narratology.

Two very short and simple narratives are enough to demonstrate one way in which the same underlying story can be expressed differently:

Narrative 1 — John ate a sandwich, and then he died.

Narrative 2 — John died after eating a sandwich.

While these sentences may not be of compelling literary value, the basic technique — changing the order in which events in a given temporal sequence are related — is important to the aesthetic

and rhetorical effect of more complex narratives and to ones of more literary interest. Even in these examples, many readers will find some higher-level differences in these two narratives: perhaps one provides a wry humor, or one suggests causality more strongly (although neither indicates this explicitly), or one suggests more of a plodding progression, and ending, of life. Even when the information conveyed is the same, the way it is told can be important.

The content/expression distinction, and anything like it, is notably absent in the architecture and knowledge representations of computer systems that generate narrative, as is discussed in the next chapter.

2.2.2 *Events and Existents*

The content plane can be seen to have two fundamental sorts of elements: *events*, which are things that happen, and *existents*, which are the entities in the story (Chatman 1978:43-145). Actors, physical objects, and places are all existents, for instance, while any change in the state of these is an event. An event may be caused by some actor within the story, or it may be a *happening* with no agent, such as “there was an earthquake.”³ This concept allows the content plane to be understood as being partitioned into (a) the state of all that exists in the story world, and (b) the changes in that state. The distinction between *events* and *existents* serves as the basis for the enriched world models of the proposed system.

2.2.3 *An Extended Example of Narrative Variation*

Before connecting these basic concepts from narratology to IF, it is useful to characterize how the main sorts of narrative variation manifest themselves in narratives of the usual sort, using a short but slightly more detailed example. Here is a simplified sketch of one chronological sequence of events as they happen in a particular story world, on the content plane:

Event 1 – A stranger assists the heroine.

Event 2 – The stranger falls in love with the heroine.

Event 3 – The heroine falls in love with the stranger.

Event 4 – The stranger loses his soul.

Event 5 – The stranger attacks one of the heroine’s friends.

3 A happening can be assigned a special agent such as “nature” or “the cosmos” if it is desirable for each event to have an agent.

Event 6 — The stranger attacks another of the heroine's friends.

Event 7 — The stranger attacks another of the heroine's friends.

Event 8 — The heroine tries to kill the stranger.

The existents mentioned here are the stranger, the heroine, and the three friends — all of them characters — along with the stranger's soul. Although there is some disagreement over how fine-grained events and existents should be in studies of traditional narratives, a more detailed telling would probably correspond to additional events and existents. For instance, if the heroine attacked the stranger with a sword in an attempt to kill him, and the stranger defended himself with a sword, the swords would be existents; particular moves during the swordfight might be events.

This sequence of events is obviously more complex and potentially interesting than is the sequence of John's sandwich-consumption and his demise. The underlying reversal in the nature of a character has repercussions in this sequence that may be compelling and perhaps emotionally wrenching.⁴

One way of narrating these underlying events would be to simply read the text provided in the event list, reading down from the first event to the last: *A stranger assists the heroine. The stranger falls in love with the heroine. The heroine falls in love with the stranger. The stranger loses his soul. The stranger attacks one of the heroine's friends. The stranger attacks another of the heroine's friends. The stranger attacks another of the heroine's friends. The heroine tries to kill the stranger.* This narrative is a chronicle in which events are arranged 1 2 3 4 5 6 7 8.⁵

But there are plenty of other ways to tell a narrative based on this content. For instance: *The heroine fell in love with a stranger. Then, the heroine tried to kill the stranger. The stranger had attacked three of her friends, you see.* This one has events arranged 3 8 (5 6 7), where the parenthesis are used to indicate that several events are grouped and told as one. There are several differences in this second telling. It is no longer in chronological order, for one thing. The heroine's trying to kill the stranger is narrated right after her falling in love with the stranger, and then the narrative looks back (in what is called a flashback, retroversion, or analepsis) to the earlier events in which the stranger attacks the heroine's friends. This and other sorts of variation in order are covered in

4 Those interested in how a basic plot like this one plays out over a lengthy narrative can watch season 2 of the US television series *Buffy the Vampire Slayer*.

5 Some argue that a mere report of events does not provide narrativity (see Ryan 2006:81), but this chronicle offers a useful contrast here; it is worthwhile to consider it a narrative.

chapter 6. Furthermore, these three events are not narrated individually (in which is called singulative narration) but all at once, a variation in frequency. The type of frequency in which several events are represented with one narration is called iterative. Frequency is dealt with in section 7.2. It is also notable that some events are omitted: This second narrative does not mention that the stranger also fell in love with the heroine or that the stranger lost his soul. This is a variation in speed, an extreme one. In this case, these events were narrated as quickly as possible, so that nothing at all was said. Of course, omitting information like this, or, more generally, providing more or less text, can have an effect on the reader's sympathies. Telling something in the narrative when some of the characters do not know it can heighten suspense and have other effects. Variation in speed is discussed in section 7.1. This narrative is also in the past tense, consistent with a telling that is happening after the fact, while the original one is in the present tense, as if it were being told simultaneously with the events. Also, there is a subtle difference at the very end. By concluding with "you see," the narrative makes reference to the one to whom the story is being told, the narratee. There is no explicit mention of the narrator (the "I"), but there is this "you," which reminds the reader that this is a narrative being told and makes the events less direct. The change in time of the narrating, and the explicit reference to the narrator, are variations in narrative voice, a category discussed in section 7.4.

Besides order, speed, and frequency (collectively referred to as narrative tense) and narrative voice, there is one more category of variation: narrative mood. The perspective from which the story is told, also called the focalization, is an important part of mood. The focalizer of the story determines what information is available to narrate, but does not have to be the "I," the narrator. This example content told from the perspective of the heroine could be quite different than the same content told with the stranger as a focalizer, even if neither of them are made into the narrative's "I". For instance, if the heroine was the focalizer, we might have access to her emotional state and thoughts, but not to the stranger's; the situation would be reversed if the stranger was the focalizer. Mood is covered in section 7.3.

Some more elaborate forms of narrative variation exist, and have been implemented in interactive fiction as part of this project, using lower-level elements of narrative discourse. The discussion of these is in chapter 9.

2.2.4 *A Mapping to Interactive Fiction*

A systematic way to add narrative variation is provided by the following view of IF:

- The story level, or the content plane, is seen as analogous to the world model or simulation in IF.
- The discourse level, or the expression plane, is seen as analogous to the interface in IF, which accepts commands from the interactor and provides replies.
- Existents and events at the story level correspond to actors, things, rooms, and events that are represented in software and are simulated in a world model.

This is not the only correspondence that can be drawn between narratological elements and IF elements, or even the obvious one. In narratology, the content plane is properly considered to underlie any representation, whether it is a text, a diagram, or a set of data structures in a computer program. So a mapping could reasonably consider the software world model to be a form of expression itself — an intermediate sort of expression, hidden from interactors but accessible to certain types of analysis. In this mapping, different IF authors and developers could be seen to create different software representations that could evoke or express the same underlying, abstract system, with this system representing a deeper sort of content.

This project does not focus on how a variety of different interactive fiction world models can represent the same underlying content, however. The question is, rather, how a given world model can be narrated differently within interactive sessions. Since theories of narrative have considered the question for decades with regard to story and discourse, the useful correspondences to draw for this purpose are between world model and content (or the narratological concept of story) and between interface and expression (or more specifically, discourse).

Interactive fiction produces texts that describe characters and objects even when these characters and objects are not simulated, that is, when they do not have a representation in the world model. An analysis that considers the cognitive effects of IF on the interactor the way an interactor imagines certain existents and events during the experience of a session, could uncover a set of events and existents that were imagined. This cognitive content would not necessarily correspond to what is represented in software. The disjunction between these two sets is very interesting to consider and can probably tell us a great deal about the relationship between

simulation and narration in digital media. But when attempting to automatically produce narrative variation, it makes sense to consider only those existents and entities that are explicitly represented in software, since they are the ones that can be directly worked upon by the narrating module of the program. For the purposes of this project, therefore, *existents* and *events* will be used in the strong sense to indicate that which is simulated as well as (potentially) narrated. Those terms will not be used to refer to anything that might be evoked by text but which lacks an internal, computational representation.

2.3 *Other Benefits for Research and Creative Practice*

The work described here was undertaken with two main goals, one related to research progress in the generation of narrative, the other directed toward progress in the creative use of computers — specifically, in interactive fiction development.

As previously described, this project has been an attempt to make meaningful contributions *both* to narratology and to computational linguistics, to show that interdisciplinary work in this intersection can be valuable to both of the intersecting disciplines, rather than just answering questions of interest to one of the disciplines involved. This involves extending the author's previous work on narratological approaches to interactive fiction (Montfort 2003b) by showing how narratology can inform creative practice and interactive narrative generation. It is hoped that these developments will also offer advantages to those seeking to generate narrative texts of other sorts, fictional and nonfictional, literary and otherwise, based on underlying computational representations of events and existents.

The project is also part of an effort to improve the state of the art in interactive fiction development so that authors will be able to program more complex and nuanced forms of narrative expression. The IF system developed will be, when released, the first system to allow authors to easily manipulate the telling of their worlds, allowing a number of literary techniques to be integrated into IF. This will have near-term benefits for existing IF authors and those who are now thinking about authoring games. The system should also eventually (because of the modularization involved) make the translation of interactive fiction, and new sorts of collaboration between different authors and programmers on individual projects, easier. It should also invite the participation of those whose literary interests do not lie at the level of creating

interesting events, but at the level of relating events in an interesting way: poets interested in narrative, fiction writers in the postmodern American tradition, Oulipian writers who appreciate narrative variations such as the ones in Raymond Queneau's *Exercises in Style*, and many others interested in narrative discourse.

Currently, IF is downloaded and read by at least thousands; millions experienced IF during the height of commercial distribution in the 1980s. With public readings, festivals, and other efforts to broaden the appeal of IF, and given the many international communities engaged with IF, and considering the continued use and spread of digital systems that use text interfaces, from cell phones with SMS to search engines, it is plausible to expect that several thousand people will interact with works created in nn within two years of its completion. These people will include students, poets, writers, editors, game developers, and others. Hopefully, they will appreciate the experience and will help to bring the ideas underlying the system into other domains, offering new techniques for new media art, literature, and gaming.

3 Review of Related Work

3.1 Existing Interactive Fiction Systems

The three state-of-the-art interactive fiction systems available for widespread use are Kent Tessman's Hugo, Michael Roberts's Text Adventure Development System (TADS), and Graham Nelson's Inform. TADS, released initially in 1987, is the oldest of these; version 3 was released in 2006. Hugo was released by 1995 and is now in version 3. Inform was first made available in 1993; by 1996 the system reached the significant plateau of the sixth version, Inform 6. In April 2006, a very different-looking Inform 7, based on Inform 6 but with new capabilities and a completely different syntax, was released.

All of these systems are cross-platform and include object-oriented programming languages with fairly complete capabilities for advanced IF development, either built in to the language or available through libraries. These three systems each offer some sort of multimedia capability, are compiler-based, and have well-developed, extensible parsers and world models. With the exception of Inform 7, all are C-like in their syntax. Inform 7, unusually, uses a very English-like syntax, one which has few antecedents — Metafor, perhaps, which builds program sketches from natural language to help beginning programmers (Liu and Lieberman 2005); the WordsEye system for building 3D scenes from natural language descriptions (Coyne and Sproat 2001); and the less graceful and less prose-like COBOL and SQL, which mimic natural language to some extent. Besides these three systems, there are research systems which have not been used to write and release full games, such as FrOZ; these are covered in section 3.4. There are also IF development systems which may be more welcoming to some authors or may facilitate more rapid development for some, such as ADRIFT and the older system Alan, but the “big three” systems are the most advanced and their capabilities are most interesting to consider.

While the big three systems all abstract input from simulation by providing a parser distinct from the world model, they generally do not abstract the output from the world model to any extent. Again, the only significant exception is seen in Inform 7, which does allow for some processing of the simulation's output prior to the actual printing of text. A rule can specify, for

instance, *Before printing the name of a woman, say "Ms "*. This facility allows control over certain sorts of expression at a higher level, allowing the enumeration of a list of items to be done in a certain specified order within a description. The introduction of any sort of post-processing at all is a significant step. Inform 7, however, does not provide for a full model of the expression plane and the dialogue, does not have first-order representations of events, and lacks an independent facility for narration which would allow, for instance, variations in order or frequency during the telling, so that two events could be aggregated or so that one might be narrated in flashback. Inform 7 is interestingly innovative, but it leaves much to be done in abstracting the simulated world from the output that describes it and that narrates the simulated events.

Even though existing IF systems have not provided facilities for varying the way that underlying events are told, there has been plenty of significant progress over the thirty-year history of IF. Will Crowther's initial formulation of rooms with exits, objects that could be carried, and proto-actors provided a very useful a starting point. Important advances in simulating actors, containers, and vehicles began in the late 1970s and were discussed in the literature (Lebling *et al.* 1979). These have been refined thanks to the development of many commercial IF systems, which were followed by several free systems. The properties of an IF object and workings of an IF system were defined and documented particularly very well by Nelson during his many years of work on Inform (Nelson 2001). This project's focus on the level of narrating has only been possible because of the great amount of IF system development work that has been done on the level of simulation, uncovering simple, powerful ways to create textual virtual environments.

3.2 Story Generation

Story generation is a thread of research that has engaged cognitive questions as well as questions of narrative and literary art. Generating stories is not an inherently interactive endeavor and it does not require the same sort of simulation that interactive fiction does, but it is an obvious place to look for work on modeling the content plane and the expression plane of a story independently, in ways that allow for narrative variation. As it happens, however, the focus of story generation research has been elsewhere, at least until very recently.

A significant early story generator is TALE-SPIN (Meehan 1976), which used a conceptual

dependency representation (Schank 1975) to generate and narrate the actions of characters in a simulated world. The idea was to generate events which were themselves interesting; once these were generated, they were narrated in an unvarying, direct way. There was at least one story generation effort that was contemporaneous with or earlier than TALE-SPIN: a project on automatic murder mystery novel writing (Klein *et al.* 1973). Numerous other systems followed these, each introducing some new techniques and reconceptualizing the story generation problem in some way, as explained in (Montfort 2005):

Michael Lebowitz's 1984 UNIVERSE refined this approach and enhanced the representation of characters (embellishing certain stereotypes) to generate soap-opera narratives. MINSTREL (Turner 1994) was a similar system to generate Arthurian tales; it was able to get 'bored' and move on to other topics. A recent automatic storyteller is BRUTUS (Bringsjord and Ferrucci 2000), a system that uses a formal model of betrayal and has sophisticated abilities as a narrator.

Since the development of BRUTUS, the system MEXICA (Pérez y Pérez and Sharples 2001) has been created to embody a cognitive model of the writing process and use abductive reasoning to create explanations for complicating events. Almost all of these systems, however, work at the level of content to attempt to produce a sequence of interesting events. BRUTUS is an exception in this regard, but by taking a lexically-oriented approach of building stories from grammars that govern the text on different scales, it deals with expression and content at once rather than considering these independently. Another exception, and an earlier grammar-driven system, was TELLTALE (Correira 1980). As four researchers from Spain and Germany reported (Gervás *et al.* 2006), however, none of this work, and no other discernible work in story generation, deals with variation in the narrative discourse independent from the content plane:

The use of inflexible techniques for Natural language rendering of automatically generated narratives might as well be due to the fact that very few attempts exist to make Natural Language Generators fit for (literary) narrative input. The only Natural Language Generator that explicitly aims at this goal is STORYBOOK (Callaway & Lester, 2002). However, STORYBOOK uses a proprietary input representation, the so-called *narrative stream* format, and, to our knowledge, there are no interfaces to the output of implemented Story Generators. The input to STORYBOOK, then, is mainly encoded by hand.

Narrative discourse techniques such as large-scale ellipsis, flashback, repetition, summary, or changes in perspective are not used explicitly or purposefully in Story Generation. In our research, we have not yet encountered any system that would include a narrative discourse middleware able to produce variation at this stage.

Previous work on narrative prose generation (Callaway and Lester 2002) has resulted in an ability to produce high-quality narrative text, and this will inform some of the additional work to be

done in microplanning in the current project. But this work incorporated only a basic narrative planner, focusing on the later stages of generation (and the generation of dialogue, which is not considered in the current work) rather than the earlier content ordering and structuring that is necessary, for instance, to achieve the typical sorts of variations in order that are seen in narratives. Aspects of this narrative generation research, along with recent work on generating aesthetic text that incorporates the ability to use analogy and metaphor in describing entities (Pereira *et al.* 2006), could certainly be integrated into the architecture for varying the narrative discourse and generating narrative text that is proposed here.

Notably, an architecture for dealing with narrative level has been described (Lönneker 2005). Narrative level deals with the nesting of stories and narrators, and is aspect of narrative not addressed in the current project. This proposed architecture does model the content plane and expression plane independently. If research using this other system is fruitful, the insights that are gained there could be incorporated into an interactive fiction system to provide better facilities for representing existents and events across levels of narration and simulation.

3.3 Interactive Drama and an IF Research System

The most significant official academic project that was closely related to interactive fiction development was the Oz Project at Carnegie Mellon University, headed by Joesph Bates. While Oz was an “interactive drama” project, indicating a different conceptual framework and slightly different aesthetic goals, the early text world systems of the Oz Project share many qualities with text-based interactive fiction. They have underlying simulated worlds, a textual exchange, and the ability to understand commands that are addressed to a character. These systems also had a different emphasis — one that highlights realistic and dramatically compelling interactions among characters. They meet all the formal criteria for IF, however.

The first technical reports from the Oz Project were published in 1989. After initially working with all-text systems, the project grew into graphical systems and developed the Woggles in 1992. The new system used simple animated graphics and was a turn away from concerns with linguistic expression. In 1996, Bates founded the company Zoesis and turned to working further on the Woggles side of the project in this commercial context. Michael Mateas continued to work at CMU as the last remaining member of the project. Mateas, who was developing the elaborate

interactive drama *Facade* with Andrew Stern, completed his dissertation and a playable version of *Facade* in 2002 (Mateas 2002). The system was publicly released in 2005. With the last project member having finished at CMU, the original Oz Project concluded.

From the beginning, the Oz Project had been concerned with achieving new sorts of variation in narrative discourse, although the variations that were targeted were more elaborate and less straightforward than variations in order, speed, and frequency. Instead, variations that involved a complete sense of subjectivity were considered. An early technical report contemplated how IF might situate the reader “inside [Chief] Bromden’s mind” in a computer version of *One Flew Over the Cuckoo’s Nest* or allow the player to see from the perspective of Stephen Daedalus in *A Portrait of the Artist as a Young Man*. The report considered these along with an array of cinematic techniques for presenting narratives, arriving at some formulations that are similar to the ones in this project. For instance, the report proposed a component that “slows down the flow of time and magnifies details,” or employs what Chatman calls *stretch* to tell events in a longer-than-usual narrative time (Smith and Bates 1989). *Stretch* is considered “feasible as a deliberate experiment” in *Narrative Discourse* (Genette 1980:95). While Genette ends up omitting *stretch*, it is included in the capabilities of nn. The Oz report contains useful ideas, but, as is described further in section 4.1, its overall approach is not as systematic as Genette’s treatment of narrative discourse.

Mark Kantrowitz’s natural language generation system GLINDA was developed as a part of the Oz architecture (Kantrowitz 1990). The system, a realizer that was inspired in part by PAULINE (Hovy 1988), took a semantic representation of a sentence called a group and transformed this group using selection, organization, and combination rules. While it was developed to help “produce a variety of views of the reality that Oz simulates and to produce vivid dialogue” by “providing an interface of ‘knobs’ for controlling the style of the natural language output,” (Kantrowitz 1990:1) it worked only on the level of phrases and below this level, and so did not support larger-scale sorts of narrative variations.

Mateas’s work on subjective avatars (Mateas 1997) put into practice some of the ideas from the previously-discussed Oz report. His *Fastfood World*, one of several of the Oz Project’s IF-like “text worlds,” allowed the user to play a food service worker and to experience his workplace from this character’s perspective. Much of the subjective effect was accomplished through the

description of sense data which was collected into groups by the avatar and then sent through a natural-language generation system. The emotional state of the character would determine what would be narrated and how, so that the fearful background noise of the fryer would be mentioned initially and thereafter would only be mentioned if the character was particularly fearful. Some work was done toward generating a “stream of thought,” a free indirect discourse that would comment on the events transpiring (Mateas 1997:6).

In *Fastfood World*, Mateas went directly for the harder goal of accomplishing a sense of subjectivity — as seen in an internal, character-bound narrator — and sought to do this based on an AI model of emotional states, called Em, and a framework for representing the world and characters’ sensory perception of it, called Tok. Looking at cognition and these sorts of deep, psychological representations constituted a different approach than formalizing Genette’s concepts of mood and voice, which are based on functions of the text. While it seems sensible from a literary perspective to tackle some of the fundamental and easier aspects of narrative variation first, it was also reasonable for the Oz Project to look elsewhere. The project sought to create “highly interactive” experiences in which the opportunity to interact was continuous, not divided into turns. Since this meant that interaction with the system should always be “live,” and since it was based on the framework of drama rather than fiction, it suggested a direct mapping between time in the simulated world and time in the user’s experience. With such a mapping, the simpler sorts of narrative variation that change the relationship of simulated time to real time would become at least unwieldy and, in some cases, perhaps not sensible at all.

Nicolas Szilas has described and is developing an architecture for interactive drama that features a “virtual narrator” between the story world and the “theatre” in which events and existents are represented to the user. “The Virtual Narrator acts as a filter between the Narrative Logic and the Theatre. It chooses, among all possible actions, which ones are worth being displayed to the user” (Szilas 2001:73). Szilas’s Virtual Narrator does not vary the narrative discourse or expression plane, however. It is not actually restricted to determining what is expressed, as the previous description suggests, but determines what actions among many *possible* actions will happen in the underlying story world: “The Virtual Narrator is thus able to manage the course of actions in the interactive drama, by focusing on the effect to the user rather than the behavior of the characters” (Szilas 2001:74). Even if the Virtual Narrator did determine

those actions whose expression was to be suppressed, this would only allow it to accomplish ellipsis and not any of the numerous other sorts of narrative variations.

A recently developed research IF system is FrOz (Koller et al. 2004). Like the Oz text worlds, which are alluded to in the name of the system, FrOz presents a typical interactive-fiction framework and uses natural language generation to provide output rather than simply presenting pre-prepared orthographic strings. FrOz uses this generation facility to appropriately generate pronouns, and it incorporates a sophisticated ability to resolve referring expressions. The system uses description logic to represent both the state of the world and the interactor's knowledge. It uses a theorem prover to reason about these representations. While innovative and a good starting point for improved understanding and generation of referring expressions, there was no attempt in this work to improve upon or replicate the overall quality of textual output in commercial interactive fiction of the 1980s (Koller et al. 2004:212) or to provide the additional abstraction that would allow for high-level narrative variation.

3.4 *Commercial Games*

The many commercial efforts in interactive fiction resulted in several technical improvements. Ones that were documented early on included code compression (S Adams 1980, discussing his and Alexis Adams's *Pirate's Adventure*) and virtualization (Blank and Galley 1980, discussing the Z-Machine). Parsing and the world-model in the commercial *Zork I*, the essential elements of which were largely carried over from the mainframe *Zork*, were also described in some detail (Lebling *et al.* 1979, Lebling 1980). Infocom touted its parser's ability and branded this component "the Interlogic parser." Other companies followed suit, Synapse with the BTZ (Better Than Zork) parser and Melbourne House with its parser, which was called English. But none of these companies had any sort of well-developed text generation system to deal with output from the world model as the parser deals with input, much less a system that was given a name by the marketing department and was made into a major selling point. The work that was done along literary lines, including Robert Pinsky's *Mindwheel* (Synapse, 1984) and Thomas Disch's *Amnesia* (Cognetics Corporation, 1986), focused mostly on the symbolic use of space and objects, rich simulation, and twists in the interactive situation rather than variations in the telling of events. In *Amnesia*, Disch did make extensive use of analepsis in the form of interactive flashback

sequences, but these were hard-coded as game segments; events were not rearranged in the narrating dynamically. Douglas Adams and Steve Meretzky offered the interactor the chance to focalize and command different PCs in Infocom's *The Hitchhiker's Guide to the Galaxy* (1984), although something like this had already been done at that company by Michael Berlyn, who provided robot PCs with different perceptual capabilities in *Suspended* (1983). While parsers were continually being built and refined for use in new interactive fiction, general-purpose narrators were not being developed in the same way; narrative variation of different sorts occurred, when it did at all, as part of a one-off experiment.

Massively multiplayer online roleplaying games (MMORPGs), exemplified by *World of Warcraft*, *Everquest*, and *Ultima Online*, have long incorporated a separate, client-side module for displaying the user's view of the world. The LucasArts system *Habitat*, which was developed for the Commodore 64 starting in 1985, provides one early example of this sort of architecture (Morningstar and Farmer 2003). The view was abstracted from the underlying world model because of the scarcity of bandwidth; this bottleneck remains important in today's systems. A similar client/server split has been used in virtual worlds that are not easily characterized as games, from *AlphaWorld* to *Second Life*. As with graphical interactive drama, the real-time, highly-interactive nature of MMORPGs and other virtual worlds preclude variations in order, speed, and frequency.

A recent attempt to create a MMORPG with a "literary orientation" incorporating "narrative techniques from the novel and other traditions ... such as voice, point of view, characterization, plot, and so on" is *TriadCity* (SmartMonsters, Inc. 2006). A note about the project on the site quotes from Brian McHale's *Postmodernist Fiction* and cites Janet Murray's *Hamlet on the Holodeck*. The work done so far emphasizes the content plane, where underlying characterization and plot reside. Differences between *TriadCity* and other MMORPGs are said to include the inclusion of alternatives to violence, the permanence of death, and the ability to create complex bots (or "automata"). While the more difficult matters of voice and point of view are mentioned on the site as interesting, simpler and more fundamental variations in order, speed, frequency, mood and voice are not.

Other real-time multi-player games, including first-person shooters such as *Unreal Tournament*, the *Quake* series, *Counter Strike*, and the older original *Doom*, also situate the

presentation functions on the client in a separate module and allow the user some control over how the world is displayed: resolution, frame rate, and gamma can be set, for instance. Even these games, therefore, show a rudimentary separation between presentation and the underlying simulated world. Again, however, the real-time nature of such games more or less prohibits narrative variations such as analepsis or repetition of the same event multiple times, at least during game play. (Some interest in narrative variation of this sort is seen in replay and rewind features of several video games.) While these systems allow for the visual display to be controlled and offloaded to a client computer, they offer little in the way of narrative capability.

3.5 Other Interactive Natural Language Generation and Dialog Systems

Many natural language generation systems are non-interactive, generating weather reports based on sensor data, for instance, or generating printable patient information based on medical records. Other systems allow some form of user interaction other than natural-language input to direct what will be generated. Examples of these include ILEX, HIPS, and M-PRIO, which were used to develop prototype museum guides (Androutsopoulos et al. 2002). Finally, there are dialog systems which accept some form of natural language input and provide generated natural language in return. Systems that take this approach range from the early SHRDLU (Winograd 1972) through TEXT (McKeown 1985), in which a two-stage pipelined architecture and schemata were developed, to more recent systems that answer questions or allow other interactions. The interest in spoken dialog systems has particularly increased; recent years have seen research advances and many fielded systems.

The systems just mentioned differ from almost all of the previous ones in that they have well-developed natural language generation capabilities. Some of these interactive systems also offer well-developed discourse models. They are sure to provide some improved ways of generating descriptions (static statements about existents) in interactive fiction in the context of the discourse. However, they seldom produce any active statements outside of a limited range: ones about the state of the system or application (such as “Just a moment, I am processing your order”) rather than about the represented knowledge, for instance, and a number of other statements provided by canned texts. The museum guide prototypes are examples of propositional and descriptive systems which have well-developed discourse models but which

rely on representations of entities, facts, and relations rather than a history of events to generate language. Interestingly, the very early system SHRDLU does model what has happened in the world and is capable of retelling what has happened and what previous states of the world were like. However, SHRDLU does not provide for the sorts of narrative variation that are explored in this project.

4 Steps toward a Potential Narratology⁶

4.1 *Theorizing Interactive Fiction*

Interactive fiction (IF), a category that is typically represented by the text adventure or text game, has literary, gaming, and other important aspects. Well-known text-based interactive fiction includes *Adventure* (1977), *Zork* (1977-78), *A Mind Forever Voyaging* (1985), *Knight Orc* (1987), and *Curses* (1993). In the first book-length discussion of interactive fiction (Montfort 2003a) the form is introduced in detail, its important historical precursors and cultural contexts are discussed, and a figurative way to think about its poetics and aesthetics is described. In this chapter, this focus is on particular ways that narratology can inform a rigorous theory of interactive fiction, a theory that remains sensitive to the many-faceted nature of this new media form.

Systematically relating interactive fiction to “game” and “story” requires more than the *ad hoc* application of terms and concepts from literary theory, narratology, and gaming. Although humanists and scientists can be prodded toward insight by offhand approaches, deeper insights and more substantial progress require a methodological framework, a way to evaluate results, and (if more than one person is to participate) some sort of common language and understanding about the nature of the topic under consideration. To build a theory of interactive fiction that is useful in deeply understanding how interactive fiction is experienced, and how better sorts of works can be created, a stronger approach than that of the theory-bag is necessary, one which distinguishes those elements of interactive fiction that result from it being

- a text-accepting, text-generating computer program;
- a potential narrative, that is, a system which produces narrative during interaction;
- a simulation of an environment or world; and
- a structure of rules within which an outcome is sought, also known as a game.

Interactive fiction was, for a long time, almost entirely neglected in academic discussion. In

⁶ This section consists of a slightly revised version of my article (Montfort 2003b). The only substantial changes involve the introduction of the concepts of unrecognized inputs and clarifications, some further development of the nature of puzzles as requiring “non-obvious” actions, and the addition of a section (4.11) offering a typology of IF outputs.

the IF community, discussion has touched on many important aspects of interactive fiction, but no theory of this sort has been developed. Marnie Parker's "Iffy Theory" is an attempt to categorize people's taste in interactive fiction (Parker 2000). It is not about aesthetics or poetics as it does not explain, for instance, how one "auditory" IF work might be better or worse than another or what the elements of such a work are. Graham Nelson's "The Craft of Adventure" consists of advice about how to write interactive fiction well, as its title suggests. It discusses many related topics in depth but does not present the beginnings of a systematic theory. Nelson calls it "a string of grits of wisdom and half-baked critical opinions," (Nelson 1995) which greatly understates the importance of this article, but the collection, insightful as it is, does not offer a framework for a new theory.

One interesting attempt to offer such a framework is "Towards a Theory of Narrative in Interactive Fiction" by Sean Smith and Joseph Bates, a result of research at Carnegie Mellon's Oz Project which was discussed in the previous chapter. This report was an attempt to formulate interactive fiction in terms of cinema, based on "an art-film text taken at random from the shelves at CMU's library" (Smith and Bates 1989:6). No distinction was made between techniques specifically tied to time-based and visual effects and those generic to narration in any medium (Chatman 1975:299-300). While the paper does describe a series of techniques for interactive fiction that is inspired by cinema, the mappings between film and IF techniques are arbitrary and unsystematic. The ideas may be of interest to IF creators, but this essay does not actually propose one coherent theory.

Roger Carbol's "Locational Puzzle Theory" is interesting in that it attempts a strict definition of certain elements of interactive fiction (Carbol 2001). Unfortunately there are numerous difficulties with the approach. To begin with, Carbol defines a game only as "a collection of objects, in the object-oriented programming sense," which does not distinguish games from non-games, as any definition should. Furthermore, "object" is not defined by Carbol as it is in any thorough discussion of object-oriented programming, but as simply "a collection of properties." The impulse to define puzzles precisely and examine their nature is a good one, but the confusion in this approach — on the one hand between a software development methodology, objects in the IF world, and narration, for instance, and on the other hand between location in the space of the IF world, the awareness of the interactor, and the properties of programmatic objects — is far too

profound, and the resulting distinctions between classes of puzzles are no better than have already been devised in less principled attempts at classification (Rees 1993).

Emily Short's essay "What's IF?" makes several points of interest, although it does not define interactive fiction well enough to distinguish it from chatterbots and other programs (Short 2001). The concept of the *benchmark* as an unique action that makes progress toward an ending is a useful one, although the discussion in "What's IF?" is still somewhat preliminary, with *action* not defined, for instance, and with the supposedly formal benchmark being defined with appeal to the interactor's anticipation and other possibly interpretive factors. The discussion of puzzle has interesting aspects but does not conclude with a definition of puzzle that can be applied consistently by other theorists. This essay is a good effort to not only define qualities of a puzzle but also place puzzles in the overall context of an IF work. Difficulties with this essay's approach make it clear, however, that a theory that carefully distinguishes formal aspects from those related to interpretation will be very helpful.

Since a work of IF can be implemented in different ways and function identically, there is another space besides that of the interactor's interpretation which a formal theory should not enter. Definitions of the elements of an IF work from a theoretical perspective should be done without making reference to a program's specific data structures, functions, objects, and so forth, considering the program instead as a black box that accepts input and generates output. (The clearest justification for this is seen in cases where two programs that are identical from the standpoint of the interactor are implemented in radically different ways — for instance, first using a functional programming language and then using a procedural one. Different objects can of course also be used in two different object-oriented implementations.) It may happen that sensible programmers developing IF works have found it convenient to encapsulate certain fundamental elements as discrete entities in code. Those studying IF while using a theory of this sort should not need to refer to the internals of a program, however.⁷

Taking this view of a formal theory of IF, the chapter considers the nature of interactive

⁷ This is not an objection to reverse-engineering programs, looking at their source code, or otherwise considering the code level and the implementation of new media systems. Such analysis is essential for full understanding of digital media and can reveal aspects of practice and computing that would be difficult or impossible to see otherwise. The point here is simply that it is possible to consider how a work of interactive fiction functions separately from how it is implemented, and that it is appropriate to do so when conducting an analysis at the level of form and function.

fiction as *program*, *potential narrative*, *world*, and *game*, describes how the perspective of the person interacting can be represented, and starts in on the difficult issue of how to conceptualize the puzzle.

4.2 *Interactive Fiction and the Interactor*

A work of interactive fiction is, among other things, a computer program that accepts text input from a user and produces text output in reply. This user of an IF work is the *interactor*, following the terminology of the major academic effort in interactive fiction so far, the Oz Project; the term has been adopted by others (Murray 1995:161). It is synonymous with *player* as that term is usually used in the IF community, but *player* has other meanings related to games and drama while *interactor* has a history of being used only to refer to the person who interacts with an IF work or similar program. In the case of a work of IF that has no multimedia elements at all and uses only text for a medium, *text* simply refers to a string of words in the ordinary sense. However, *text* can also be considered semiotically to be any set of signifiers; thus IF works (and perhaps other works as well) that contain graphics, sound, or video can be accommodated in this way. Using *text* more specifically, to mean “strings of words,” *interactive fiction* indicates a category of text-based works, works that can contain other media elements but where text and textual exchange are central. *Computer program* could also be generalized to include other sorts of text machines in the broader *Cybertext* sense — written-out instructions that a person could follow, for instance, or Scott Adams mimicking his *Adventureland* by uttering the output it would give in reply to someone’s spoken input (Hoy and Jerz 2001). For the purposes of this essay, only computer programs in the usual sense need to be considered as interactive fiction, although, again, the theory presented here should be extensible to other types of systems.

Rather than state, as Short does, that “IF *tends* to represent, in some form, an environment or imagined world whose physical space we can explore,” (Short 2001) it seems better to say that a simulated world, the IF world, is essential to interactive fiction. The only counterexample Short advanced was Andrew Plotkin’s 1997 *The Space Under The Window*. This is a work of hypertext implemented in Inform; instead of clicking on a word as would be typical on the Web, typing one of the words displayed causes the appearance of a new lexia, in George Landow’s sense, indicating a section of hypertext. Plotkin refers to this work as “Not standard interactive fiction”

(Plotkin 2001). None of the theoretical discussion that Short develops in her essay applies to this work, which clearly seems better considered as hypertext than as interactive fiction. Requiring a simulated world does not mean that any particular code is required in a work of IF. Whether a work simulates a world or not can be determined from outside, by an interactor studying the work.

Since a simulated world and textual description of events in it is entailed by a program's being interactive fiction, an IF work is also necessarily a generator of narratives. The distinction between what can be simulated and what can be narrated is particularly important to understanding the workings of interactive fiction; although the *potential narrative* aspect of interactive fiction is produced based on events in the *world*, there may be things that are narrated during an interaction but are not simulated.

It is standard to refer to IF works as "games," but a work of IF is not necessarily a game (Giner-Sorolla 1996). A work can present a world which is pleasant to explore, but which has no quest or intrigue. There may be no final reply that is a "winning" one, perhaps no final reply at all. Because of this it makes more sense in theoretical discussion to refer to a *work* of IF, rather than using *game* as the generic term for everything in the form. Even in the case of works that are actually games, using the former term can help to signal that it is interactive fiction from all relevant perspectives, rather than interactive fiction *only* as game, that we are principally interested in. The advantage of using a term like "work" is most clear in the case of certain IF works that do have no optimal outcome (that is, they cannot be won), do not keep score, and contain no puzzles. Ian Finley's simulated gallery opening *Exhibition* provides a simulated space in which the player character can look at paintings while chatting with four characters who have very different perspectives on the artist and his work; there is no way to win or lose it. Calling this a "game" is unfair to *Exhibition*, which is not actually a game. Calling *Exhibition* a game is also unfair to IF works such as Dave Anderson's *Hollywood Hijinks*, which simulates a treasure hunt in a mansion and has a very definite and explicit goal. Works of this sort clearly are games. But in a careful discussion, the generic use of the term "game" to refer to every interactive fiction work would denude that term of its particular meaning. The term "game" is the norm in casual discussion, and is admittedly used by this author in those contexts. Another theorist and author refers to her own (clearly non-game) work by making reference to "a game like Galatea" (Short

2001). “Work” has real advantages as a term, however, in discussions where precision is essential.

4.3 *Sessions, Interactions, Traversals*

As computer literature pioneer Rob Wittig describes, while it is commonly thought that the reading of a book proceeds as “the reader dutifully trudges the linear track prescribed by the author,” this is certainly not always the case. A reading of a book may involve browsing it in the bookstore, reading in short bursts in different places, skipping ahead to see if it gets any better at the end, looking through bits in the middle to then figure out what happened, and giving up without actually reading everything (Wittig 1994:81-83). It is difficult today to understand much about the heavily-studied processes of reading without appreciating that “readings” may not be done in the intended sequence and may not be total. The nature of interaction and interactivity in interactive fiction, which has been studied hardly at all and which in general allows for no “total reading” of the book sort to be done, will be even harder to theorize without making distinctions between aspects of interactive fiction as computer program; ways in which IF works are world, game, and potential narrative; and the interactor’s own interpretation and experience.

A *session* is what happens during the execution of an IF program. The session begins when an IF program starts running. It ends when the program terminates. The text that results (both text typed by the interactor and text produced by the program) is the *session text*.

An *interaction* describes a series of continuous exchanges of texts between the program and the interactor. “Continuous” does not have a formal meaning, nor is it a property of the text or program. The interactor’s sense of continuity and unity is what makes a certain experience a single interaction; different interactors may have different opinions of what an interaction is. The text (from both interactor and program) that corresponds to an interaction is an *interaction text*.

The experience of interaction belongs to the person involved.⁸ The session, on the other hand, is a property of the program and its execution. Still, interactions and sessions often correspond: an interactor starts the IF program, reads and types for a while, perhaps saves (allowing the current state to be restored later on) or perhaps arrives at a conclusion, and then

⁸ Or to the people involved. It is common for several people to interact with one IF work at the same time, although this reality is seldom mentioned in discussions of interactive fiction. This document, however, does not deal with how multiple interactors can experience a single IF session together. There is also be no discussion of the interesting fact that an interactor could actually be a computer program rather than a person.

terminates the program. However, one interaction may take place over many sessions, because the interactor may terminate a program and then start it again immediately, interacting with the program repeatedly in what is to her a continuous interaction. Similarly, an interactor can start a session (and an interaction), go on vacation for a week while leaving the computer and the program running, and then return to have another, different interaction that is part of that same session. Of course, the point of many works of IF is to win them, that is, to proceed towards a certain goal or outcome; “winning” can be seen as one analogue to having “read the whole book.” (This is not the only such analogue, though.) Winning cannot be described in terms of *session* or *interaction* alone.

A *traversal* is what happens in one or more sessions, and one or more interactions, when the interactor “completes” a work of IF by going from the beginning until no more can be narrated. The full definition of traversal is given in section 4.5; to define the term exactly is it is necessary to describe more about IF as simulated world and potential narrative. The traversal is mentioned here because of its relationship to *session* and *interaction*. Of course the text corresponding to a traversal is called a *traversal text*.

4.4 *Cycles, Exchanges, and the IF World*

Anything contributed by the interactor, from a press of the spacebar to a long typed text, is an *input*. The texts produced by the program are *output*. If the program outputs some text that the interactor originally typed, that is nevertheless output, just as whatever the interactor types (even something previously output by the program) is input. A *cycle* is one input and all the output that follows it until the next input. The *initial output* is whatever output is produced before the first opportunity for input; this is before the first cycle. All of this is defined formally with regard to an IF work’s nature as a computer program. Pressing the space bar in response to “[MORE]” is an input, for instance, even though it normally provides the interactor no opportunity to influence the course of the narrative that is being produced. It is simply because an interactive fiction work is an interactive computer program that it has input and output.

In the sense that scholars of the story and of narrative (that is, narratologists) use the terms, a work of IF is not a narrative. An IF work is an interactive computer program, but not directly a *narrative*, “the representation of real or fictive events and situations in a time sequence” (Prince

1980:180). Similarly, interactive fiction is not a *story* in the sense of the things that happen in a narrative, or more precisely, “the content plane of narrative as opposed to its expression or discourse; the ‘what’ of a narrative as opposed to its ‘how’” (Prince 1987:91). In everyday speech, of course, “story” also refers to a particular genre, the type of thing people expect to hear when they say in conversation “so, tell me the story” or that a child expects to hear after asking to be read a story. Interactive fiction is not precisely this sort of story, either, although there may be a “frame story” provided in the documentation or there may be a certain type of story which is always generated in successfully traversing the work. An IF work is always related to story and narrative in their narratological sense, even if a particular work does not have a “story” in this ordinary sense.

The distinction narratology makes between story and narrative has been noted in various ways since Aristotle, who distinguished the argument, *logos*, and how it was arranged into plot, or *mythos*; the Russian formalists also distinguished the material of the story or *fabula* from how it was told in the *sjuzet* (Chatman 1975:295). Interactive fiction has the potential to produce narratives, usually as a result of the interactor typing things to effect action in the IF world. In fact IF works are *potential literature* in the sense of the Ouvroir de Littérature Potentielle (Workshop for Potential Literature, abbreviated Oulipo) (Mathews and Brotchie 1998, Motte 1986), and specifically they are potential narratives.

IF works also present simulated worlds. These *IF worlds* are not merely the setting of the literature that is realized; they also, among other things, serve to constrain and define the operation of the narrative-generating program. IF worlds are reflected in, but not equivalent to, maps, object trees, and descriptive texts. In fact, the IF world is the content plane of interactive fiction, just as story is the content plane of a narrative. The interactor typically types what one or more *player characters*, who exist within the IF world, are to do. The nature of the player character, and other sorts of characters, is discussed in greater detail in section 4.6.

An input that refers to an action in the IF world is a *command*. In narratological terms, a command is *diegetic* (Genette 1980:227-234, Cadre 2002). This command is usually in the form of an imperative to the player character. It does not have to refer to a physical action. Commands include *think*, any input directing the player character to speak, and any input directing the player character to examine something or otherwise sense something about the IF world.

Commands that do not succeed are still considered commands, as long as they are understood by the parser and interpreted as attempts at action. The input given to clarify a command (such as *kill the troll* What do you want to kill the troll with? *the sword*) is considered part of the command being clarified. An input that refers to several actions (for instance, *take all*) consists of the several commands into which it is decomposed by the parser.

Other inputs that refer to the program rather than the simulated world, such as those that save, restore, quit, restart, change the level of detail in the room descriptions, or address some entity that is not part of the IF world — to ask for hints, for instance — are *directives*. A directive is, in narratological terms, *extradiegetic* (Genette 1980:227-231). Commands and directives are two distinct sets; all inputs that are recognized by the program are one or the other. Directives include what Graham Nelson refers to as “meta” actions in Inform (Nelson 2001:90). Based on this, “meta-command” has been previously suggested to refer to actions outside the game world (Olsson 1997), but this term has the potential to confuse a narratological study of IF, since “meta” has already been used by Genette in the opposite direction — to refer to narratives within narratives rather than to refer to the level of narration itself. To avoid confusion the term “meta-command” is left, in this discussion, to refer only to its specific meaning within Inform programming, and “directive” is used for all inputs that do not refer to the IF world.

There are some inputs that are neither commands nor directives. Any input that is unrecognized, such as a typo or a statement too elaborate to parse, is in this category. It seemed expedient at one point to classify these *unrecognized inputs* as directives (Montfort 2003b), but work on an IF development system has shown that the modules for handling these two types of inputs should be different, since the function of these two types of inputs in the interaction is quite different.

Considering all inputs rather than just text entered at the prompt, it is still easy to classify recognized inputs into directives and commands. Pressing the spacebar when “[MORE]” is displayed to indicate that additional text is available is a directive, for instance, while typing a number to select one of several conversation options is a command. *what is a grue?* in *Zork* (1979, Tim Anderson, Marc Blank, Bruce Daniels, and Dave Lebling) appears to be a directive, since there is no one within the IF world to whom this question is addressed; the information is apparently related to the interactor outside the IF world. On the other hand *plugh* in *Zork* is a

command, because it refers to the player character speaking the word “plugh,” and it results in a hollow voice within the IF world saying “Cretin” in reply.

Outputs that follow input from the interactor and describe anything about the IF world and events in it (including the inability of the player character to enact a particular action as commanded) are *replies*. Whether the text is a direct result of what the interactor typed or whether the event it describes occurred because of a timed or random event, it is considered a reply, as long as it describes something about the IF world. All other outputs — that is, all outputs that do not describe the IF world — are *reports*. “[MORE]” and “[Press space to continue]” as they usually appear are reports, as are “Are you sure you want to quit?” “Your score is 0 out of a possible 100, in 2 moves.” and “Brief descriptions.”

	Extradiegetic	Diegetic
	Interactor	Player Character
Input	Directive	Command
<i>e.g.</i>	<i>QUIT</i>	<i>PICK UP THE PHONE BOOTH</i>
Output	Report	Reply
<i>e.g.</i>	<i>Are you sure you want to quit?</i>	<i>You find nothing of interest there.</i>

Table 1. Recognized inputs and the outputs that correspond to them may be diegetic or not.

An *exchange* is one command and the reply that follows it; the reply in this case includes all references to the IF world in all the output, up until the next command is entered. As command and reply correspond to input and output, so exchange corresponds to cycle.

The following excerpt from a session text of *Zork* presents two exchanges, in bold:

```

>open the mailbox
Opening the small mailbox reveals:
A leaflet.

>ear the leaflet
I don't understand "ear".

>eat the leaflet
Taken.
I don't think that the leaflet would agree with you.

```

In the first exchange, the player character is ordered to open a mailbox. This is accomplished and the result is narrated: a leaflet is now visible. Next there is an input that is not a command, since it is not understood to refer to the IF world. This is an unrecognized input that produces a

clarification, “I don’t know the word ‘ear’” — revealing the limited vocabulary and brittle nature of interaction in early interactive fiction, problems which have only been mitigated in part. That cycle does not constitute an exchange. Finally there is a command for the player character to eat the leaflet. This results in the player character taking possession of it but not actually eating it. The reply seems bizarre in context; an understanding of the distinction between the diegetic and the extradiegetic, and between the command and directive, helps to explain why. “I don’t think that the leaflet would agree with you,” coming at this point in this session text, makes it seem as if the extradiegetic “I” in the previous report (the “I” who cannot understand certain words and translate them into actions) is now somehow within the IF world, counseling the player character not to eat a piece of direct mail. Further implications of this sort of transgression, and other sorts, are discussed in section 4.8.

4.5 *Initial Situation to Final Situation, Prologue to Final Reply*

The IF world can be described before the first opportunity for a command. It usually is. Such a description is the *prologue*. The term is used here much as it was in the PrologueComp, a 2001 writing contest announced on rec.*.int-fiction (Myers 2001), except that, strictly speaking, any of this initial text that does not describe the IF world is not considered part of the prologue. This concept is similar to that of the *overture* (Nelson 2001:270).

The state of the IF world after the prologue, when the first opportunity to enter a command is presented, is the *initial situation*. A single IF work may have multiple initial situations, but because of how the initial situation is defined these cannot possibly be determined by the interactor’s input. This is because the first input that can influence the world in any way is the first command; the opportunity to enter this command comes after this initial situation. Different initial situations might be determined by randomness (Short 2001), by the presence or absence of a particular file on the computer’s hard disk, by the date and time, or by any other factor besides interactor input. The initial situation refers to the state of the IF world, not how that state is described. A work of IF may begin immediately with a prompt, describing nothing about the IF world. Jon Ingold’s 2001 *All Roads* begins with a quotation and a menu but does not state anything about the IF world or the player character’s situation. Thus, it has a *null prologue*. Similarly, the 1998 *Bad Machine* by Dan Shiovitz begins with just a prompt and has a null

prologue. Nevertheless, like all IF works, these have an initial situation — this situation is simply not described before the first prompt for input. As commands are provided by the interactor, the replies reveal what this initial situation was.

The *final reply* is that reply after which the narration of events in the IF world cannot be continued. This text indicates what is usually called an ending (Short 2001). After the final reply either the program terminates or the only option is to input a directive. The state of the IF world that is described in the final reply cannot be changed by any commands made after the final reply. In traditional interactive fiction, the final reply usually narrates either the player character's death or ultimate triumph. A final reply is not required for a work to be interactive fiction, and some works, by design, do not produce a final reply. An unfinished or bug-ridden work might also not produce a final reply at all; it might instead only manage to produce a final report that is an extradiegetic error message, explaining what caused the program to crash.

By convention, some directives, such as *quit*, *restore*, and *restart*, are allowed after the final reply. Neither *restore* nor *restart* allow the narrations of the IF world to continue, however, after a true final reply; they revert the IF world to some other saved state or to an initial situation. Similarly, *undo* in this situation does not allow a narration to continue; it simply restores the previous state of the IF world and allows the narrative to continue from that point. Adam Cadre's 2000 *Shrapnel* achieved its effect by presenting what seemed to be final narrations while actually continuing to narrate events in the same IF world in reply to subsequent commands, suggesting a transgression. The transgression is between what will be called different *courses*.

A series of exchanges that are part of the same narration, and are presented along with all the directives and reports embedded in it, constitutes a *course*. The earlier excerpt from *Zork* describes a course, for instance. In Andrew Pontius's *Rematch* and Sam Barlow's *Aisle* there can be no courses longer than one exchange. The following session text, from Emily Short's 2000 *Metamorphoses*, illustrates how — because of certain directives — a single session text can contain several courses. It also shows how an exchange can be part of more than one course. Exchanges, which have been numbered, are in bold:

1 >**get the rock**
Taken.

2 >**put the rock in the water**
Anything you dropped in there, you would be unable to retrieve.

>undo
Shore of An Underground Lake
[Previous turn undone.]

3 >**hit the bell**
You slap ineffectually at the bell.

4 >**hit the bell with the rock**
The peal is deep and resonant; the surface of the lake stands up in ripples; the darkness grows (if that is possible) more dark. Even when the sound has died and the water stilled, you find yourself waiting.

>undo
Shore of An Underground Lake
[Previous turn undone.]

5 >**listen to the bell**
You hear nothing unexpected.

1-2 is a course; nothing occurs after exchange 2 because that command is undone. 1-3-4 is another course. To quote this course we simply include everything up through exchange 4; 2 is now considered as directive because the UNDO directive was input after it, rendering that input hypothetical and meaningless within the IF world. What was the reply to 2 can be considered a report: because of the effect of UNDO this text now tells the interactor, outside the IF world, what would have happened had the command “put the rock in the water” been issued at that point in time — or, in a work of IF that does not depend on time or chance, what will happen if that command is then entered. Similarly, 1-3-5 is a course. Since any portion of a course containing at least one exchange is also a course, 1-2, 1-3-4, and 1-3-5 are only the longest three courses of fifteen in this session text.

Typing *restore* and restoring an earlier situation brings one to the end of an earlier course, where the *save* directive was issued. This allows a single course to extend across several sessions. A course can also extend across several interactions.

Can the same situation recur within a course? This depends on the nature of the IF world. In a world in which time always progresses, one cannot return to the same situation within a course — it will be later, so at least one aspect of the situation will have changed. But if time does not exist or if its laws are different, it may be possible. In fact, it is only impossible for a situation to occur twice in a course if an irreversible event occurs after every command. The progression of time is a special case of this. Note that keeping a count of how many “moves” have been made may or may not pertain to the IF world. If events always occur in the IF world after a certain

number of moves have been made, this is relevant to that IF world, but the number of moves made may just be provided (in a report) for the interactor's information. The player, of course, may not be stepping in the same stream twice when a situation recurs, since she may have a different level of knowledge the second time. But "situation" refers only to the state of the IF world, not to that of the interactor.

The state of the IF world after a true final reply is a *final situation*. So a *traversal* of an IF work is the course extending from a prologue to a final reply, and from an initial situation to a final situation. A *successful traversal* ends with a final situation that corresponds to winning; this seems consistent with what is meant by *playthrough*, (Short 2001) a term that has only been used on rec.arts.int-fiction fairly recently (Schmidt 1999) despite its longer history of use pertaining to video games. Since that term has been used in video gaming to refer to something more like a *traversal* in general, or to refer to the completion of a level, the terminology presented here seems preferable.

4.6 *Player Characters, Non-Player Characters, and Other Persons*

A *character* in interactive fiction is a person who is simulated within the IF world. A character's actions as narrated can differ depending upon the input provided. The term as it pertains to interactive fiction derives not only from dramatic use and from discussion of the novel, but also from the specific use of the terms *player character* and *non-player character* in the prototypical fantasy roleplaying game, *Dungeons and Dragons*. These terms have a similar special meaning in interactive fiction.

A *player character* or PC is a character directly commanded by the interactor. Any other character is a *non-player character* or NPC. The interactor may request that an NPC do something, or even command an NPC to do something, but such a request or command will always be done via the PC, who is the one directly commanded. NPCs are the anthropomorphic entities who can take actions in some way within the IF world — similar to the PC-like entities called *actors* (Lebling *et al.* 1979) — but who are not directly commanded by the interactor. An actor does not have to be anthropomorphic, but this is a requisite for an NPC. An adventurer-like freedom of action or ability to act is not required in either case.

There are also *other persons* who are mentioned but who are neither PCs nor NPCs. (Since the

terms *player character* and *non-player character* seem to complete the set of characters, these other persons are better not called characters; besides, in the study of narratives the term “characters” only refers to those people who actually exist within the story, not those who are simply mentioned.) Marshall Robner, the man whose death sets up the initial situation in Mark Blank’s 1982 *Deadline*, is not a character in that work of IF. Lord Dimwit Flathead is not a character in *Zork I*, either, since he is mentioned but not simulated. In Brian Moriarty’s 1985 *Wishbringer*, the dragon Thermofax appears alive (albeit in a daydream) in the prologue, but it is not possible at any other point during an interaction for Thermofax to be mentioned again in a reply, and thus no input causes his actions to vary and he is not simulated. Thermofax is a person, but not a character. Three scientists who appear at various points in an interaction as if they were in the room with the player character in Ian Finley’s 1997 *Babel* are also not characters, since they can be recalled by touching objects but are not simulated in the IF world; no actions can influence what happens (or rather, what happened) to them, and they cannot undertake any actions in the simulated IF world.

The idea of a character (including player characters and non-player characters) in interactive fiction is analogous to the idea of a character in a narrative, defined as “an EXISTENT endowed with anthropomorphic traits and engaged in anthropomorphic actions; an ACTOR with anthropomorphic attributes” (Prince 1987:12). The difference is that a character in interactive fiction must be an existent who acts within the IF world. Being a part of the simulation, rather than being a part of the story that the generated narrative tells, is essential for a character in interactive fiction. Since people may disagree about what traits are sufficiently anthropomorphic to allow an entity to be a character in a story, there are sure to be some similar disagreements about whether something is a character (or indeed, whether it is even in the broader anthropomorphic category “person”) in interactive fiction. But the category “character” in interactive fiction is similar to that category in narrative, and should be as useful. The presence of entities that cannot easily be seen as anthropomorphic or not, as in Dan Schmidt’s 1999 *For a Change*, has an interesting effect, in part, because it tends to defy the easy categorization that we would like to make when thinking about characters.

Aside from the issue of how anthropomorphic a person has to be, there may be dispute about what constitutes “simulation,” and therefore whether a person exists as part of the

simulated world and should be considered an NPC. Sean Barrett gives the case of the Implementors in *Enchanter*, who appear as a result of the player character casting a spell, then immediately disappear (Barrett 2002). They have a sort of existence within the IF world, but there is no opportunity to interact with them. Therefore, although they are narrated and their narration is the result of a command, they are not simulated in the way that the thief, the robot, or the troll is in *Zork*. An opportunity for the interactor's input to influence the behavior of a person — not simply to cue an appearance — would seem to be important in designating this person an NPC. Thus, the Implementors are other persons and not NPCs in *Enchanter*.

4.7 *World, Rooms*

As has been discussed already, a defining characteristic of interactive fiction is the simulation of a world. This is one aspect that distinguishes an IF work from, for instance, a chatterbot like ELIZA/DOCTOR (Weizenbaum 1966).

The IF world is divided into discrete locations known as *rooms*, which have also been called *locations* and *areas*. Like other essential elements of the form, rooms are defined independent of their implementation. A room is a simulated place from which a certain set of elements in the IF world can be sensed, manipulated, or otherwise acted upon. A room quite often contains *objects*; of course portable objects may be present or absent in different situations and objects that are present may be configured differently (for instance, may be open or closed). A different configuration of objects does not make for a different room. Rather, if a command is required to move the player character in space before certain other objects can be manipulated, those objects are said to be in a different room. Rooms, like characters, are simulated and are part of the IF world; they are not just mentioned in some of the narrations that are produced.

Shade, for instance, is aptly described as “a one-room game set in your apartment,” (Plotkin 2001) even though the player character *can* be commanded to move between the futon, the main room, the bathroom nook, and the kitchen nook. There is, by the definition presented here, only one room, because all the actions that are possible in one part of this apartment can be conducted from any other part of it, with the movement between parts of the apartment automatically entailed. The only exception is that the interactor must command the player character to stand up initially, but this is part of waking up rather than being a restriction on moving around in general.

After this, any action in any location is possible with a single command, even if the player character is back on the futon.

Even if there were works of IF that allowed the interactor to type a command like *move three centimeters left*, represented the position of the interaction is a seemingly-continuous way, and thus described an environment not broken into discrete rooms the way that traditional interactive fiction is, there would still be certain sets of actions that were possible at all the different potential locations of the player character. Thus, this definition of *room*, although possibly less useful in this circumstance, would still apply.

Rooms are *adjacent* if the player character can move between them as a result of a single command that represents a single action in the IF world. Opening a door usually changes the adjacency of rooms. By this definition, *End of Road* and *Inside Building* in Will Crowther and Don Wood's 1976 *Adventure* are adjacent, as are *Inside Building* and *Y2*, since a magic word will move the player character between these rooms immediately. However, even though the robots in Michael Berlyn's 1983 *Suspended* can be commanded to move to any room from any other room, all rooms are not adjacent to all other rooms because movement between rooms occurs as a series of discrete actions, each of which is simulated in sequence over time. A robot's movement may be interrupted along the way by some obstacle or by a new command that countermands the previous one; the whole trip is not atomic, as it is when moving from one room to an adjacent one.

After a player character has been to every room, the IF work has been *fully explored*.

4.8 *Diegesis, Hypodiegesis, and Extradiegesis*

Up to now "IF world" has been used as if there were a single world for each IF work. Actually, there may be many worlds in a given IF work, just as there may be several stories told in a single text, including *hypodiegetic* ones nested inside the main *diegetic* one. (The "frame story" of the *1001 Nights* is diegetic, for example, while the stories Scheherazade tells are hypodiegetic.⁹) IF worlds, like the stories in a text, may be linked in certain ways. In Steven Meretzky's 1985 *A Mind Forever Voyaging* there are six simulated future worlds in which Perry Simm is the player

⁹ Genette uses the term "metadiegesis" instead of "hypodiegesis," but admits that in his usage, "this term functions in a way opposite to that of its model in logic and linguistics" (Genette 1980:228). Other narratologists have used "hypodiegesis" to refer (less confusingly) to narration at this same level, so that term is adopted here.

character; these occur in a framework in which PRISM, a sentient computer, is the player character. The world with PRISM is *diegetic*, while the worlds with Perry Simm are *hypodiegetic*. Commands that refer to action in such a world can be called *hypodiegetic commands*. In *A Mind Forever Voyaging*, an hypodiegetic world can be reached by putting the player character into Simulation Mode, one of several modes that are available. As Perry Simm, the player character then walks around a simulated version of the city Rockvil. Typing *north* in this mode provides a hypodiegetic command (it is an instruction for the simulated human being Perry Simm to go north), while *record on* is a command of the usual sort (it is an instruction for the computer PRISM, in the frame world, to begin recording what Perry Simm is seeing).

Suspended presents an interesting case in which the player character is in partial suspended animation in a cylinder, and only a few commands (such as *wait*) refer directly to actions of the PC. Most commands are hypodiegetic commands issued to robots, who, although they are described by the generated narratives as being in the same physical space, an underground complex, are really in a different IF world. The robots, unlike the immobile human player character, can be told to go to different parts of the complex, can sense things, and can manipulate the environment to effect repairs. They exist and act in the IF world of this underground complex. The human “controller,” fixed in the canister in the middle of a large room in the complex and unable to take any physical action at all, is most clearly seen as being part of a different (but linked) IF world. Rather than seeing the robots (who are under the complete command of the interactor) as non-player characters, it makes sense to see them as player characters in a hypodiegetic world, similar to Perry Simm in one of the simulated futures of Rockvil. That the top-level world can be breached by a robot in the second-level world, who can be commanded to open the cylinder, ripping wires from and killing the player character in the frame world, can be seen as an instance of fatal *metalepsis* (Genette 1980:234-237), a transgression between different levels of story or between story and narration. This fatal variety, specifically as encountered in interactive fiction, has been called *dyslepsis* (Aarseth 1997 :118); of course a sort of dyslepsis can occur in narrative also, as in Julio Cortázar’s short story “Continuidad de los parques” (“Continuity of Parks.”)

Reference to the nature of interactive fiction as a program is no novelty. When Don Woods first expanded *Adventure* to create the canonical work of interactive fiction, he added a segment

that would be encountered at the end of a successful traversal; in this segment, the “closed” cave was fairly explicitly presented as a computer program that was not running. This is an example of one other type of metalepsis. Another clear and memorable instance of metalepsis early on in the history of the form is in Steve Meretzky’s *Planetfall*: The robot character Floyd (within the IF world) comments amusingly on the use of the *save* directive, which is extradiegetic and which Floyd should not know about. In *Planetfall*, the awareness of metalepsis allowed humorous use of it; the unintentional metalepsis shown in the *Zork* session text in section 4.4 is, instead, awkward.

Understanding the basics of diegesis, hypodiegesis, and extradiegesis allows us to make more sense of the seeming polyphony of voices in which statements are made in the computer-generated text of interactive fiction. “There are at least three identities involved in play: the person typing and reading (‘player’), the main character within the story (‘protagonist’), and the voice speaking about what this character sees and feels (‘narrator’)” (Nelson 2001:368). Nelson states that this narrator speaks the prologue, but notes that “in some games it might be said that the parser, who asks questions like ‘Which do you mean...?’ and in some games speaks only in square brackets, is a fourth character, quite different from the narrator” (Nelson 2001:373). These different speakers in the computer-generated text are what have led others to identify the narrative voice not “as a singular speaker but, rather, as a composite, mechanical chorus coming from both inside and outside the intrigue envelope” (Aarseth 1997:120).

Just as a work of interactive fiction can have many worlds, it can have many different narrators — which need not all correspond neatly to each of the worlds. For instance, at different times, different narrators might report the events that transpire in a single world. The voice of the parser (and of other parts of the program, such as those responsible for the ability to *save* and *restore* a particular situation) is extranarrative, and need not correspond to any of these narrators. Similarly, a voice that reports on hypodiegetic events (those that happen in a world within the main IF world) is hyponarrative. The numerous voices evident in even a simple work of interactive fiction are not a undifferentiated confusion or chorus, but typically correspond to different functions in interactive fiction which can be separated. Even in those cases where different voices are confused (as with the example from *Zork* given earlier) the particular voices which are being confused, intentionally or unintentionally, can be identified.

4.9 *Winning and Losing*

Many IF works have a goal that is explicitly presented or that becomes clear during interaction. Such works often indicate during their final reply whether or not this goal has been achieved. By several definitions, works of this sort, as with any rule-based activity engaged in for an outcome or for symbolic rewards, are games (Aarseth 2001, Zimmerman 2000, Zimmerman 2001).

Reaching a final reply that indicates the achievement of the IF work's goal is *winning*, and a traversal that ends in such a reply is a *successful traversal*. Similarly, reaching a final reply that indicates failure is *losing*, which concludes an *unsuccessful traversal*.

It seems the first work of IF to problematize the concept of "winning" was Michael Berlyn and Patricia Fogleman's 1983 *Infidel*. The final reply in *Infidel*, after completing the final task and achieving the highest possible score, includes the text "You will never get out of this pyramid alive. You earned this treasure. But it cost you your life." Despite the attainment of the maximum score, the goal of *Infidel* was clearly not to perish inside its pyramid, having collected all the treasure. But the goal — to plunder the pyramid and escape — could not be achieved; it was possible to attain the top score and solve all of the puzzles, but only possible to win this sort of Pyhrric victory. Still, interactors could state that they "won" *Infidel* after getting to this final reply. Later works, including *Exhibition*, *Aisle*, and Emily Short's *Galatea*, offer no optimal final reply; it would be bizarre for an interactor to claim to have won one of these. The 2001 work *Schroedinger's Cat* by James Willson does not even produce a final reply, so it is impossible to traverse at all. However, it can be solved in a certain sense, since it presents a world that the interactor can theorize about, experiment with, and understand; this notion of solution is discussed in section 4.10.

In Michael Gentry's 1998 *Little Blue Men*, in contrast, it is possible to win after entering only a few of the most obvious commands. (*Little Blue Men* can be won in 10 commands; an interaction that results in a successful traversal might take only two minutes.) The optimal score is achieved in this outcome, and the final reply includes the text "**** You have learned to love yourself *** // In this game, you have finally managed to love life." *Little Blue Men* is a much more intricate and complex work than such a victory would suggest, however. A different choice of commands brings the player character into an office environment that holds many puzzles and conceals something bizarre and horrifying.

An IF work has been *won* after a successful traversal, when a winning final reply is produced and a winning final situation reached. Since *Adventure* and *Zork* there has been a tradition of “the last lousy point.” Because of this and for other reasons, many IF works can be won without achieving the full score. Winning, besides not necessarily corresponding to attaining the maximum score, also does not particularly correspond to *full exploration*. It also may not correspond to the *solution* of the work.

Although IF works are always called games, and almost all of them are games, their nature as game has hardly been explored at all. A common idea is that the author competes against the player in the “game” of interactive fiction, but this makes no sense when considered in the context of other games. The developers of Monopoly, from Elizabeth Magie to Charles Darrow, do not compete with the people playing Monopoly. Will Wright and his team do not compete with a person playing *The Sims*. Nor is the computer the opponent in interactive fiction, any more than a computer version of solitaire opposes the player. In interactive fiction, the computer serves as a referee rather than an opponent (Solomon 1984:20). (If the computer provides hints it may be acting in a different role, that of a second.) “An Adventure game is an example of what a games theorist would call a cooperative game. If there are many players, as is often the case, they function as a team” (Solomon 1984:21). The myth that interaction in these sorts of games is solitary and is done by a lone interactor contributes to this misunderstanding of the form.

From the standpoint of game theory, the typical interactive fiction game differs from a game like chess not only because the players in chess oppose one another but because in that game total information about the game state is always available to players. The state of the game (or the state of the IF world) is known only in part in interactive fiction, and, furthermore, the workings of this world (and of the particular interface to it) are also unknown. Thus “the discovery of the rules, through trial and error, is one of the principal attractions of the game. The mark of a well-designed game of this type is that the rules reveal a consistent style, and are not merely arbitrary” (Solomon 1984:20). The nature of IF as game is too complex a topic to explore further in the current discussion, but clearly it is necessary here as well to recognize what type of game interactive fiction works can be and what aspects of a game help to make it interesting. It is worth noting that the perspective of game theory does support the figure of the riddle as a way of understanding interactive fiction, although the riddle may not formally be the same type of

game. The text of a riddle itself is completely known to a riddlee, but solving a riddle requires that the workings of the riddle's world be explored and understood, that its rules be discovered.

4.10 *Puzzles and Their Solution*

One way of understanding the relationship between the literary and the puzzling aspects of interactive fiction is by reference to the riddle, a figure that — unlike “puzzle,” “problem,” “game,” “world,” and many other commonly-invoked figures, can actually help to explain how the literary and puzzling aspects of the form work together (Montfort 2003b, Montfort 2004). The riddle is seldom invoked directly as a figure — at best, it is discussed as one type of puzzle that might be presented. But the figure of the riddle is consistent with some discussion of the puzzle in the IF community. Gregory Cox suggested two requirements for a puzzle: “a puzzle has to have an objective” and “a puzzle can't be obvious” (Cox 1999). This is quite similar to a definition of the riddle that has been advanced: “Every proper riddle must fulfill two conditions: the first is its social function as a competition between the riddler and riddlees; the second is its literary form, which must be difficult and enigmatic, yet containing the clues necessary to decipher it” (Pagis 1996:81). A similar definition of puzzle is a good start, but it leaves several questions open. Is a puzzle posed to the player character or to the interactor? Does a puzzle have to be “required” for a successful traversal in order for it to be considered a puzzle?

This section uses the canonical Crowther and Woods *Adventure*, the first known 350-point version, to discuss puzzles in depth. If theorists can agree about how many puzzles *Adventure* has and what they are — or even if they can disagree and articulate exactly how they disagree and why — this will be a good sign that the concept of a puzzle can be sensibly discussed as it pertains to IF works in general.

It seems possible to fruitfully discuss puzzles as formal elements of an IF work. In fact, it makes little sense to seek the puzzle in the mind of the author. What if the author is persuaded that it is a puzzle for the interactor to figure out how to type *go north* when the player character is in a room where a doorway is clearly described as being to the north? The author's belief does not, by itself, make this a puzzle. Similarly, we should not simply believe an author who denies that a certain intricate and difficult-to-discover series of required actions constitutes a puzzle.

It will also not do to rely too much on a specific interactor's state of mind and level of

intelligence. Clearly, since puzzles are constructed to challenge people, a definition must refer to the thought process of the interactor in some way. Still, it makes little sense to consider that Graham Nelson's 1993 *Curses*, for instance, actually contains more puzzles when a novice sits down at the computer than it does when an expert begins to interact.¹⁰ Also, puzzles should remain puzzles even if a particular interactor knows how to solve them. However, a puzzle does need to be presented as a challenge to the *interactor*, not necessarily to the player character. It is the interactor's effort at figuring out a puzzle, not any labor on the part of the player character, that is important. This is seen most clearly in part of Jeff O'Neill's 1987 *Nord and Bert Couldn't Make Head or Tail of It* and in this author's 2000 *Ad Verbum*, where the solution of puzzles relies on expressing a command properly rather than actually determining the correct action that the player character should perform. Although the IF world is essential to puzzles, puzzles are ultimately posed to the interactor outside the level of the IF world.

There is no necessary relationship between the score and the solution of puzzles. This is seen easily in a work such as Andrew Plotkin's 1995 *A Change in the Weather*, which has puzzles but no score. In *Adventure*, score has little to do, directly, with solving puzzles; it is mostly tied to picking up and dropping treasures. Driving away the snake, which clearly seems a puzzle, does not, in itself, earn the interactor any points. Yet *Adventure* awards 25 points for "getting well into the cave," although nothing special needs to be done to get that far. Crowther's original version of *Adventure* did not keep score (Peterson 1983:188).

There is also no requirement that anything immediate happen in the world when a puzzle is solved: the player character may only later visit another part of the world to see the result of solving a puzzle. Solving puzzles does not always unlock new parts of the IF world, or unfold some larger space; a solution may restrict rather than enlarge a player character's, and therefore the interactor's, options. As a result of collecting all the treasures in *Adventure*, for instance, the cave closes and the player character is teleported to a new and much smaller location.

A puzzle is a challenge in a work of IF that requires a non-obvious set of commands in order

¹⁰ Since *Adventure* was the first work of IF, this case is unusual; people did not know anything about how to interact, and just discovering how to move around and get into the cave was challenging. Even in this case, figuring out how to operate the work of IF in general, and how to move the player character about, is best not considered as a puzzle itself, although it may be essential to the pleasure of (or disappointments with) interaction. A difficulty that can reasonably be considered a puzzle is seen when the general operation of an IF work differs from the standard operation of interactive fiction, and thus presents a special challenge, as in Carl Muckenhoupt's 2001 *The Gostak*.

to be met. Non-obvious refers to a hypothetical, typical interactor encountering the work for the first time; puzzles do remain puzzles, in this formulation, after an interactor discovers how to solve them. Unlocking the grate with a ring of keys, found in plain sight a few rooms away, is not a puzzle, since it is obvious that keys unlock things. A series of obvious actions (open a box, remove the key from inside, unlock the door) remain obvious, but an action is non-obvious when an interactor must move beyond routine and do something out of the ordinary to understand the world and how to proceed. Looking beyond the obvious might require close reading to uncover hidden senses of a character's speech or of descriptions of things, conducting experiments (for instance, by putting different objects inside a machine and activating it to figure out what the machine does), or attaining understanding of the nature of something described in figurative language. While unlocking a door with a key is obvious, recognizing that something unusual is a key goes beyond the obvious.

This is not enough of a definition to classify every challenge as obvious or not, but this criterion at least begins to suggest some way of identifying puzzles, one that does not refer to the author's intentions and the interactor's specific knowledge and aspirations. Any typical interactor with the appropriate language skills, typing and computer interaction skills, and basic sorts of common knowledge would, in this formulation, be able to determine what is or is not a puzzle simply by studying the IF work in question, without needing to interview the author or take a survey of other interactors. The other factors essential to the determination of "obviousness" should be not the mindset of the author or of a particular interactor, but the culture or subculture within which the work was published — along with the conventions of interactive fiction.

There is no requirement that a puzzle's challenge relate to any other elements of an IF work in order for it to be a puzzle. It simply has to be presented as a challenge. While the typical way of doing this is to make the solution to a puzzle a requisite for a successful traversal, puzzles can be presented in some other way. Formally, the *solution to a puzzle* is the series of commands that meet the challenge of a puzzle. A *solution to a work of IF* is a series of commands that result in a *successful traversal*, with puzzles solved along the way. The typical walkthrough, of the sort often found online, records a solution to a work of IF. It is important to note that "solution" has not only a formal meaning but also a meaning that refers to an interactor's interpretation, operation,

and understanding of an interactive fiction work. An interactor who, by interacting with a work, comprehends the entire system of the IF world — why it is arranged as it is and why it functions as it does — has solved the interactive fiction work in this sense.

The puzzles in *Adventure* are:

- Driving the snake away
- Getting the gold nugget out
- Getting the emerald out
- Lighting the dark room
- Killing the dragon
- Creating a bridge
- Dropping the vase safely
- Watering the beanstalk twice
- Oiling the door
- Opening the oyster
- Replacing the troll's treasure
- Feeding the bear to calm it
- Deploying the bear against the troll
- Finding the way through the Pirate's maze
- Finding and purchasing lamp batteries in the other maze
- Blasting out of the repository
- Dropping the magazine at Witt's End and leaving the area

The last of these presents what is probably the most questionable case. Puzzles do not have to be required for a successful traversal of a work in order to be puzzles, according to the definition advanced here; they do not have to be tied to any benchmark or other plot element. In the case of dropping the magazine at Witt's End and leaving the area, this is a puzzle because *Adventure* clearly presents a challenge to the interactor: to get the last lousy point, independent of successfully traversing and winning *Adventure*. If the interactor had 350 points beforehand and dropping the magazine gave the interactor 351 points — and there was thus no way to know beforehand that an extra point could be obtained — this could be referred to as an *Easter egg* but would not be a puzzle. A challenge would not have been presented initially.

This last lousy point also demonstrates that solving a puzzle does not have to relate to anything meaningful in the IF world. Since this puzzle does not, it may make for a lousy puzzle, but the meaningless commands to drop the magazine and leave Witt's End are nevertheless the solution to an actual puzzle. The typical method of solving this — which involves reverse-engineering the program and actually reading through the resulting assembly language to figure out where the last point is assigned in the code — is certainly challenging for the interactor, despite the lack of any relationship to the IF world.

Finding the batteries in the other maze is a puzzle since it is presented as a challenge, issued rather directly when the lamp runs low. It is not required for a successful traversal, however. In fact, buying a lamp battery deprives the player character of one treasure and the possibility of gaining the full score.

Collecting the remaining treasures and depositing them in the building is not a puzzle or set of puzzles, because, although the pirate might steal a treasure to thwart its being moved, in general these actions are no more difficult than picking up other objects and moving them around — they just happen to be scored. A series of actions that is required for a *successful traversal* but is not a puzzle can be considered a *task*. In *A Mind Forever Voyaging*, there are no puzzles presented in the initial Rockvil simulation in Part I, only a list of tasks. The interactor is challenged to carry out these tasks, and it is enjoyable to explore and experience the IF world while doing them, but they do not require the interactor to do anything non-obvious.

4.11 *Typology of Interactive Fiction Outputs*

When considering how to develop natural language generation for interactive fiction, it is important to figure out what text needs to be generated and for what purposes. What are the functions of the texts that constitute interactive fiction output? Without determining what texts do in interactive fiction, it will be very difficult to construct a system that can generate text to cover all of these functions.

First, broadly speaking, there are texts that pertain in some way to the content plane and texts that pertain instead to the operation of the computer program. These have already been distinguished as replies and reports in section 4.4. An example of a reply, taken from *Adventure*, is: "You are standing at the end of a road before a small brick building." An example of a report,

from the same work of interactive fiction, is: “Do you really want to quit now?” This section continues this discussion. The goal at this point is to determine how the text of replies and reports actually occurs in an IF work’s output and to characterize the nature of these texts.

With few exceptions, a turn is usually either devoted entirely to a reply or a report. If the player inputs a directive, a report is almost always provided. A rare violation of this rule is seen in *Planetfall*, where, with humorous effect, the robot Floyd metaleptically says “Oh boy, are we going to try something dangerous now?” after a directive to save the game is entered. In the case of a command, a text that is entirely a reply is the typical response. An exception to this is seen when the traversal reaches an end and, after a reply is initially provided, a report offers the interactor the opportunity to restart, restore, or sometimes undo the previous turn. Another exception is seen when a score notification, often set off by special typography, is inserted into a text that is otherwise a reply. Specifically, the words “Your score has just gone up” are frequently used in these messages, which are often set off in brackets.

4.11.1 Reports

Reports refer to the program level and not the simulated world. They are extra-narrative, as are paratextual elements such as running headings in book or the indications of where stories jump to (for instance, “See OIL on A15”) in newspapers and magazines. Reports are often formulated according to standard conventions. In any interactive fiction piece, there are only a small set of possible reports. In general, this set is not extended nearly as easily or as often as is the set of replies; there are only a limited number of things that can reasonably be done at the program level, and these things tend to be the same whether the work has an adventure, mystery, science fiction, or contemporary milieu and regardless of the intricacy or simplicity of the interactive fiction world.

Since reports are essentially extra-narrative, addressed to the user or operator of the program rather than the narratee or player, generating reports is not part of the problem of narrating a simulated world. As a practical matter, there is not much reason to generate reports as opposed to just providing them as canned text. They might be expressed in a language-independent form for multi-lingual output, but the semantic content of a report does not need to be expressed differently when different sorts of narrative variation are being used.

It is possible to vary how explicitly the “I” and “you” of an interactive fiction are indicated in a report, just as this can be varied in narrative:

Do you really want to quit now?

Shall I end the game now?

Do you really want me to end the game now?

But in these three texts the “I” and “you” do not represent the narrator and narratee. Rather, they represent an interactive program and the user, or interactor. The same messages could be provided during a completely non-narrative game, or — with the word “program” substituted for “game” — for any other sort of computer program. Since reports are essentially extra-narrative, there is no obvious or straightforward way to vary signs of the narrating when expressing reports. The use of “I” and “you” in reports, when these terms are also used in the framework of narrating, may of course have an effect on the way the interactor perceives the narrative, which also uses “I” and “you,” but there is no direct connection to be considered in generation.

The extra-narrative nature of reports, and the fact that they are almost always produced apart from narrative expression, warrants the production of reports from canned texts via a separate module that does not deal with narrating. It may be ideal for this module to share the surface realization component of the Narrator, so that reports (along with replies) can be generated from representations that are more or less abstract. This could facilitate multi-lingual generation or generation across modalities at some point in the future. This work does not bear on the research questions involved here, however, so it has not been included in the project at this stage.

4.11.2 Replies

Replies can contain diegetic and non-diegetic texts. Diegetic texts consist of simulative texts (representing events and states that are simulated by the game) and non-simulative texts. Non-diegetic texts consist of texts pertaining to the interface, texts pertaining to the narrative, and other texts.

Simulative texts are either representations of events (active statements) or descriptions of existents (static statements) in the simulated world. A SENSE event, as described in subsection

5.1.3, explicitly calls for static statements about what the commanded character senses. There are other appropriate times to generate static statements, but for an initial formulation, it is reasonable to say that events of all sorts can occasion the expression of active statements, while SENSE events in particular can also occasion static statements.

Non-simulative texts are of the same two types. The only difference is that they represent events and existents that are not simulated and that lack a representation in the world model. An example can be seen in Meretzky's *A Mind Forever Voyaging* in a childhood scene which is replayed in a virtual reality joybooth: "The roller coaster whooshes out of the tunnel, and you scream, not from fear but from the exhilaration of the experience. Rav, scrunched next to you in the seat, laughs and laughs and laughs and Frita laughs and laughs and hands you some cotton candy and suddenly the roller coaster is gone and the park is gone[.]" Among other things, this text represents the event of the roller coaster leaving a tunnel, the player character (as a child) screaming, and two other characters, Rav and Frita, laughing. But Rav, Frita, and the roller coaster are not simulated; this text is an orthographic string that does not correspond to any events or existents that are part of the *A Mind Forever Voyaging* simulation. If this expression were to be generated from something other than canned text or an orthographic string, it would be necessary to either add the underlying events and existents as part of the simulation or come up with some alternate way for generating it.

Non-diegetic texts are sometimes better seen as pertaining to the interaction than to the narration. For instance, in Emily Short's *Bronze*, when attempting to enter the crypt without having performed the proper action, this text is provided: "There must be some preliminary, a matter of spiritual etiquette perhaps, to establish yourself as the master of those below." This does not appear to directly narrate anything, simulated or otherwise. It can be read as a comment on the interactive fiction world, as an indication that this is a world in which etiquette and dominance are important and one in which spiritual beings exist. While the text has this function, it also functions as a hint, letting the interactor know that something must be done before the crypt can be entered. This text implies that the crypt is not forever inaccessible, but must be appropriately unlocked by some sort of "preliminary." The "you" of this text is the player character. If the function of this text as a hint is taken to be the most important, the text is mainly addressed not to the narratee hearing the story of what has happened so far, but to the interactor,

who must figure out what commands to type in to solve the game.

In interactive fiction, texts often represent events or describe existents rather directly, but sometimes texts do refer to the narrative in an important way. For instance, in the sitting room in Michael Gentry's *Anchorhead* the command "examine eyes" can be entered, referring to the eyes of a figure in a portrait. The opportunity to do this occurs after the player character has had a dream. This text results: "They are the eyes from your dream, without question." Clearly, the dream has to have occurred as an event in the content plane for this text to be appropriate. The eyes in question, or at least a figure of some sort, should also have been described in the narrative when the dream occurred; the description of eyes or someone with eyes could not have been elided for this later expression to function in the same way. While the generation of this text would depend on the discourse history, this text is nevertheless occasioned by the description of an existent in the world.

Again, it is possible to have texts referring to the program or game level, that is, reports (such as score notifications) included in an output along with a reply. In all cases, the report is couched as something pushed onto, and then, in some cases, popped off, the stack of diegetic conversation. In the case of a score report, this program-level utterance being popped off the stack after it appears means that it cannot be referred by to a pronoun in a command. In the case of a "restart, restore, undo?" text at the end of a traversal, only a directive is possible in response, since the diegetic conversation has ended. For these reasons, it is not necessary to determine how to textually integrate reply and report, even though they sometimes occur together; they can be produced separately.

Many statements which, read literally, would not appear to be in the previous categories can be usefully seen as representing events or existents; at least, that is one of the ways in which these statements function. *Shade*, for example, can produce the sentence "It's amazing how much lack of sleep feels like a hangover, only without the preceding party." While this is explicitly just a proposition (and one that signals the existence of a narrator by suggesting that the proposition is amazing), it functions to represent that the player character feels hung over due to lack of sleep. It is not best understood as a generic statement about the world. In Jason Devlin's *Vespers*, a cask is described as follows: "It is big and empty. The Eucharist often gets larger in time of stress." The proposition about the generally observed enlargement of the Eucharist is, again, directly tied to

the description of an existent — as well as to the current stressful situation. In Gareth Rees’s *Christminster*, examining the punt results in this text: “A punt is a flat-bottomed wooden boat, about five yards long and one across. It is propelled on shallow and slow-flowing rivers with the aid of a long pole.” Although strictly speaking this text is a definition of the term “punt,” it clearly refers to the punt that exists in the simulated world of *Christminster*. It functions a description of this existent phrased as a definition, and, for instance, it warrants the interactor to ask about or refer to a long pole. The typical long pole may be absent here, as there is nothing promising that it is there, but this object has nevertheless been introduced into the discourse. In all three of these cases, it would be fine to generate such sentences as simulative, static statements rather than attempting to generate them as texts pertaining to the interaction, as texts pertaining to the narrative in general, or as other texts.

Past the prologue, it is actually somewhat difficult to find descriptive texts in interactive fiction output that do not refer to some specific existent. However, replies in interactive fiction do describe small and large sets of existents. One example: Part of the prologue of Joe Mason’s *In the End* is “Everyone in the room (and there are many, for Jon had many friends, right up to the last) sits dry eyed and stolid.” It is also possible for part of a reply to comment on the general situation in the interactive fiction world, which can be considered a description of the set of all existents.

Finally, just as a narrative can refer to what does not happen in it and things that do not exist in it (Prince 1988), it is possible for interactive fiction replies to comment on what does not occur and what is absent from the interactive fiction world. Perhaps the most famous example is the “no tea” that, amusingly, is implemented as an object in Douglas Adams and Steve Meretzky’s *Hitchhiker’s Guide to the Galaxy* and which (along with the “tea” object) plays an important role in a puzzle. The general problem of generating such disnarration from an underlying world model is not an easy one, and might require the addition of a “non-world model” that represents all the absent events and existents that are of interest and might need to be mentioned. (This would be an extension to the set of multiple world models, supporting multiple focalizers, that will be described.) Descriptions of individual existents and sets of existents, up to and including everything in the world, nevertheless account for a great deal of descriptive texts, so a system that is capable of generating texts grounded in these provides a useful starting point.

Any text not in these categories — something which does not refer to the diegetic world

(simulated or not), to the interaction, to the narrative, or to the program — should be possible to insert using canned text, as it should not be necessary to change it grammatically when world-specific text is recast.

4.11.3 *Generating Diegetic, Interaction, and Narration Texts*

Since “other” text can be canned and reports can also be canned (or the generation of these can be dealt with after this initial project focused on narrative discourse), there remain diegetic texts, texts referring to interaction, and texts referring to narration. All diegetic texts can be generated via proposed expressions (PEs) that are based on simulated events and existents. The other two sorts of texts, referring to the interaction and the narration, can also be generated based on the current turn’s PEs and the past history of PEs.

First, non-simulative texts can be generated by adding events and existents for the specific purpose of text generation, even if there is not opportunity to interact with these. In the case of *A Mind Forever Voyaging*, existents representing the roller coaster, Rav, and Frita could be added at the appropriate chronological time, during the game’s “backstory,” along with the appropriate events. Rav and Frita’s conversational abilities would not have to be filled in by the developer; it would simply be recorded that they laughed at a certain time, decades before the simulated world that contains the joybooth. From this, the appropriate text could be generated, just as text describing simulated, interactive events is generated. In this way, all diegetic text can be consistently generated, and even generated consistently in different narrative and discourse contexts. In both simulative and non-simulative cases, static statements can be generated based on a sensory perceptions and active statements based on events that occur in the immediate area.

A text referring to interaction is traditionally warranted by some description of an event or by a static statement describing an existent, both of which will occur in some PE. In the cases of a “novice mode” introducing IF, an adaptive hint system, and even an antagonistic anti-hint system trying to fool the player, whatever statements are produced should typically refer to some recent event or something in the immediate environment. In rare cases they will not — “You remember that you should have checked that area down south” — but these sorts of statements can still be conditioned on a particular past sequence of events, such as “20 turns after a particular event, if the PC has not visited to the specified area, generate this statement ...” The same sort of

preconditions for inserting a PE are used in both cases, whether the reference to the interaction is to happen immediately or after some delay. Both previous expression and previous simulation need to be taken into account. If the PC has acquired the object that the hint is about, it is probably a poor idea to generate the hint, even if this acquisition was not narrated for some reason.

Since the initial PEs (corresponding to events; prior to adding signs of the narrator, narratee, and narrating) provide a skeleton from which the narration of these events will be generated, they are also appropriate as a basis for generating expressions that refer to the narration.

4.12 Further Steps

This section has begun the discussion of the elements of interactive fiction from a theoretical standpoint, drawing mainly on narratology. The discussion here has only explored a few of the most important implications of clearly distinguishing the simulated from the non-simulated, the IF world from the text that describes it, and the diegetic from the extradiegetic. Perhaps the few points that have been made are at least adequate to demonstrate that a better perspective on IF can result from making such distinctions, however.

While this document will extend the initial theory presented in this section in a few ways, it will also be valuable to have more discussion of the nature of interactive fiction, beyond what is covered here, that

- focuses on specific works in giving examples of what the elements of the form are;
- distinguishes between elements in terms of their being formal or interpretive; and
- makes strong and meaningful claims that can be evaluated by others and, if useful, built upon by others.

As should be clear from the title of this section and from the discussion of puzzles, the intention here is to help begin a strong theoretical discussion of interactive fiction rather than to conclude it. There is much that remains in considering the nature of puzzles and how they fit into an IF work overall, relating to its aspects as program, potential narrative, world, and game. Clearly, people in the IF community are beginning to think about theorizing interactive fiction in a stronger way than can be done in casual discussion. There are certain to be benefits for interactive fiction if this more involved discussion of the nature of the form continues.

Approaches from other fields of study (narratology, to be sure, but perhaps also fields including game theory and game studies) can result in a better understanding of interactive fiction. These sorts of studies should also help IF authors and developers of IF systems advance the state of the art.

5 A New IF Architecture

Identifying the essential form and functions of interactive fiction provides a foundation, but a system for the development of interactive fiction, supporting research into interactive fiction and narratology, also requires an architecture informed by this perspective of potential narratology. The architecture of nn draws on well-established techniques for simulating an IF world and a standard pipeline architecture for natural language generation. It adds worlds representing the perspectives of focalizers, a division between content (and simulating) and expression (and narrating), and a modularized way of carrying out other IF functions.

In nn, representations of the world state, focalizer knowledge, and the discourse state are stored in models — namely, the IF Actual World, the Focalizer Worlds, and the Discourse Model, covered in sections 5.1, 5.2, and 5.3. These are updated by modules, which also handle input and produce output. The most important modules from a research standpoint are the Simulator and the Narrator, with the Recognizer also having some relevance. Section 5.4 describes the modules and their relationship to one another at a high level. The workings of the Simulator are covered in 5.1 alongside the IF Actual World. The Narrator is discussed in detail Chapters 6 and 7.

5.1 Elements of the World Models

As described in the next section in the overall context of nn's modules, the Simulator module manages the entire simulation, updating existents and adding events. This simulation is based on the existents and events of the IF Actual World, which is the base, authoritative model within the interactive fiction itself. (It is the "actual" world from the standpoint of interactive fiction actors.) This IF Actual World model is one of several world models, however; the models also include a Focalizer World, based on perceptions and experiences, for each actor. Each Focalizer World represents one actor's theory about the actuality, or "reality," that the IF Actual World encodes.

The central function of the Simulator as it acts on the IF Actual World is modeling the necessary physical aspects of the interactive fiction. There are various techniques that can be used in the system to model the interior mental states of characters in lightweight ways, but this sort of representation is not the purpose of the current Simulator. The Simulator was developed to be

able to flexibly and easily represent things such as the physical movement of objects and the configuration of a space. Detailed emotional and mental models of characters are compelling topics for research and have been rather well-studied; while they may be integrated into nn at a later point, the internal motivation for characters' actions was not a focus of this research project. Characters in the first-generation nn system simply act in whatever ways game authors have programmed them to act; no matter what they do, the narrator module must be able to narrate their actions, and everything else that transpires in the simulated world, appropriately.

The IF Actual World and Focalizer Worlds each have their own tree of existents and their own list of actions, connected by causality and marked as happening in time, with each action being defined in terms of a small, closed set of primitive events. Inasmuch as the Simulator's events are language-independent and primitive, they relate to the primitive actions of conceptual dependency (Schank 1975). While the concept of primitive event is useful, the Simulator does not borrow much else from conceptual dependency theory. The set of events is fixed; this set and how it is employed in the Simulator strongly emphasize the physical world rather than trying to account for things like mental states. Events are not related by conceptual dependencies, but by causal entailments; preconditions are checked before events can occur and complete postconditions are available. The idea of this representation is not to encode natural-language knowledge but to run an underlying simulation of a story world at an appropriate level of detail.

Sequences of events are wrapped in a higher-level action, which helps to connect the lowest-level representations to intentional, interpretable behavior. The Simulator takes as input an intended action which may or may not succeed; it outputs whatever actions are simulated. These may be different from the intended action if the intended action does not go as planned; the output may also contain additional actions performed by other actors and actions that are a consequence of other actions, rather than being a result of commands the player has entered.

In explaining the Simulator, the following discussion takes a bottom-up approach, focusing first on the core simulation that is run on the IF Actual World. The representation of existents is explained first, then that of events and actions, including those that represent changes in the state of these existents. After this, the system's model of perception is described. This section concludes by describing the way that existents and events are propagated to Focalizer Worlds using this model of perception, and the way that Focalizer Words are defined to enable retrospective

narration rather than simulation.

5.1.1 Modeling Existents

Existents represent matter — physical, persistent objects in the simulated world. They are either things, actors, or rooms. A special existent called COSMOS is privileged in this model. The children of COSMOS are the rooms, representing distinct physical spaces. The children of these rooms may be either things or actors, and these may have their own children, and so on; in *nn*, existents form a tree. The basic idea of a graph of existents is fairly standard in IF. An object tree is used, for instance, in the popular system Inform (Nelson 2001), which employs one of the capabilities of the 1979 Z-Machine to represent objects in this way.

These are the three types of existents in *nn*:

Thing. Inert objects, if they do not need to be represented as individual parts or in finer-grained ways, exist in the simulation as *things*. Things may be of any size and may be stationary or movable.

Room. A physical location of a reasonably small size (such that everything in the area can be easily seen, touched if it is not inside something else, picked up if it is portable, and manipulated if it is manipulable) is represented as a room. These discrete clumps of virtual space are convenient for the implementation of IF even when the space to be simulated is not inside a cave or building. Rooms are conventional in interactive fiction, as previously described. Rooms, unlike actors and things, may have passages leading off in different directions, to other rooms. Visibility is not restricted to the current room in *nn*, although by default the existents that are manipulable and those that are visible are both defined by containment in a room.

Actor. The only existents that can initiate actions are actors. For this reason, the special existent COSMOS is an actor: happenings, such as “it started to rain,” are represented as being initiated by COSMOS. Actors need not be anthropomorphic characters; they are simply existents that can generate actions, either in reaction to events or spontaneously. The player character is an actor, but a laser printer that spits out a printed sheet is also an actor.

Existents are instances of the Thing, Room, or Actor classes, or of a subclass of one of these

that has been augmented in some way. For instance, a different act method can be added to a subclass of Actor, allowing all actors instantiated from this subclass to perform arbitrary actions according to this new method. Existents have several standard properties:

called — a tuple containing, first, the article that is used when the existent is first mentioned.

The last element is the standard noun or noun phrase used to refer to it. Optionally, in between these, it is possible to add one or more elements containing modifier phrases which can be used by the Narrator for variety.

parent — a pair containing the relation and the parent of the existent, present for every existent except COSMOS. This attribute is described in detail in the next subsection.

adjs — a list of words which are understood as adjectives referring to the existent. Every adjective that appears in *called* should, of course, be in this list, so that the interactor can refer to the existent as the computer-generated text does. An element of *adjs* by itself is not adequate to designate an existent, but it can be useful, or even necessary, in distinguishing between two existents that are similar.

nouns — a list of words which are understood, by themselves, as referring to the existent.

The nouns in *called* should be in this list.

appearance — a string-with-slots that tells what the physical appearance of an object is. The “description” in Inform and TADS is similar but often contains other information and commentary. These should not be encoded in appearance because this string-with-slots may be used to describe what other actors besides the focalizing actor see and it may be used when recounting past events.

qualities — a list of qualities that may have to do with an existent’s function and appearance or with the materials it is made from. *adjs* and *nouns* are extended so that words appropriate to this quality are included. This allows an existent, for instance, to be designated as having the DEVICE quality, so that “device,” “mechanism,” “machine,” and “apparatus” can all be used to refer to it, without any of these needing to be explicitly added to *called*.

state — a dictionary of all sorts of information about the physical properties of an object, some of which is standard and some of which may be the result of author-added states. This includes whether an object is open or closed and whether it is locked or unlocked, if

applicable. It also includes how visually prominent an object is and how much light it is radiating. For actors, the gender of the actor and a script of things to do can be stored here.

5.1.2 The Tree of Existents and Relations

The tree representing existents and their relations to one another is a quintuple (c, R, A, T, E) where:

c is the root existent, COSMOS, which is an actor.

R is the set of all rooms, $r_1 \dots r_n$.

A is the set of all actors except c (COSMOS), $a_1 \dots a_n$.

T is the set of all things, $t_1 \dots t_n$.

E is the set of child-parent-relation triples (x, p, l) with:

$$x \in R \cup A \cup T,$$

$$p \in c \cup R \cup A \cup T / x \cap \text{desc}(x),$$

$$l \in L, \text{ the set of all relations.}$$

This set E corresponds to all of the edges of the tree, each edge connecting some x with its parent p and annotated with the relation l . As specified here, the set of all existents $X = c \cup R \cup A \cup T$, with R , A , and T disjoint and none containing c .

Since each existent except COSMOS has a unique parent, there is exactly one child-parent-relation triple for every existent in $R \cup A \cup T$, so $|R \cup A \cup T| = |E|$. Every room has COSMOS for its parent and is in the IN relation: $r \in R \Rightarrow (r, c, \text{IN}) \in E$. Actors and things, on the other hand, may not have COSMOS as their parents: $x \in A \cup T \wedge (x, p, l) \in E \Rightarrow p \neq c$. This guarantees that actors and things will always have a room as an ancestor and will ultimately be contained in some room.

The special root existent, COSMOS, has no parent and is at level zero in the tree. All the rooms, and only the rooms, are at level one. Actors and things are at the lower levels in this model.

The edge connecting an existent and its parent is annotated with a particular relation. Each existent has some unique relation to each parent. The standard relations are IN, OF, ON, or PART:

IN — Describes a typical thing sitting in a room, which can be picked up and moved about.

Also, IN is used when a thing is inside a container of some sort. Finally, IN is also used for the relation that rooms are in with COSMOS.

OF — Indicates a possession of an actor. If an actor owns something and is carrying it around, it is OF that actor.

ON — Used for a thing that is worn by an actor or placed on some supporting surface such as a table.

PART — A fixture in a room will be in the PART relation with the room. Similarly, some feature of the landscape will be a PART of the outdoor room in which it appears. A PART cannot be picked up or moved.

A child can change its relation to a parent while still remaining a child of that parent. For instance, the player character taking off a coat moves the coat from the ON relation to the OF relation, but the coat remains a child of the player character.

These are the standard relations provided, and many existing pieces of interactive fiction should be implementable using only these. The set of relations is open, however, so an author might add one such as ORBITING or IMPLANTED for use in a specific game. The parent-child relation indicates physical proximity and some sort of dominant or controlling relationship, such that an action done to the parent will have implications for children. Typically, an existent will move if its parent is moved — shoving a desk to another room moves everything in or on the desk, and an actor walking to another room means that the actor's possessions move to the new room as well.

All the existents other than COSMOS have a single parent which is neither the existent itself nor any of its descendants — that is, the parent p of x is in the set $c \cup R \cup A \cup T / x \cap desc(x)$ — so the graph of parent-child relations is acyclic (a tree) and spans all existents. Restricting existents to having a single parent provides a reasonable model of an environment in many ways and makes the traversal of the graph (to find the parent container or the room in which an existent is located) simple.

Relaxing this restriction to allow multiple parents, or a parent that is a set of existents, could be helpful in some cases. Modeling doors, a tricky case in interactive fiction, illustrates the value of allowing sets of existents to be parents. A door divides two rooms, typically — for example, a porch and a foyer. If the door is unlocked, it can be opened from either room, so it should be

accessible from both. If it bars the way from one room to another, it does the same in the other direction. If it were to be torn off and carried away, it would no longer be the child of either room. Other passageways, even ones that cannot be locked and unlocked, opened and closed, similarly connect two rooms and are best represented as children of both. There are other things (such as a long rope, a classic example of a tricky object to implement in interactive fiction) which might be located in multiple rooms and best modeled as having multiple parents. Multiple parents could even enhance the representation of small things that can be supported by two or more things — such as a board that can be placed on two cinderblocks to form a shelf. The situation of a large object being carried by two people, or a larger group of people, could also be better represented if multiple parents were allowed. The loss of the tree property might not be inherently problematic, as long as the graph remains acyclic. However, the current system was implemented using the assumption that existents (except COSMOS) have exactly one parent and that they are, therefore, in at most one room. Because of this, and because this issue is not central to the problem of narrative variation, a possible extension to multiple parents has been left for a later stage of development.

Non-hierarchical relationships are not directly represented. This seems to be a reasonable limitation of the model, because the relations that are represented are ones that indicate one existent's containment of, support of or influence over another. They are part of the model so that the effects of actions on one existent can be easily determined for subsidiary existents. Whether a thing is in the same container or room can be computed easily in this representation scheme, and does not have to be encoded directly as an attribute of the object. Siblings can be easily found by ascending to the parent and requesting the list of the parent's children. All objects in physical proximity can be determined by ascending to the room and requesting all the descendants of the room.

Existents can have state in this scheme thanks to the *state* attribute. A lamp, for instance, might have a Boolean-valued LIT state with true corresponding to on and false to off. An oven could have an integer- or real-valued TEMPERATURE state that corresponds to how hot it currently is. Some standard states are provided, as with relations, but the set of states is open for authors to augment. The standard states include OPACITY (a value between 0 and 1) and LIGHT-LEVEL (a value between 0 and 1), to represent to what degree an object is transparent or opaque

and to what degree it radiates light. In both cases, a default value is provided so that the IF author does not have to continually indicate that things are opaque and that they do not provide light; a fully lit world where light sources are never an issue will be provided if the defaults are used. Another standard state is PROMINENCE, indicating how visually obvious an object is. Small objects that do not have an unusual color will have low prominence (near 0), while huge and very evident things, such as monuments or large geographical features, will have a prominence approaching or equal to 1. There are also the Boolean states OPEN and LOCKED and the state KEY, which holds a string corresponding to the thing necessary to unlock the thing with the KEY state. If OPEN is specified, an object can be opened and closed. If OPEN and LOCKED are also specified, an object can be locked and unlocked as well. KEY is optional (there are things, such as bathroom doors, which can be locked and unlocked without keys) but if it is present, it will additionally specify that a thing is necessary to change the LOCKED state and will specify what that thing is.

5.1.3 Modeling Events

The model, tailored as it is to represent the physical world, represents detailed, low-level occurrences as events. Any step of a higher-level action that might succeed in one context and fail in another is represented as an event. Events are considered as atomic representations of occurrence. They may seem unusually specific if narrated, but they are very well-suited to the simulation of the model world and useful in checking how an action progresses within this simulation.

There are five basic events, one special event, and two special non-events. Only two of the five basic events change the state of the world. All events have some arguments in common that deal with when they begin, how long they take to complete, and what caused them; these are not mentioned below. The event types are as follows:

MODIFY changes the state of an existent. A light is switched on or off by this type of event.

An existent x — specifically, one of its states in the *state* property — is changed by MODIFY, but the typology and nature of relations in the tree (represented by the set E) are not changed. A MODIFY has the agent, the object, the state, the old value, and the new value as arguments.

CONFIGURE changes the world tree in some way, either by modifying a relation or by moving an existent to have a different parent and to be in some relation with that parent. All sorts of physical movement of existents are represented by CONFIGURE events. The typology and nature of relations in the tree (represented by the set E) are changed by CONFIGURE; specifically, one child-parent-relation pair (x, p, a) is changed. A CONFIGURE has as arguments the agent, the object, the old relation, the old parent, the new relation, and the new parent.

MISC represents some event that, by itself and as far as the simulation is concerned, has no effect on the world. Speaking a word or waving something around provide examples. If the actor is alone when doing these things, nothing may be entailed. If some communicative purpose is accomplished, a MODIFY event may be entailed that changes the state of the actor listening. A MISC often represents something that, although it does not matter to the simulation, should be narrated, either as part of a larger action or by itself. The arguments of MISC are the agent, the verb characterizing the event, the object, the preposition relating the event to the indirect object, the indirect object, and the utterance. Only the first two are required; the utterance will only be used if the MISC is an action to vocalize something.

IMPEL represents application of some amount of force to an existent in a particular direction. Almost any physical action begins with an IMPEL, but one specific example would be pushing a desk. By itself, the IMPEL event does not change the world state (for instance, if the desk is too heavy to move the IMPEL event would succeed but would have no consequence), but the world state will change in the next step if IMPEL entails a MODIFY or CONFIGURE event. The arguments are the agent, the object, the direction, the force, and the manner. The direction is sometimes irrelevant and can be omitted in these cases.

SENSE represents focused sensory attention; reports to the player character about what can be perceived are provided because a SENSE event occurred. The SENSE event cues an update of existents in an actor's focalized world but does not itself carry out this update in the same way that the CONFIGURE and MODIFY events do the work of updating the IF Actual World. The SENSE event relies on determinations of visibility when sight is the

sense used. The arguments are the agent, the object, and the sense which is being attended to.

A few special cases — first, a special event:

BUNGLE represents a failed attempt to do something. The failure to do something is important both because the attempt may have implications in the simulation (it may entail some additional actions) and because it needs to be mentioned by the narrator. BUNGLE is a real action in the world, but it cannot be intended, only caused by the failure of some other action. BUNGLE does not cause anything to happen by itself; MODIFY or CONFIGURE need to be entailed by it for that to happen. A BUNGLE event has as its argument the event type which was attempted unsuccessfully along with all the arguments for that failed event.

And next, two non-events which are treated like events by the system:

REFUSE represents the player character's refusal to do something. For instance, if there is a solid wall to the east that can clearly be seen, a *walk east* command from the interactor would usually be met by refusal by the player character. The refusal certainly needs to be mentioned when focalizing and commanding the character who has refused to do something, but this non-event differs from a BUNGLE. It cannot cause anything to happen in the world — it is the same as if the player character had simply done nothing. As with BUNCLE, REFUSE has the refused event as an argument; the reason for refusal is another argument.

CONCLUDE indicates that some conclusion has been reached and that neither the simulation nor the narration should proceed beyond this point in time. When CONCLUDE is encountered by the Narrator, nothing else is narrated and the program terminates. CONCLUDE has no arguments.

To focus on one of these event types, IMPEL, and how and when a particular event of this sort might be instantiated and used: When a player gives the command "touch the dumbwaiter" the first event in the intended action would be IMPEL (*pc DUMBWAITER - 0 touch timestamp cause*). Most physical actions begin with IMPEL events since it is necessary to exert force to accomplish much of anything in the physical world, although there are exceptions such as looking around. "pc" stands for whatever the tag of the current player character is; this is the actor. The object is

that existent with the tag "DUMBWAITER". The "-" that is provided as the next argument means that the direction of force is irrelevant or not applicable. The "0" means that there should be the minimum amount of force that can be applied; the PC is simply trying to sense the dumbwaiter and not to do something to it such as opening it (which would involve at least a small amount of force) or destroying it (which would involve more). An argument indicating the manner in which the force is applied occurs next, "touch." The arguments taken by different events vary, but the last two are always the *timestamp* and *cause*, as shown here.

Events are not only put together as a consequence of user input. An actor can generate an event like this one automatically without creating the string "touch the dumbwaiter," by simply instantiating a complete, assembled action with this event within it. Alternatively, an autonomous actor can use the same command-to-action mapping that is used when interactor input is converted into an action. This involves generating an already-recognized command, such as (TOUCH DUMBWAITER), which functions as an intended action, as described in the next section.

5.1.4 Modeling Actions

Events, by themselves, are useful for tracking the world and the changes taking place in it, but are often too detailed, and too disconnected from high-level concerns to be suitable for narration. They are also not ideal when it comes to commanding a character with an input string. They represent an occurrence, encapsulating information about the time of that occurrence in the story world and its causality, but do not represent intentionality. To provide a model at a higher level than that of the event, the *action* is introduced.

The Simulator receives an *intended action* from the recognizer, which is a semantic representation of what is intended by a command. For instance, "pick the yellow cake up", "GET BUNDT" and "take cake" might all correspond to the action (TAKE CAKE1), where CAKE1 is the identifier of an object. There is no guarantee that this action will correspond to something possible or something that the commanded character will be willing to attempt; this is simply a formal representation of what the command is understood to indicate as the interactor's intention. The term *intended action* relates to the first of philosopher Alfred Schütz's concepts of action, *Selbstverstehen*. This sort of action is described by Meister: "In the mode of *Selbstverstehen*,

or subjective understanding, action is the design which precedes a person's behavior and which he imagines is already completed" (Meister 2003:62).

Intended actions can include sequences of events, which are marked as essential or peripheral. These are distinct primitive events which are all intended for the same purpose, the completion of the intended action. The essential events must succeed in order for the intended action to be successful.

For example, the action GO WEST, called in ROOM1 when ROOM2 is to the west, would set up two events: IMPEL (pc pc - 50 *timestamp* cause), CONFIGURE (pc pc IN ROOM1 IN ROOM2 *timestamp* cause). The IMPEL event may fail, perhaps because the player character's shoes are superglued to the floor. Or the CONFIGURE event may be reached but may fail, perhaps because the player character is able to impel himself but an invisible wall of force blocks westward progress. In either of these cases, the overall GO WEST action does not succeed. Some events may entail actions (for instance, looking around at the new room one has entered). The success or failure of these entailed actions does not affect the original one; nothing that happens after an event succeeds or fails can change that. Events are generally all essential, although there are exceptions, some of which are discussed next.

More generally, an action contains a sequence that may include not only events, but also other actions, each of which is also essential or peripheral. This allows for general, obvious activities with their own event structures to be planned and incorporated as part of an action. For instance, if a player character is moving between her bedroom and her living room, to the north, the GO NORTH action might entail an OPEN DOOR action, an IMPEL event, a CONFIGURE event, and (if this is the sort of player character who typically closes doors after herself) a CLOSE DOOR action. In this case, the CLOSE DOOR action, if included, would be marked as peripheral: a failure to close the door would need to be remarked on, since it would be unusual, but it would not prevent the entire GO NORTH action from succeeding.

As the system is developed further, there is nothing to restrict actions that are included within an action from containing, in turn, their own actions. A door might need to be unlocked before it is opened, so the action of unlocking it might be included at the beginning of the action representing the opening of the door. There are several reasons to restrict the depth of this tree that are related to good interactive fiction design and to practical concerns in both narration and

computation, as discussed at the end of subsection 5.2.3. In a system with a more complete model for the mental states and motivations of non-player characters, a hierarchy of actions might begin at a very high level (“go on a quest,” “get revenge”) and be decomposed several times into other actions that are still fairly high-level (“search the village for information,” “interview the librarian”). In the current system, non-player characters have their behavior hard-coded in a less systematic fashion, while the interactor must figure out what the player character’s overall plan for action is. By allowing nested actions, actions that are “obvious” can be included, perhaps even planned, to facilitate interaction, while very high-level commands can be declined.

5.2 IF Actual World versus Focalizer Worlds

The focalizer of a narrative, or of a portion of a narrative, is, very loosely, the character who the narrative follows. A focalizer is similar to what has been called, in the context of the novel, a “central consciousness” (James 1972). Genette builds on work from the tradition of New Criticism (Brooks and Warren 1943) in distinguishing the “character whose point of view orients the narrative perspective” from the narrator (Genette 1980:186). A

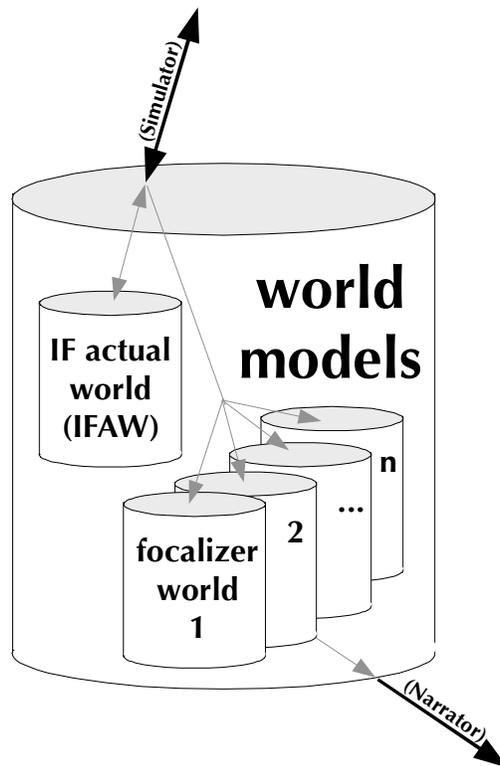


Figure 1: Interior structure of the world models component. The simulation runs based only on the IF Actual World, but the Simulator also updates all of the Focalizer Worlds. The Narrator then bases its narration on one of these Focalizer Worlds — in this example, that of focalizer 2. An external, “omniscient” Narrator would base its narration on the COSMOS’s Focalizer World, which models everything in the IF Actual World. The Joker’s ability to read from and reset the world models is not represented here.

narrative does not have to be told by a character for it to be based on what that character sees, experiences, and knows.

In narratology, narrative itself is conceived of as being independent of medium and not restricted to language (see Ryan 2005). At the same time, there are clearly special capabilities that different media afford, or at least different strengths that they each have. While the movie camera can show events from the viewpoint of a particular character, text affords certain special abilities as regards focalization and the subjective perception of a focalized character.

As Shlomith Rimmon-Kenan argues: "The indeterminacy and abstract nature of language ... renders literature more amenable [than theatre or cinema] to the representation of dreams, hallucinations, and the like" (1989:162). ... Since nonverbal media cannot express modalities, they impart an implicit existence to whatever they represent. Language is unique among semiotic codes in its ability to assign propositions to private domains, to express the forking paths of plans and projections, to represent possible worlds in their alterity, to state what motivates agents, and even to express causal relations. ... (Ryan 1991:266)

To begin to exploit some of these special capabilities of language, the IF system developed here incorporates in its world models not only a representation of simulated reality but also a simple representation of what each of the potential focalizers knows about that world. While the existing representation is limited to focalizers' knowledge about the simulated world, this set of per-focalizer worlds is drawn on a basic insight from Marie-Laure Ryan's application of modal logic and the theory of possible worlds to the concept of fictionality (Ryan 1991). Ryan describes a fictional world not as a defective or limited version of reality, but as the new center of a re-centered possible worlds system. This allows for explanations of complex phenomena, such as how fictional characters can lie and be mistaken relative to the world projected by the fictional text. It even provides a model of how texts that project their own distinct fictional worlds can exist within a fictional text, and what ontological status these texts-within-texts have. As Ryan explains:

In contrast to modes of expression that refer to the non-actual in a hypothetical mode, such as if...then statements, fiction includes both factual and non-factual statements. The former outline a textual actual world (TAW), while the latter allude to the virtualities of the fictional system. The contrast actual/non-actual is thus reinscribed within the textual universe. ... Through the concept of playful recentering this proposal reconciles the indexical theory of actuality proposed by Lewis with Rescher's absolutist view. From the point of view of the 'actual actual world' the worlds of fiction are discourse-created non-actual possible worlds, populated by incompletely specified individuals; but to the reader immersed in the text the TAW is imaginatively real, and the characters are ontologically complete human beings. (Ryan 2005)

To take the first step toward the representation of more complex phenomena such as these, an

interactive fiction's central reality — the world which is simulated by the Simulator in this project's architecture, and which is actual from the standpoint of actors within the IF — is represented in software as, following Ryan's terminology, the "interactive fiction actual world." In this text, "IF Actual World" is used to refer to the part of the interactive fiction system that is a model of this world. In a state-of-the-art *circa* 2006 IF system (TADS 3 or Inform 7), there is one interactive fiction world model. This single model represents the IF Actual World; there is no separate representation of how focalizers perceive or imagine it. There is IF in which *ad hoc* code and text provide for focalization changes of some sorts, such as ways of seeing the same world through different characters' eyes in Michael Berlyn's *Suspended* and Sean Barrett's *Heroes*. But there is no standard way of representing a focalizer's view of the IF Actual World.

The world models of nn include not only the model of the interactive fiction actual world but also additional worlds representing focalizers' individual perspectives on and beliefs about their environment and what has happened in it. These Focalizer Worlds derive from the IF Actual World and, because they are implemented as individual worlds rather than as subsets of the IF Actual World, they can, in theory, deviate in arbitrary ways. This allows for characters to be completely wrong about the world, not just limited in their information about it. The rule for updating the Focalizer Worlds that is currently implemented only allows for a partial but otherwise accurate view of the IF Actual World, since the update is based on that "reality." However, a different update rule could set up Focalizer Worlds that have existents not in the IF Actual World and which represent hallucination, mistaken belief, misperception, and misrecognition. Whatever the state of a Focalizer World, however, it cannot directly influence the IF Actual World. Focalizer Worlds are theories about this core "reality" or actuality which may indirectly influence the IF Actual World by being the basis for an actor's behavior, but which do not affect it directly.

5.2.1 *Updating Focalizer Worlds*

The Simulator runs the IF Actual World, updating the existents in it in accord with events that transpire. Updating the Focalizer Worlds, which represent theories of this actual world that are held by characters, is a different sort of task. It would have been reasonable in some ways to implement a separate module for maintaining the Focalizer Worlds, because this functioning

does not include running a simulation and updating existents based on events, but rather, determining which events and existents (in what states) are perceived by which potential focalizers. However, the different modules of the system — including the Recognizer, Simulator, and Narrator — all deal with an entire turn at a time. The perception of events, and the integrating of that perception into a Focalizer World, is more appropriately done after each of the actions in the simulation transpires, so that actors are represented as perceiving what is happening during a turn. Because updates to the Focalizer Worlds are best done contemporaneously with the the simulation, the Focalizer Worlds are updated as the Simulator runs, each time the IF Actual World calls the *do* method of an action to have that action performed and its consequences recorded. Focalizer Worlds are updated directly with actions, which are higher-level and more meaningful than individual events.

5.2.2 *Use of Focalizer Worlds for Narrating*

The main motivation for the Focalizer Worlds is their use in narrating: They allow generation of narrative from a particular actor's perspective, limiting the information that can be used to what that actor has perceived during the simulation and what that actor knows to begin with. Practically all IF systems support one type of focalization — limiting what is told to the perspective of the player character. In nn, any actor can be focalized at any time, and events that have already transpired can be retold from the perspective of any actor.

A simple case of varying levels of knowledge can be demonstrated in a slightly modified version of *Adventure*. In the modified game, the adventurer can be focalized, as usual, while other actors can also be focalized, including the pirate, who has been moved to an above-ground location and acts a bit differently. The things that have happened can also be recounted in this version. Using this IF work, it is possible to see how two different characters can recount what they have seen, including some overlapping events and some that are different. Here are the exciting events that transpired over the first 10 turns, from the adventurer's perspective:

The pirate waved.
Meanwhile, you went to an unknown location.
Then, you looked at the building's interior.
You were inside the building, the well house for a large spring.
You saw that the building's interior contained the brass lamp, the bottle, food, and the glinting keys.
Then, you took the glinting keys.

Then, you picked up the brass lamp.
Then, you conveyed yourself to the end of the road.
Then, you looked at the end of the road.
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.
You saw that the end of the road contained the pirate.
Then, the pirate waved.
Meanwhile, you went to an unknown location.
Then, you looked at the valley.
You were in the valley in the forest beside the stream tumbling along the rocky bed.
Then, you conveyed yourself to an unknown location.
Then, you looked at the slit in the streambed.
At your feet all the water of the stream splashed into a 2-inch slit in the rock.
Downstream the streambed was bare rock.
Then, you headed over to an unknown location.
Then, you looked at the area outside the grate.
You were in a 20-foot depression floored with bare dirt. Set into the dirt was the strong steel grate mounted in concrete. A dry streambed led into the depression.
You saw that the area outside the grate contained the strong steel grate.
Then, you unlocked the strong steel grate.
Then, you opened the strong steel grate.
Then, you went through the strong steel grate.
Then, you went to an unknown location.
Then, you looked at the area below the grate.
You were in a small chamber beneath a 3x3 steel grate to the surface. A low crawl over cobbles led inward to the west.

Here are the same 10 turns recounted with the pirate focalized. In this version of *Adventure*, the focalized character automatically becomes the narratee, the “you” of the narrative, so the pirate is referred to as “you”:

You waved.
Meanwhile, the adventurer walked to the building’s interior.
Then, you waved.
Then, you waved.
Then, you waved.
Meanwhile, the adventurer walked to the end of the road.
Then, you waved.
Meanwhile, the adventurer conveyed himself to the valley.
Then, you waved.
Meanwhile, a dwarf headed over to the end of the road.
Then, you waved.
Then, you waved.
Then, you waved.
Then, you waved.

The pirate (evidently under the impression that he is on a parade float or part of a Disney attraction) simply waves all the time, standing at the end of the road. He sees the adventurer enter the building, emerge from the building, and go south to the valley, but isn’t aware of anything else the adventurer is doing — acquiring things in the building, unlocking and opening the grate, or going to any other location. The pirate also sees, after all of the adventurer’s

perceived actions, that the dwarf walks to the location where the pirate is standing. The adventurer, who is elsewhere, doesn't notice this, so it is not narrated in the first recounting.

The Focalized Worlds not only model the knowledge arising from different perceptions of actors; they also have the ability to model differences in what actors know to begin with. In this modified version of *Adventure* a new character, the dungeon master, has been added. This character can also be focalized. Initially, the adventurer, located at the end of the road, knows nothing about the layout of the world and about what physical spaces are where. As the adventurer is commanded to move around, her knowledge increases and messages such as "South is the slit in streambed" are output, reflecting this. Messages indicating the player character's increasing knowledge of the geography of the world are indicated in bold below. (The pirate was moved from the end of road in this instance of the game.)

Welcome to Adventure!!

>look

You look at the end of the road.

- End of the road -

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.

>go east

You head over to the building's interior.
Then, you look at the building's interior.

- Building's interior -

You are inside the building, a well house for a large spring.
You see a shiny brass carbide lamp, a glass bottle, food, and some metal keys.
West is the end of the road.

>go west

You walk to the end of the road.
You look at the end of the road.

- End of the road -

You are standing at the end of the road before the small brick building. The small stream flows out of the building and down a gully.
East is the building's interior.

>go south

You walk to the valley.
Then, you look at the valley.

- Valley -

You are in a valley in the forest beside a stream tumbling along a rocky bed.
North is the end of the road.

>go north

You head over to the end of the road.
You look at the end of the road.

- End of the road -

You are standing at the end of the road before the small brick building. The small stream flows out of the building and down the gully.
South is the valley. East is the building's interior.

The dungeon master, on the other hand, has a full copy of the physical locations of the world in her world model. Focalizing the dungeon master and looking around reveals where every exit leads. If the interactor commands the dungeon master to move to a new location, the dungeon master's complete knowledge of the world will result in a description of where every exit leads from this new room. Messages indicating the player character's pre-existing knowledge of where everything is located are indicated in bold below.

Welcome to Adventure!!

>pfm pc dm

<<< Focalization is set to dm. Commanded character is set to dm. >>>

>look

You look at the awkward canyon.

- Awkward canyon -

You are in an awkward sloping east/west canyon.
West is the bird chamber. Down is a debris room.

>go down

You walk to the debris room.
You look at the debris room.

- Debris room -

You are in a room filled with debris washed in from the surface. A low wide passage with cobbles becomes plugged with mud and debris here, but an awkward canyon leads upward and west. A note on the wall says "MAGIC WORD XYZZY."
You see a black rod.

West is the awkward canyon. East is the cobble crawl.

5.2.3 Potential Use of Focalizer Worlds for Planning

The current system provides for a basic model of what characters know as distinct from actuality, but it does not model possible worlds based on an accessibility relation or use modal logic to reason about possibility. An intriguing possibility for the future would involve combining artificial intelligence work on using modal logic to model character knowledge and possible actions (Moore 1980) with work on using possible worlds to explain the workings of fiction and narrative (Ryan 1991).

Planning based on focalizer worlds could, of course, enable more intelligent planning for non-player characters (the actors who are not being focalized and commanded) and allow such characters to act in more flexible, appropriate ways. For instance, actors could be assigned the goal of preventing or facilitating a particular action by the player character, or a particular change in the state of the world, and could plan to achieve this goal based on their own Focalized World, not based on the IF Actual World. This would mean that a player character might have to conceal things from or reveal things to others to avoid being thwarted or so that another actor knows enough to help. A non-player character also could be set up to have a poor sense of direction and could be unable to remember the map correctly. This would cause the character to get lost (in a realistic way) while assiduously seeking a particular location.

There are a few reasons to believe that planning based on Focalizer Worlds could benefit even the interactor's relationship to the player character. While an IF author would not usually want to automate the player character and have him or her carry out high-level goals, a planning system could fill in certain routine tasks so that the interactor would not have to type very detailed, low-level commands to accomplish things. A simple example: With appropriate use of a Focalizer World, a player character who knows the layout of a building would be able to go to any room via a known route when commanded to do so, while such a player character would not take a route that have never been traveled before, and the same player character in an unknown space would not be able to navigate it automatically.

It is generally useful to have actors capable of doing a small series of routine actions as part of a higher-level command — to have the necessary actions entailed based on a command, when all the means to do those actions are available. As mentioned in the section 5.1.4, if a player

character is standing by a closed but unlocked door that leads west, it is convenient to have *go west* entail the (OPEN DOOR) intended action, so that the player character first opens the door and then proceeds through it. (And, depending upon what sort of character we have, a (CLOSE DOOR) intended action might be entailed, too.) Furthermore, if a locked door is to the west and the player character has the key — and knows that this is the key to the door — it is reasonable to also entail an (UNLOCK DOOR WITH KEY), before the (OPEN DOOR). If the key is in the player character's pocket, it is reasonable to entail another intended action before and of these, (TAKE KEY). Of course, the planning process should be constrained in some ways to prevent arbitrarily long sequences of actions being entailed by simple commands, but there are certainly advantages to a system like this. Using individual rules for what may be entailed, Emily Short's *Bronze* is able to entail the equivalent of (TAKE KEY), (UNLOCK DOOR WITH KEY), and (OPEN DOOR) actions as described here, making movement through the virtual space less tedious.

If a system were to allow these sorts of entailed actions to be determined automatically, it could greatly simplify IF development and improve the IF interface. However, basing such entailment on a single world model rather than on a character's particular Focalized World would be untenable. Consider the case of a player character who is standing outside a cave for the first time and (unbeknownst to this character) needs to say "open sesame" to have the cave open and to be able to go north into the cave. If the planning process simply searched for what needed to be done to open the door, it would find that a simple speech act suffices and would automatically generate (SAY "OPEN SESAME") as an intended action. This means the system would solve a puzzle rather than letting the user solve it; also, it would present the false idea that the player character knew how to cause the cave to open. The standard scheme for planning is not exactly too smart, but it is using the wrong knowledge base: an omniscient one rather than one that represents the player character's knowledge. If the planner is pointed to the player character's focalized world instead, it will be unable to find this action, because the player character does not initially know that (SAY "OPEN SESAME") in this location has as a postcondition that the cave is open.¹¹ The knowledge that (SAY "OPEN SESAME") causes the

¹¹ Some authors might want to use this facility to prevent the interactor from guessing the password, "open sesame," based on the player character's Focalized World, which would indicate that the PC doesn't know the password yet. From a game design standpoint, this would be an extremely tedious thing to do — it precludes experimentation and guessing as a way to solve the puzzle — but it could be done. A more interesting idea would be varying the narration so that guessing the password results in different text being generated than does

cave to open is only acquired when the player character tries this action.¹²

In brief, if it is desirable to have the player character able to automatically do obvious things (such as opening the front door when commanded to leave the house) but not able to do things that are non-obvious and unknown (such as typing in the correct password when the player character doesn't know that password), a model of what the player character knows is extremely useful. The Focalzier World for this character is a basic form of such a model. The alternative to modeling what the player character knows — laboriously coding every “obvious” thing that the PC can do — is time-consuming, is error-prone, and does not scale up well to more complex interactive fiction worlds.

Developing these sorts of automatic reasoning abilities and possible world models of knowledge and action would involve extensive work in the actor class and in the Simulator. Realizing such a system for planning based on enhanced Focalizer Worlds was outside the scope of the first phase of development. Having multiple world models that represent different characters' perceptions and knowledge at least provides a basis for explorations of this in the future, however.

5.2.4 Determining What is Perceived

Characters' perception of the world, or their inability to perceive the world, is dealt with like any other part of the simulation. Perception itself is therefore carried out within the Simulator, and is based on the IF Actual World. When an action is done in the IF Actual World, the IF Actual World instance checks to see which of the Focalizer Worlds should be updated with the action and with new existents, or existents having a new state, because of that action.

For now, a simple model of perception has been implemented. When an actor successfully performs a SENSE on a room or some other existent, that existent (along with all of its visible descendants) is transferred into the actor's focalized world. If an actor perceives an action in any way — which does not require a specific SENSE event, only proximity and visibility — that action is similarly transferred from the IF Actual World to the appropriate Focalizer World.

learning it diegetically and saying it.

12 Of course, actors could make incorrect assumptions about causality and build focalized world models that are not correct. Perhaps someone else, unbeknownst to the player character, is pressing a garage door opener that is the real cause of the cave's opening just as the (SAY “OPEN SESAME”) action completes.

However, the existents that participate in that action are transferred only if they are visible to the actor, so an action may have blank existents in it. For instance, this text is generated in *Lost One* when the player character sees, in a distant location, the trash collector (who has the default prominence, 0.5) picking up the candy wrapper (which has prominence 0.3): “Meanwhile, the trash collector picks up something.”

The candy wrapper is too small to see from far away, but the player character is able to see the trash collector and that he was taking something. The underlying action looks like this in the IF Actual World:

```

ACTION 4 by COLLECTOR: ['TAKE', 'COLLECTOR', 'CANDY-WRAPPER']

45-50-55 [COLLECTOR-ACTING]
/// ACCESS.COLLECTOR.CANDY-WRAPPER
e5: impel (collector candy-wrapper - 10 lift)

45-55-65 [e5]
/// RELATION.CANDY-WRAPPER.IN.PLAZA-SW-INNER
/// ACCESS.COLLECTOR.CANDY-WRAPPER
/// ALLOWED.CANDY-WRAPPER.OF.COLLECTOR
e6: configure (collector candy-wrapper in plaza-sw-inner of collector) -core-
\\ \\ RELATION.CANDY-WRAPPER.OF.COLLECTOR
```

It looks the same in the player character’s Focalizer World. However, in that focalized world, the object identified by the tag CANDY-WRAPPER has been blanked out — its attributes have all been cleared and the player character knows nothing about it except that it is an object. (The tag “CANDY-WRAPPER,” like all such tags, is an identifier for use by the programmer, and never printed except in debugging messages. Tags are arbitrary; this one could have been “OBJ2367” or any other string. A meaningful tag name like this one does not indicate the knowledge that this object is a candy wrapper.) Instead of blanking out an existent entirely, a better model of partial knowledge might allow for the specification of what attributes or qualities can be seen, while others that cannot be perceived are blanked. The current “all or almost nothing” model is intended as a useful first step toward a richer system that would allow these sorts of partial perception.

5.2.5 Zero Focalization

Genette describes the case of *zero focalization*, “the narrative with omniscient narrator [which] Pouillon calls “vision from behind,” and which Todorov symbolizes by the formula *Narrator* >

Character" (Genette 1980:189). The system allows for narrating from a standpoint that is greater than any actor's individual knowledge, that is, for narrating with zero focalization. Specifically, since COSMOS is an actor, COSMOS is a potential focalizer and has its own Focalizer World. (This does not just duplicate the IF Actual World, although its contents completely correspond. As described in subsection 5.2.1, the Focalizer Worlds are different from the IF Actual World in providing the ability to roll back and that they lack the ability to simulate events.) It is possible to narrate while focalizing COSMOS, in which case everything that happens in the IF Actual World has the potential to be narrated. This is really not just the situation of *Narrator > Character* but the situation of *Narrator = Fictional World*, just about as omniscient as is possible, although the narrating of a particular turn is still focused on a range of times corresponding to the current turn, and particular actions and events may be elided, as with any sort of focalization. It would

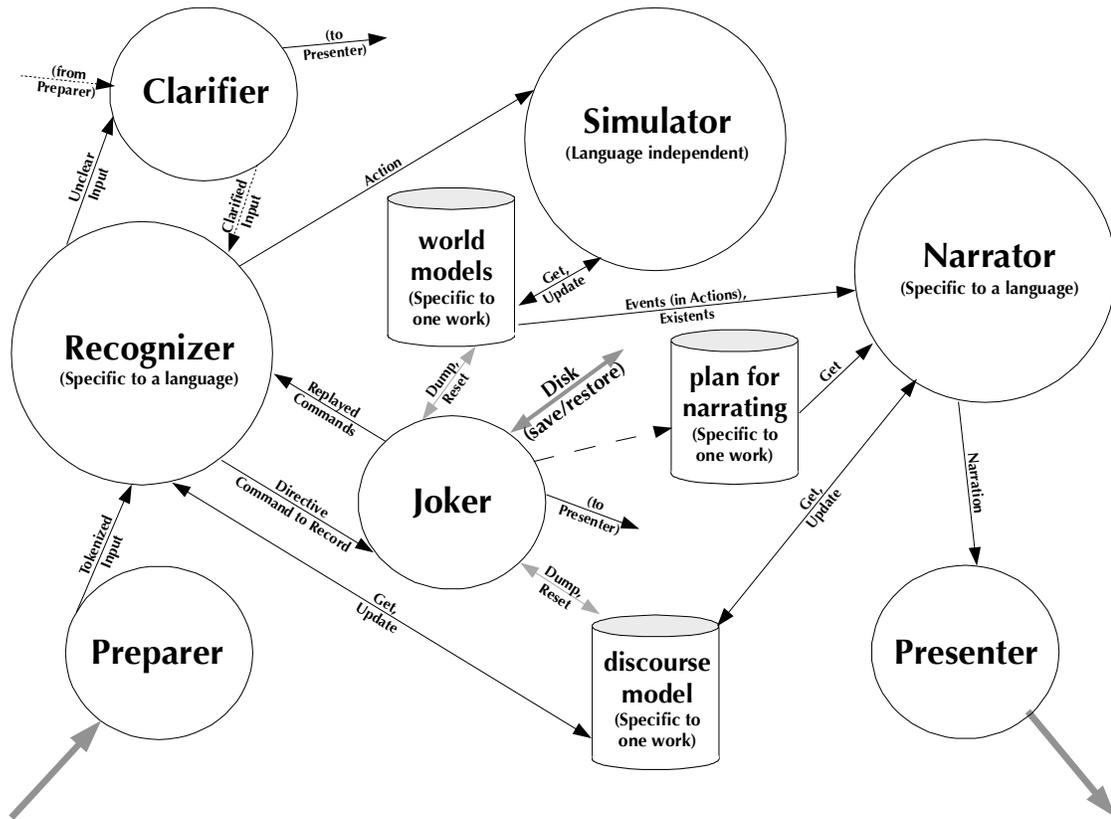


Figure 2: The system architecture, with input on the lower left and output on the lower right. The smaller circles are simpler modules less relevant to the proposed research. Of the major modules, only the Simulator can write to the world models (the IF Actual World and Focalizer Worlds), while only the Recognizer (which parses and understands input) and the Narrator (which does all the preparation of output except for a final formatting step) access the Discourse Model.

be possible to construct other Focalizer Worlds that correspond to sets of actors and represent situations where *Fictional World* > *Narrator* > *Character*, although this has not yet been undertaken.

5.3 *Discourse Model*

The Discourse Model deals with the interactor — what language he or she types and reads, what text has been printed earlier in the session to create the current context, and so on. The interactor's knowledge (a single interactor per session is assumed) is considered to be separate from the knowledge of the player character. The currently implemented Discourse Model is rather vestigial. It includes a given-new list and all the mappings from states to state names, directions to direction names, and so on. The facility for printing numbers in English (for example, as "sixty-nine thousand one hundred and five") is also part of the Discourse Model. The given-new list is the only part of the model that is updated during a session. Clearly an enriched model that also maintains information related to topic, salience, centers, and other important features would be an advantage for interaction. The Discourse Model could even help to inform the Clarifier about the interactor's expertise and his or her difficulties in typing to the system. The minimal model that has been implemented is still useful, however, in testing where the Discourse Model should fit within the architecture.

5.4 *The Modules*

An interactive fiction development system needs to provide separate models for simulation and narration if authors are to be able to achieve narrative variation. The architecture presented here (sketched in figure 2) abstracts the representation of the world from the representation of the discourse and the function of simulation from the function of narration. It encapsulates the different functions of an IF system, as described in chapter 4, in different modules, specifying what each can access and in what way.

5.4.1 *Preparer*

This component is a very simple tokenizer. The module was given a more generic name than "Tokenizer" because this architecture is not meant to be specific to text-based interactive fiction. The Preparer might carry out different functions if there were a different modality being use for

input, such as speech.

Input: A string (textual input from user) read from the input stream.

Output: List of strings (tokens).

5.4.2 *Recognizer*

The Recognizer is a simple parser that uses an easily extended semantic grammar. The Recognizer accepts strings; it outputs either a recognized command and corresponding intended action for the Simulator, a recognized directive and canonical form of this directive for the Joker, or an unrecognized input for the Clarifier.

Input: List of strings (tokens).

Output: A pair, the first element of which is the type of input and the second element of which is the command, directive, or (in the case of an unrecognized input) a list of all rule matches, which may be empty.

5.4.3 *Joker*

This component carries out program-level functions such as saving, restoring, and restarting the game as well as undoing a turn. It does this by accepting directives (non-diegetic inputs), updating the world models if necessary, and producing the appropriate replies (non-diegetic outputs). For development purposes (both for development of the system itself and to assist in the authoring of specific interactive fiction works) the Joker has been provided with the ability to modify the plan for narrating. This is indicated with a dashed arrow in figure 2.

Input: List of strings (directive).

Operation: May update the IF Actual World and Focalizer Worlds, may read from or write to save or replay files.

Output: Pair of strings, the first of which is report text (program-level) and the second of which is reply text (diegetic).

5.4.4 *Clarifier*

Currently, this component does little more than producing an “I don’t understand...” reply for anything not parsed. It also indicates if an input is ambiguous and seems to refer to more than

one command. In a released system, the Clarifier would have to be extended to help clarify ambiguous or unclear inputs, probably interactively and with reference to the Discourse Model.

Input: List of matching rules (possibly empty).

Output: A clarifying message, which is printed using the Presenter.

5.4.5 Simulator

The Simulator manages the entire simulation, including the states of existents (a door is open or closed), the configuration of existents in the world (the player character is in the kitchen), and events that may change these. The Simulator was described in detail in the first part of this chapter, alongside the world models that it works upon. The Simulator is meant to be completely independent of the natural language that is used for input and output.

Input: List of strings (command), current commanded actor.

Operation: Updates the IF Actual World and Focalizer Worlds.

Output: Indices of the actions that were just done during the turn.

5.4.6 Narrator

Without making any changes in the simulated world, the narrator produces discourse-level expression for the interactor to read. The narrator uses a standard three-level pipelined architecture for text generation.

Reply Planner. Content selection and ordering is done here, according to the current narrative specification and based on events and existents.

Microplanner. This level incorporates referring expression generation and along with finer forms of aggregation, determining when phrases will be combined into a single sentence. Speed is considered here.

Surface Realizer. Detailed sentence plans are realized as language at this stage and are converted into orthographic strings for output. If the representation used by the Microplanner was sufficiently abstract, generation could be done at this stage in languages other than English. The current representation of sentences embeds English strings and is otherwise language-specific, so the Microplanner would have to be modified as well to create a Narrator for a different language.

The Narrator is central to the research project discussed here, so the workings of the Narrator

are covered in much greater detail later, in chapters 6, 7, and 8 . The preference for different narrative variations is formalized into a plan for narrating, which is an input to the Narrator.

Input: Indices of the actions that were just done, plan for narrating, Focalizer Worlds.

Operation: Updates the Discourse Model.

Output: String (with diegetic or reply text) to be printed using the Presenter.

5.4.7 Presenter

The final processing of text for display is done in this component. For now, this simply involves formatting it for a terminal window of a particular width. Other presenter modules could be developed to output in HTML or other formats, including non-textual ones.

Input: String.

Output: Formatted string, written to the output stream.

5.5 What an Author Writes for a Particular Work

The author of an IF work will generally define existents (and their arrangement in the tree) and will specify the plan for narrating. Some existents are needed so that events (which require existents as dependents) can transpire, so that there will be something to interact with and

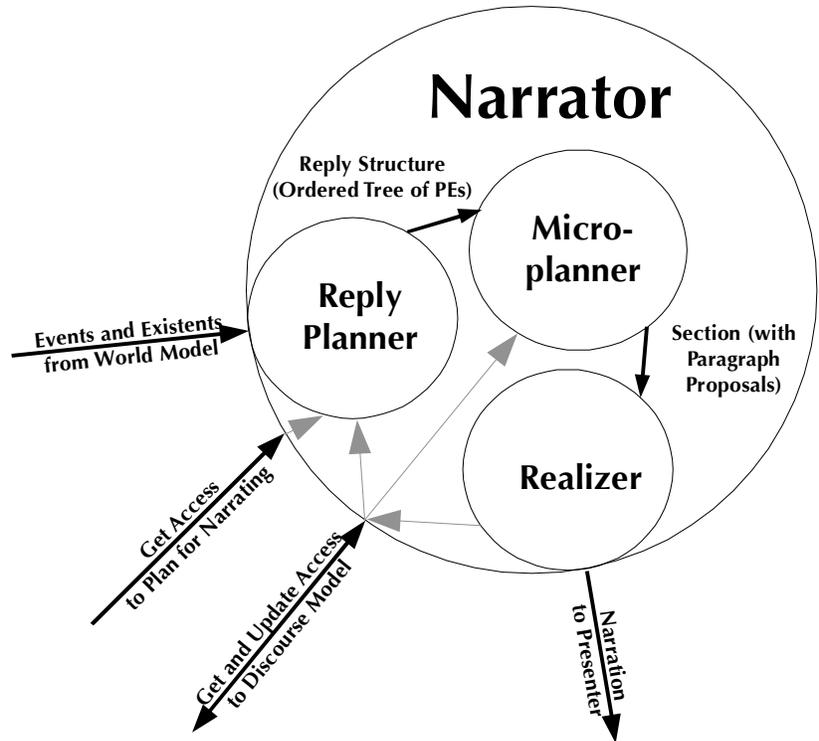


Figure 3: The architecture of the Narrator module. A standard three-stage pipeline is used. The Reply Planner first converts events, existents and a plan for narrating into a reply structure, an ordered tree of proposed expressions that represents what is to be narrated and in what order. Next, the Microplanner accepts this and uses discourse information to do aggregation and referring expression generation, converting the RS into a list a paragraph proposals. The Realizer renders each sentence in these as text and outputs them, updating the Discourse Model as it does this.

something to narrate. There should at least be one room (so that things and actors can be contained in it) and one actor (since there needs to be a focalized and commanded character). Defaults can supply most of the plan for narrating, although a typical nn game will probably want to specify its own plan for narrating to take advantage of nn's capabilities. The following code defines a complete and very simple game in nn, to illustrate how one actor, one room, and the essential parts of the plan for narrating are defined:

```

"""One Room

A simple (and very uninteresting) IF piece"""
__author__ = 'Nick Montfort <nickm@nickm.com>'
__version__ = '0.1'

from Models.discourse import Message, command_lists, compass,
main_directions, directives, relation_names, text_to_relation,
state_names
from Models.world import Actor, Thing, Room
from Simulator.events import *
from Narrator.microplanner import Paragraph
from copy import deepcopy

frontmatter = (
    'One Room',
    'An extremely simple piece of IF',
    'By Nick Montfort',
    ''
)

message = Message()

existents = [

    Actor( 'PERSON',
        called=('the', 'person'),
        parent=('IN', 'CAVERN'),
        adjs=['my'],
        nouns=['self', 'myself', 'me'],
        appearance='a typically nondescript character',
        state={
            'GENDER': 'MALE'
        }
    ),

    Room( 'CAVERN',
        called=('a', 'cavern'),
        adjs=['enclosing'],
        nouns=['cavern', 'cave'],
        appearance='an enclosing cavern',
        exits={},
    )

]

pfn = {}
pfn['focalized'] = 'PERSON'

```

```
pfn['commanded'] = 'PERSON'  
  
person_existents = deepcopy(existents)  
  
focalizers = [  
    ('PERSON', person_existents, {})  
]
```

A typical IF work in nn will also have methods on the COSMOS existent that state how to update the plan for narrating or if another plan for narrating should be substituted under certain circumstances. Beyond this, a work may contain strings-with-slots to customize the narration of events and even the output of reports, may add to the Recognizer's grammar for mapping input texts to intended actions, may add to the mappings intended actions and actions, and may extend the basic system in many sorts of ways.

6 Reply Planning and Variations in Order

A detailed and systematic treatment of how the telling of a narrative can be considered apart from the existents and events represented in it is found in Gérard Genette's *Narrative Discourse: An Essay in Method* (published in French in *Figures III* in 1972; English translation, 1980). In this discussion, Genette covers variations in temporal relationships or *tense*: how events can be narrated in a sequence that differs from their chronological sequence (order), how the telling can move more rapidly or more slowly and narrate events for shorter or longer periods of time (speed¹³), and how events can be narrated once each, one for several occurrences, or several times for each occurrence (frequency). Genette continues to consider the narrative analog of what is called in grammar *mood*, which includes the perspective from which a narrative is told (focalization). As part of the final category, *voice*, the time of narrating and its function in narrative is discussed. The time of narrating deals with whether the narrating predates, postdates, or occurs simultaneously with the events of the narrative. While the basic concepts have been extended, revised, and further discussed in various ways (including by Genette himself, 1988), the general outline of relationships was groundbreaking and has proven essential to decades of development in narratology.

The concepts introduced in *Narrative Discourse* and developed by other scholars in the years following are often quite precise. Since they pertain to the form and function of narrative and have developed from formalist and structuralist ideas, they can be considered formal in this sense. However, although the clarity with which narratological concepts are expressed compares favorably to other sort of humanistic theories and critical approaches, they are not expressed in a way that is mathematically or computationally formal. To implement different types of narrative variation in an interactive fiction system, additional refinement and formalization of narratological concepts is necessary so that they have a symbolic expression that allows them to be worked on computationally. A group of researchers who have undertaken such work explain the challenges:

13 The term "duration" (*durée*) is used in *Narrative Discourse*, but in *Narrative Discourse Revisited* (Genette 1988) "speed" (*vitesse*) is offered by Genette as a more appropriate term that captures the essential relationship between the pace of narrating and the pace of the narrated events. For this reason, "speed" is used here throughout.

In general, most Humanities models of narrative contain formalizations only at very abstract levels, if at all. By formalizations, we mean here a representation in some logic language (e.g., predicate calculus) or other structured representation, including tables, graphs, etc. Indeed, most works dealing with narrative and not going back directly to the structuralist tradition are composed in “plain prose”. Especially, there seems to be a tendency to apply formal notions to the abstract *histoire* level only. Phenomena at *discours* level that apply to the *structure* of discourse (e.g., discourse relations) are sometimes formalized in linguistics and are usually described in words only — sometimes accompanied by tables — by literary scholars (Genette, 1980). Where models are based on the *discours* (text) layer of a narrative or include it, genuine Humanities models usually lack formality, though their descriptions might offer a variety of authentic examples. (Gervás *et al.* 2006)

As Michael Young has explained, “A central challenge of any computational approach that seeks to operationalize concepts from narrative theory is to determine appropriate methods to translate concepts derived from analysis into concrete, formal models capable of being put to use in the creation of an *interactive* virtual environment” (Young 2007, emphasis added). The additional challenges that arise in an interactive system, as opposed to a “batch” system such as a traditional story generator, are those relating to the interface, the way the underlying simulation is affected by the interactor, and the additional interactive (rather than exclusively narrative) functions of the textual output.

It can be seen as a deficiency of computers that they demand such precise, rigid definitions. There can be benefits for the study of narrative, however, when systems are formalized for the computer. Marie-Laure Ryan has described some of these in discussing a computational system for the analysis of texts:

As Jan Christoph Meister writes, ‘most narrative theories are developed and tested on the basis of a single illustrative example which later provides an ideal breeding ground for sweeping generalization, dogmatic theories, and vitriolic polemics’ (p.201f.). The computer provides a potent antidote against this kind of theoretical laxity. It forgives neither imprecise definitions nor inflated claims, and it possesses the uncanny ability to confront us with situations that we had not anticipated. Since computer programs are designed to run on a variety of inputs, they offer a reliable test of the generality of theoretical concepts. An idea based on a single example is likely to make the program crash, or to make it produce garbage, on any text other than the one that it originally described. (Ryan 2003:x-xi)

Meister’s quoted statement may seem extreme. At any rate, narrative theories are not unique in growing from the seeds of single examples. Meister continues by noting that “the remarkably persistent popularity of Genette’s theories is due not least to the fact that he keeps his concepts and categories general enough to be suitable for further use, despite demonstrating them with just a single text” (Meister 2003:203), making an exception for the main narrative theory that informs this project. Nevertheless, the generality that a computer implementation requires does

offer something new to the development and testing of narrative theories, and can help to refine even Genette's general and already useful concepts. This is true not only for analytical systems such as the one Meister developed; it can also apply to generative systems such as mn, particularly when they are developed to narrate in different ways and to have the capability to use many different simulated worlds as the basis for the telling.

To achieve the sort of formalization necessary for implementation in a computer system, the inputs to the Reply Planner, the internal operations of the Reply Planner, and the output from the Reply Planner to the Microplanner are defined in detail here, alongside the concepts of narrative discourse that these operations are based on. The lower-level work of the Microplanner, and the lowest-level work of the Realizer, are discussed in the next chapter.

The input to the Reply Planner consists of a Focalizer World, a set of indices to actions indicating what in that Focalizer World has transpired in the most recent turn, and a plan for narrating. The Focalizer World, of course, includes existents, along with the capability to roll back to a point in the past and see what existents were like at that point. It also includes events with causal connections between them and with temporal information about each.

The Reply Planner uses this input to build an ordered tree called a reply structure (RS), with proposed expressions (PEs) as leaves. The PEs indicate how the narration of an event, the description of something in the content plane, or the creation of some non-diegetic text is to be done. For instance, a standard transformation to produce the sort of narration often used in existing IF would result in chronologically-ordered nodes being placed in an RS of depth 1 in which each event is marked with the default speed (.5). For every PE that is in the output RS, some text will be generated — all content selection is done in the Reply Planner, and nothing selected at that point may be elided at a later stage. The details of what text is generated from PEs are handled by the Microplanner and the Realizer. To whatever extent this is possible, the Microplanner will output a longer abstract paragraph or sentence representation when the speed is slower (and length should be greater) and a shorter one when the speed is faster. In the last stage, the Realizer, the abstract representations provided by the Microplanner are converted into strings of English for formatting and output.

6.1 Narrative Tense

Varying the representation of events in the content plane, or how “story” is expressed, involves being able to output different signs of the narrated for the same underlying events. About three decades ago, it was asserted that “Gérard Genette’s elegant analysis of the time-relations between story- and discourse-time must form the basis of any current discussion. Genette distinguishes three categories of relations: those of order (*ordre*), duration (*durée*), and frequency (*fréquence*)” (Chatman 1978:63). Genette’s foundation for the discussion of story-time and discourse-time has remained very helpful to theorists of narrative, has been used in understanding the workings of time in cybertexts (Eskelinen 1998), and has also been at least an inspiration in even the “anti-narrativist” consideration of time in video games (Juul 2001). While there are some difficulties with analyzing narratives while assuming an underlying chronological order for events in the content plane (see Adams 1999 and Herman:211-261), for the task of generating narrative based on a known, underlying simulation, these problems do not truly manifest themselves.¹⁴ For these reasons, Genette’s treatment of time serves as the basis for the generation of narrative variation.

The current project concerns itself only with generating a narrative (sometimes a very short one) given a representation of events and existents in the simulated world and information about which of these have transpired since the last input from the interactor. For this reason, the discussion that follows restricts itself to the reasonably well-established framework of *narrative time*, although there is also the important question of how this time couples with a larger framework of *interactive time*, which is essentially the same as what Juul has called *game time* (Juul 2001). For instance, if the first ten turns of an interactive fiction are played in a simulated 1980, the second ten turns in 1970, and the third ten turns in 1960 — all in the same simulated world — many issues arise that are not covered in the upcoming discussion of retrograde narration, because it is not just the narration but also the opportunity to interact and command a character

¹⁴ The argument against chronology may show that the deterministic, chronologically ordered content plane as it exists in this interactive fiction system is inappropriate and that, for instance, it would be better to have a store of events from the distant past which are not organized chronologically but which can nevertheless be inserted in analepses. Such events could be accommodated by a weaker version of the current system, however, in which timestamps are not required or can be approximate. Removing the requirement that all simulated events are associated with a definite point in time would not mean scrapping the current simulation and model of events; it would just require certain modifications. However, it may be that simulation requires a definite chronology not present in narrative, so that the generation of a narrative from a simulation will always necessarily be based on events with a definite chronology.

that moves backwards in time. For now, the problem being considered is that of narrating with game time always moving forward, as is conventional in interactive fiction and video gaming. The matter of the relationship between interactive time and narrative time is very important from a design and aesthetic perspective, and will be important in future work done from a new media and creative perspective. The current project's focus on narrative time, however, is meant to serve for a wide range of interactive fiction designs and to supply at least an adequate starting point for work on other sorts of interactive designs.

6.2 Pure Order of Events in the Narrative

Order is a feature of all narratives, although it may not be a very significant one in narratives that are as simple as possible (narrating only one event) and in ones that provide the least possible information about time. The telling of a single event can be considered a narrative (Prince 2003:58) and one might think of such a narrative as a chronicle, but it might just as easily be considered a degenerate case of any possible order. Additionally, if there is nothing to indicate that events are chronologically related at all, nothing can be said about how the order in which they are narrated relates to their order in the content plane. But for other sorts of narratives, order is meaningful and is usually very important. The claim that “[o]rdering of events in time is one of the most fundamental characteristics of any story” (Prince 1973:23) has hardly proved controversial. It has been validated by the efforts of literary writers and those who create narratives in other media to resist this principle and to confound the attempts of readers to make sense of a narrative's chronology. From one perspective, “the straining against the ‘tyranny of time’ throughout the ages, in modernism, for example, only reaffirms and redefines the tyrant's power with each abortive rebellion” (Sternberg 1990:901). A milder comment along these lines would be that there would be no compulsion to continually subvert or play with chronology if it were not somehow a particularly powerful organizing principle.

Given several events with a known chronological relationship, there are many ways these events can be ordered in a narrative. There is also a repertoire of conventional types of orderings which have been observed in literary, conversational, and other narratives. For n events, varying only the order, and keeping events expressed (rather than elided) and the frequency singular (so that each event is expressed once), and leaving aside for now the issue of how tense, time words,

and other markers might be used or omitted to make temporal relationships more or less apparent, there are $n!$ variations in order possible. These can be grouped into several higher-level categories of interest. Here, the categorization is based on Genette's discussion of order in *Narrative Discourse*. It was noted almost three decades after the French publication this discussion that "[m]ost subsequent narratologists have adopted Genette's analysis of anachronies, either explicitly or implicitly, and none have extended it" (J-K Adams 1999:114). The variations in order that Genette discusses are:

Chronicle. Events are narrated in the order in which they occur. Simultaneous events can be narrated in any order, relative to each other, in a chronicle. This ordering, as natural as it may seem, has been called "more hypothetical than real" (Genette 1980:36). Nevertheless, stretches of chronological narrative can be seen in many works of historiography, including the Bible and Thucydides's *The Peloponnesian War* (see Sternberg 1990:921-922). Conversational stories also are often told as chronicles, for example, "The Baddest Girl in the Neighborhood" and "Eating on the New York Thruway" (Polanyi 1985:21-22, 36-37). There is a great deal that has been written and that still remains to be said about the chronicle as a type of narrative, but at least two extreme views can be rejected. First, chronicle is not always the default arrangement for events in a narrative. It may be that causality,¹⁵ the categories into which the events fall, the associations they occasion to the narrator, or many other factors are as important as the order in which they transpired and that these other factors lead to a different ordering. Second, it is not reasonable to think that the chronicle is always absent from interesting or even avant-garde narratives. For instance, in Harry Mathews's first three novels, *The Conversions*, *Tlooth*, and *The Sinking of the Odradek Stadium*, and in his most recent novel, *My Life In CIA: A Chronicle Of 1973*, chronological order predominates, although these novels (or "chronicles," as Mathews originally called them) are far from conventional.

Modern interactive fiction, while it is based on an essentially chronological simulation of events, nevertheless often deviates from the chronicle. The transcripts from interaction with *Varicella* and *Bronze* (appendices A and B) both make extensive use of analepsis. These can be

¹⁵ The importance of causality is almost certainly understated by any theory of the arrangement of events that is based on time and temporal relations. nn does represent causality, but order is dealt with from a mainly temporal perspective, which may be problematic. However, temporal relations do need to be tracked and expressed in syntax no matter what principle for ordering them is used. What is learned about the arrangement of events here would still be relevant to a system that looked mainly at causality rather than time.

seen, for instance, in the first two replies in *Varicella*, which narrate Terzio's assassination and Primo Varicella's first encounter with Charlotte, and in the first reply in *Bronze*, beginning with "You caught him staring at you once..." Of course, the events that are narrated during these analepses are not simulated; their narration is hand-written into strings of text. This can, certainly, result in a powerful interactive experience. The system described here attempts to allow such sorts of recounting, and other sorts of narrating, to be generated in much more flexible and general ways, and to use not only fixed, non-simulated events from the past but the full store of events that have happened.

Retrograde. Events or temporal sections are narrated in the reverse of the order in which they occur. Examples with different levels of granularity include the basic structure of opening lines of the *Iliad* (Genette 1980:36-37), Charles Baxter's *First Light*, the main sequence in Christopher Nolan's *Memento*, Martin Amis's *Time's Arrow*, and Alexander Masters's biography *Stuart: A Life Backwards*. Outside of literature, film, and fictional narrative, retrograde narration does make some appearances: The typical résumé has milestone events grouped by category (education, work experience) within which the events are presented in retrograde order.

Zigzag. Events or temporal sections from period 1 (the "now") are interleaved with those from period 2 (the "once") as they are narrated in order: One example is a passage from Marcel Proust's *Jean Santeuil* (see Genette 1980:37-38). There is a correspondence between the sections and a comparison of a sequence of events in the past to ones in the "now" of the narrative. It is possible to either narrate the "now" first, followed by the "once," or to do the opposite. As with syllepsis, which is discussed further on in this section, the events that are paired must be similar in some way: the character entered a crowded room then, she enters a crowded room now; someone spoke to her then; someone speaks to her now. A strict zigzag has an even number of events or temporal sequences, since there is a series of alternations between "once" and "now."

Analepsis. An event or temporal sequence is narrated that is from the past, relative to what is being narrated. An analepsis has reach (the distance backward in time) and extent (the duration of the past event or temporal sequence). There are many different sorts of analepsis even when only order is considered (external, internal, partial, complete) but all of these can be characterized in terms of reach and extent.

Prolepsis. An event or temporal sequence is narrated that is from the future, relative to what

is being narrated. A prolepsis also has reach (the distance forward in time) and extent (the duration of the future event or temporal sequence).

Syllepsis. The order of events is based on some grouping that is not chronological. For instance, in recounting a stereotypical adventure, all the encounters with monsters might be narrated, all the arrivals in new places, and then all the acquisitions of treasures. Examples from off the computer include the spatially-organized *La Vie mode d'emploi* by Georges Perec and Stendhal's *Mémoires d'un touriste*. Another example of spatial syllepsis is seen in how "Proust sometimes presents a whole series of events, all of which have occurred in the same place" (Bal 1997:99). The tree representing zigzag has the same form as the one representing syllepsis; Zigzag can be seen as a restricted form of syllepsis with the additional constraints that in each category, one event from the "once" and one from the "now" are narrated and that the overall progress is chronological.

Achrony. In this case, the relationship between the order in which events are narrated and the order in which they occur is impossible to establish, or is exceedingly difficult to establish and seems arbitrary because the occurrences of events are not "dated" and their order cannot be known. Alain Robbe-Grillet's *La jalousie* provides an example of this type of order. The full effect of achrony is usually achieved by omitting discourse markers and the typical signs of chronological relationship provided by tense and aspect, and this type of disassociation of an event from time is what Genette refers to as causing an achrony (Genette 1980:83), but these are concerns that will be taken up after the basic issue of ordering events.

6.3 *Pure Ordering of Events in the Narrating*

The order in which events are represented is something that is seen in the narrative (what Genette calls *narration*), but the process that ordered these events is in the narrating (*récit*), "the producing narrative action" (Genette 1980:27). To have a system that can generate narratives with events in several different orders, it is necessary to have algorithms for ordering events. The processes of ordering have to be discussed, not just the outcomes of such processes.

It seems meaningless to discuss the order of events in an narrative that represents only one event, but it is meaningful to discuss whether such a narrative is consistent with particular ordering processes. For instance, consider two narrators, one that has always narrated events in a

retrograde manner and one that always narrates chronologically. If both narrators produce a very brief narrative, representing a single event, we can ask whether each of these most recent narratives are told in a consistent way. In both cases, since a single event cannot be narrated in the “wrong” way in either a chronological or retrograde scheme, their narrations will be consistent with their previous behavior as narrators. If there were a narrator who always included an analepsis in narratives, on the other hand — always narrated, at some point, one event followed by some other event that happened earlier — no single-event narrative could be consistent with this characteristic. While it is not meaningful to say that a single-event narrative has a particular order, it is possible to determine that some processes for ordering can produce it while others cannot.

Another example can be seen in random narrating. Four events that are ordered 1 2 3 4 (with 1 being the earliest, 2 next, and so on) are, of course, chronologically ordered. But 1 2 3 4 is also a valid choice for a random narrator (assuming that there is some probability mass on 1 2 3 4). For a narrator that selects sequences uniformly at random, this is as probable as any other sequence. So, 1 2 3 4 is consistent with both processes. It will always be produced by the former, of course, while it will never be produced by certain other processes, such as a process of retrograde narrating.

With this in mind, it is possible to define particular processes whose characteristic output falls into the categories described by Genette:

Chronicle. Sort a set of events into chronological order. “Chronicle” will not always specify a unique order, even when a timestamp is required for each event, because the set of events may include some that are simultaneous.

Retrograde. Sort a set of events into reverse chronological order. Again, because of simultaneity, this may not be enough to specify a unique order.

Zigzag. The process of zigzagging between two related chronological sequences of events requires that two such sequences are designated. Beyond this, a rule for moving between sequences is needed. This could be as simple as “narrate a single event before switching,” or it could involve specifying that all the events in a single physical location are narrated in the “now,” then the corresponding events in the “then,” and then similarly with the next physical location. Whatever the case, the process of ordering simply involves applying this rule to the two

sequences specified.

Analepsis. This indicates an anachronism inserted into a main sequence that is presumably chronological. For this process to work, both a main sequence and the point of insertion of the analepsis need to be designated. From the standpoint of the analysis of narrative, is it sensible to discuss the reach and extent, but when generating an analepsis, the difference in time and the duration of the analepsis are not the most useful things to specify. Instead, it is more useful for a rule to specify what should be included in the analepsis based on features of events. For instance, “select the most salient event from the first time the focalizer encountered this character” or “select the most salient events that the focalizer has seen happen in this room in the past, up to three of them.” Of course, to make the latter rule useful, a rule for determining the salience of events must also be precisely specified. Given the main sequence, the point of insertion, and a fully specified rule for selecting events from the past, the process of ordering events so as to include an analepsis is straightforward.

Prolepsis. To insert a prolepsis, the same three inputs are needed: a main sequence, a point of insertion, and a rule for selecting events from the future. When some newly-simulated events are being narrated for the first time, there will not be a supply of simulated events waiting in the future. However, there are still circumstances under which a prolepsis can occur. An IF author can prepare “inevitable” events with future timestamps, representing things like the sun going down, a storm beginning, an election happening, nuclear missiles arriving to destroy the city, and so on. Also, there will generally be plenty of times in which the main sequence of events being recounted is from the past, perhaps because a

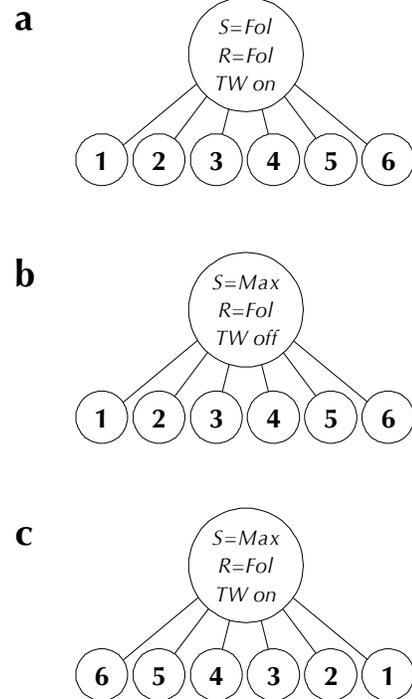


Figure 4: Three simple reply structures: (a) represents a present-tense chronicle with time words used; (b) represents a similarly-ordered chronicle, but with no time words and told in the past tense; (c) represents a retrograde narration in the past tense.

character is recounting it in direct or indirect discourse or because the top-level narrator has chosen to recount it. In such a case, there will be plenty of future events to include proleptically.

Syllepsis. Beyond the original set of events, only a sequence of categories seems essential for specifying sylleptic narrating. For instance, such a sequence might have these three categories of events in it: “the adventurer entering a new area,” “the adventurer defeating a monster,” and “the adventurer acquiring a treasure.” If all events are in exactly one category (the categories partition the set of events), the categorization will be unique. The narrator can move through each of the categories in order and, within each category, can represent each of the events chronologically. There is no reason to restrict a sylleptic narration to chronological order within categories, though. (As previously mentioned, even the prosaic résumé can be viewed as a sylleptic narrative that uses a retrograde order within its categories.) So, it is most flexible to allow any principle for ordering based on time alone (chronicle, retrograde, achrony) to be specified for ordering the narrative within categories.

Achrony. Ordering events at random seems the most suitable way to produce the type of order needed for achrony. As discussed earlier, there is always the chance that choosing an order uniformly at random will produce an order such as 1 2 3 4, which is probably not convoluted enough. To specify a narrator truly capable of “privileging confusion,” something Janet Murray has accused postmodern writers of doing (Murray 1997:133), it would be enough to use a distribution over sequences that has reduced or no probability mass on obvious, non-confusing sequences such as 1 2 3 4, so that more unusual ones would be preferred. But a narrator that orders events uniformly at random is probably confusing enough for all practical purposes.

6.4 Time and Grammatical Tense in Ordering Events

So far the discussion has only covered how events can be rearranged from a chronological sequence into a narrative one. But reordering is not best seen as simply producing a sequence. An analepsis, for instance, is not well represented by the sequence 3 4 5 1 2 6 7. The sequence of events that is in the past, relative to the main sequence — the 1 2, in this case — is embedded in a way that cannot be seen in this simple representation. When the main sequence is being told in the present tense, the 1 2 will almost certainly be told in the past. If the main sequence is already being told in the past tense, there will almost certainly be some cue that 1 2 occurs at a much

earlier time: a phrase such as “before that,” an explicit reference to the earlier date, some statement about habitual occurrences in the past, or a statement in the perfect leading into the analepsis. Even without attempting to generate all of these sorts of transitions, or many of them, there is clearly a need to designate more about the order of events than a simple sequence does. The representation should not force the tense of the analepsis to be different, but it should allow for this difference. It should also integrate the times at which events occurred into the decision about tense. Simply associating an arbitrary tense with the main sequence and another arbitrary tense with the analepsis would not accomplish this. The grammatical tense should be a result of the position of the simulated events in time — along with other essential parameters.

Genette noted that the nature of Western languages means that the temporal position of the narrating vis-à-vis the narrated has a special status:

I can very well tell a story without specifying the place where it happens, and whether this place is more or less distant than the place where I am telling it; nevertheless, it is almost impossible for me not to locate the story in time with respect to my narrating act, since I must necessarily tell my story in present, past, or future tense. (Genette 1980:215)

These tenses lead to the “three major possibilities” for the temporal position of the narrating relative to the narrated: posterior, anterior, and simultaneous narration (Prince 1982:27). While Genette deals with this in the category *voice* rather than in his discussion of *order*, from the standpoint of generating narrative and determining the grammatical tense to use, the temporal relationship of the narrator to events is as important as the temporal relationship of events to one another. They must be dealt with jointly.

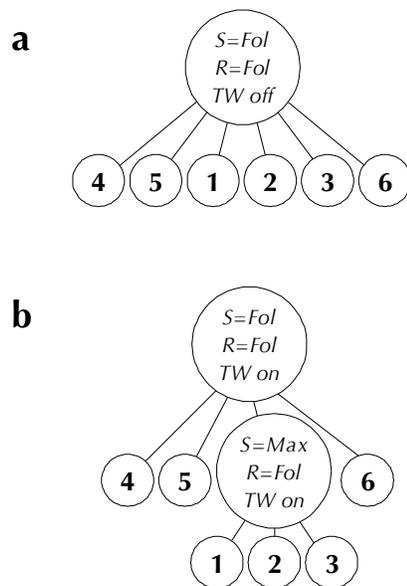


Figure 5: Two reply structures: (a) represents achrony, with events represented in the present tense; (b) generates a present-tense narration of events 4, 5, and 6 with a past-tense analepsis (including 1, 2 and 3) embedded within it. Note that the sequence of events is the same in both cases; the ordered tree representation allows for the same sequence to be generated coherently in the latter case and confusingly in the former, with both corresponding to different types of order identified by Genette.

The discussion that follows explains how the tense of a PE is necessary for realization; how this tense can be determined from three points in time assigned to the PE that are called E, R, and S; and how these points can be defined using general rules (specifically, FOLLOW, MAX, MIN, N, and HOLD) that reside in the reply structure on internal nodes.

To allow the timestamps of the events and the temporal position of the narrator to participate in the determination of tense, an ordered tree of internal nodes and proposed expressions (PEs) of events is used. This complete representation of order is the Reply Structure, a specification of what content to include, what order to include it in, and how to embed sequences. To realize a particular PE, there must be enough information about it to fully specify its syntax; in particular, this means that the system must be able to determine the tense of the sentence or sentences that are to be generated.

To determine the tense, we look to a theory of how knowing three points in time — speech time (S), reference time (R), and event time (E) — allows grammatical tense to be determined (Reichenbach 1947:287-298). Three times are identified as necessary by Reichenbach because in a sentence such as “Peter had gone,” there are three relevant points of time that are needed to explain the tense: the time at which the sentence is spoken (S, the time of speech); the time at which Peter left (E, event time), and another time which is being referred to, in this case after the event time and before the time of speech, by saying “had gone” rather than something else, such as “went” or “was going.” This last time is R, the time of reference. Specifically, “The position of R relative to S [corresponds to] ‘past’, ‘present’, and ‘future’. The position of E relative to R ... ‘anterior’, ‘simple’, and ‘posterior’” (Reichenbach 1947:297).

Absent any context and any information about the temporal position of the narrator, a particular proposed expression of an event will still have the necessary information about when the event occurred, corresponding in Reichenbach’s system to E. The Narrator would not be very helpful if it were necessary for the author to write code to determine every value of R and S for every PE. Instead, the Reply Planner uses the topology of the reply structure to assign R and S in a systematic way across each embedded sequence. Each embedded sequence has a parent, an internal node. On each internal node, a rule for determining the R and S values for children is provided. For each of R and S, the rule can be:

FOLLOW — Set the value of R or S to E. Reference time or speech time “follows” the events.

MAX — R or S are assigned to have the maximum value (*Max*), always greater than E; if both are set to *Max*, $R=S$.

MIN — R or S are assigned to have the minimum value (*Min*), so that this value is always less than E. If R and S are both set to *Min*, $R=S$.

N — Any integer value.

HOLD — Use the current rule for R or S as determined by a higher-level internal node, given this point in the parent's embedded sequence.

For example, consider a reply structure that consists of just a root (a single internal node) with one level of n PEs beneath it, their event times indicated by $E_1 \dots E_n$. Setting speech time to MAX and reference time to FOLLOW in this internal node will assign $S_1 \leftarrow \text{Max}$, $S_2 \leftarrow \text{Max} \dots S_n \leftarrow \text{Max}$, and $R_1 \leftarrow E_1$, $R_2 \leftarrow E_2 \dots R_n \leftarrow E_n$, so that throughout the sequence, $E = R < S$. This corresponds to simple past-tense narration for the entire reply.

Setting both speech and reference time to FOLLOW will similarly mean that $S = R = E$ everywhere, producing simple present-tense narration. When narrating events and moving back in time to narrate previous events, in an analepsis, the speech time can be held at the current point in the main sequence using HOLD while the reference time is set to FOLLOW, so $E = R < S$ for past-tense narration throughout the embedded sequence. Finally, in narrating some events that happened between time 500 and time 600, R can be set to 600 and S to MAX to generate representations of the events in the past perfect. A narrative sequence can be generated in any of Reichenbach's nine fundamental forms ($\langle \text{anterior, simple, posterior} \rangle \otimes \langle \text{past, present, future} \rangle$) by specifying S and R in such ways.

Each PE has its R and S values set in the Reply Planner using the rule from its parent. Once all the PEs have been defined with specific values for E, R, and S, all the necessary information is in place for the next stage of the Narrator to compute the tense using Reichenbach's formulas. The R and S rules on the internal nodes are general (they do not require that all values of R and S be computed by author-written code) but also flexible (they do not demand, for instance, that every analepsis is told in a different tense from the sequence in which it is embedded). Different types of order are specified directly using these rules, which are associated with internal nodes in the Reply Planner.

The internal nodes also carry some additional information. They have a time words (TW)

setting to determine to what extent expressions such as “before that,” “then,” and “meanwhile” will be used to link representations of events directly beneath them. Using information stored on the internal nodes and in commentary nodes at the beginning and end of sequences of proposed expressions of events, more complex effects could be achieved. Framing statements from the narrator such as “I remember” and “anyway” could be added around an analepsis, and a preface such as “I foresee” could be inserted before a prolepsis.

With this model established, it is possible to precisely define what distinguishes achrony from a random reordering of events that is related in a coherent way. In achrony, speech and reference time are equal and either remain greater than, remain less than, or follow the event times for the entire interval, so that everything is narrated in the same tense, and helpful time words such as “then” and “before that” are suppressed. If an event moved to the past is instead treated as an analepsis with the appropriate shift in tense, or if time words were generated to indicate how events are related, the result will not be as disorienting.

Three example narratives generated from these sorts of reply structures, all relating the same set of events, follow. The first two are generated from reply structures similar to the ones shown in figure 4a and 4b, although there are more events and some of the events are simultaneous. The second one is generated from a reply structure like the one in figure 5a, although the particular shuffling is different.

6.4.1 Chronological, Simultaneous, with Time Words

You look at the center of the plaza.

Your senses are humming as you view the broad, circular, encircling Plaza of the Americas. The morning has concluded. It is midday now.

From here, you see a statue and a flaneur to the north, a fountain to the east, a trash collector to the southwest, a ball and a boy to the northeast, a mime and an obelisk to the south, and some punk and a tree to the west.

Then, the punk kicks the tree.

Meanwhile, the flaneur conveys himself to the northern area.

Then, the boy throws the ball.

Then, the flaneur looks at the northern area.

Then, the mime waves.

Meanwhile, the trash collector takes something.

Then, the ball falls to the ground.

6.4.2 *Chronological, Subsequent, no Time Words*

You looked at the center of the plaza.

Your senses were humming as you viewed the broad, circular, encircling Plaza of the Americas. The morning had concluded. It was midday then.

From there, you saw a statue and a flaneur to the north, a fountain to the east, a trash collector to the southwest, a ball and a boy to the northeast, a mime and an obelisk to the south, and some punk and a tree to the west.

The punk kicked the tree.

The flaneur went to the northern area.

The boy threw the ball.

The flaneur looked at the northern area.

The mime waved.

The trash collector picked up something.

The ball fell to the ground.

6.4.3 *Achrony, Simultaneous*

Some punk kicks a tree.

A trash collector picks up something.

A mime waves.

A ball falls to the ground.

You look at the center of the plaza.

Your senses are humming as you view the broad, circular, encircling Plaza of the Americas. The morning has concluded. It is midday now.

From here, you see a statue and a flaneur to the north, a fountain to the east, the trash collector to the southwest, the ball and a boy to the northeast, the mime and an obelisk to the south, and the punk and the tree to the west.

The boy throws the ball.

The flaneur looks at the northern area.

The flaneur conveys himself to the northern area.

7 Reply Planning and Variations in Speed, Frequency, Mood, and Voice

7.1 *Speed*

For n events which can each be narrated at d different durations, nd variations in speed are possible in the overall narrative. Genette originally (1980:86) called his chapter on this category of variation in narrative time *Duration (Durée)*, but he wrote in *Narrative Discourse Revisited* “the relevant feature is the speed of the narrative, and for that reason I think today I ought to have entitled that chapter not *Duration* but *Speed*, or perhaps (since, I suppose, no narrative moves forward at an entirely steady pace) *Speeds*” (1988:34). Each individual PE can be narrated with a different speed: it can be omitted or included, for one thing, but also narrated in a more or less verbose or syntactically difficult fashion. Although it may be reasonable to maintain the same speed while a particular span of events, there is no essential requirement that the speed be maintained at all from one event to another; certainly, the same speed is almost never maintained for an entire narrative. So a single speed, unlike an order, is not associated with an overall per-game narrative preference, nor with a reply that includes many events, but with individual events by way of their corresponding proposed expressions: Each PE has a speed.

The first four variations in speed each set the PE’s length factor L (which can range down to but not including 0 and up to 1, indicating the maximum length), while two of these also add a PE.¹⁶

To discuss variations in speed, it is useful to establish a base or default speed for each type of event or temporal sequence. These default speeds can vary among event types (for instance, a narrator can prefer to narrate IMPEL events at great length while barely mentioning MODIFY events) or in more specific ways (for instance, CONFIGURE events that move an agent between rooms may be narrated very quickly, other types of CONFIGURE events very slowly). These default speeds, however differently or uniformly they may be specified, are established to

¹⁶ Note that a shorter length L corresponds to faster narration; setting this parameter to smaller numbers means a higher speed. This is not intuitive in some ways when emphasizing speed over duration, but it allows the setting of $L=0$, a finite number, rather than the setting of speed to infinity, to indicate ellipsis.

indicate basic narratorial preferences, and the base speeds remain constant throughout the session. Changes in the speed of narration are accomplished relative to these base speeds.

Explicit ellipsis. An event or temporal sequence that would normally be expressed is not, and there is explicit mention of what has been omitted. There is no PE added to represent an action, but a commentary PE is added to indicate that some part of the narrative is being omitted. One sort of explicit ellipsis that might be generated is would indicate that the actions of a particular actor are not being narrated at all. The result might be text such as “What the mime is doing will not be mentioned.” There are many other ways to elide narrative while explicitly mentioning the ellipsis, such as by naming an area in which actions are not being narrated (for example, “No mention shall be made of what happened in the southern part of the plaza”) or by calling attention to the omission of one event in a sequence (for example, “The boy picked up the ball. I will not say what he did next. The ball smashed the window.”)

Implicit ellipsis. An event or temporal sequence that would normally be expressed is not, and there is no mention of this omission. The system accomplishes this by simply not adding a PE; a length L of 0 is never explicitly set for a PE. This avoids the problem of aggregating one PE with another one that is going to be expressed, and also linking expressions such as “then,” “meanwhile,” and “before that,” being used to connect two PEs when the first of these will not end up being expressed. In reading the resulting narrative, there may be no way to determine that something was omitted, although the other events that are represented may make this evident: If a room is first described as empty, and then a character is represented as dribbling a basketball in the room, the character and the basketball must have entered the room at some point, and the narration of this event must have been implicitly elided.

Summary. An event or temporal sequence is expressed in less than the standard number of words: Set L to whatever value l , $0 \leq l < .5$, is specified.

Scene. An event or temporal sequence is expressed in the base default of words. $L=.5$.

Stretch. An event or temporal sequence is expressed in more than the default number of words: Set L to whatever value l , $.5 < l \leq 1$, is specified.

Pause. Pause indicates that text not representing an event — it may be descriptive or it may be some non-diegetic text — is inserted at some point in a temporal sequence, via a PE.

7.2 Frequency

Frequency determines how a set of events (including either one or some larger number n) will be narrated. It can apply to individual PEs, but some sorts of variation in frequency — namely, iterative narration — are not possible unless a set of similar events are chosen. So, frequency in general is not assigned to individual events alone, nor to reply structures, but to sets of similar events. As an approximation, we can use the same rules for assembling sets that are used for sylleptic narration: any similar set of events can either be narrated in a singulative fashion as one branch of a sylleptic narration, or it can be grouped together and narrated once in an iterative narration.¹⁷ The Reply Planner implements different sorts of frequency by adding one or more PEs to again narrate an event already represented in a PE or by merging several PEs into one. The way in which text is generated from such a merged PE in the Microplanner and Realizer is discussed later.

The case of narrating a single event more than once is rare, but this sort of narration is used at times. Given n events which can each be narrated between 1 and f times, nf variations in frequency for the whole sequence are possible. More commonly, a set of m like events is narrated once instead of m times. Potentially, any partition of a set of events corresponds to a possible variation in frequency (with each part being narrated once), so the number of variations of this sort for n events is given by the Bell number B_n and computed with Dobinski's formula:

$\frac{1}{e} \sum_{k=0}^{\infty} \frac{k^n}{k!}$ It is not reasonable to narrate arbitrary subsets of events with a single expression, however, so this upper bound is not very meaningful in practice.

1 to 1. (Singulative.) The default case.

n to n. (Singulative.) The default case.

1 to n. (Repetitive.) PEs are duplicated and added in the appropriate places to narrate the same event many times. This unusual case can be seen, for instance, in Alain Robbe-Grillet's *La jalousie*. (It is not clear in this novel, admittedly, whether some multiple narrations refer to the

¹⁷ As a practical matter, being able to tell multiple events with a single narration may only work with reasonably similar events, while gathering events into sets for sylleptic narration may not require as much similarity. Sylleptic narration can effectively categorize events based on something as loose as theme or location, and does not require that the expression of the events be aggregated together into a single statement. Because of this, certain events that can be narrated together using syllepsis because they happened in the same place (for instance, "the adventurer killed the dragon, and then the adventurer acquired the grail") may disappear into meaninglessness if narrated as a single expression (in this case, something along the lines of "the adventurer altered some entities").

same event or similar events occurring at different times, but Franck's destruction of a centipede, for instance, is probably something that occurs once and is narrated several times.) The technique is used for emphasis and for other effect more sparingly in other novels, such as Don DeLillo's *Underworld*. The number of repetitions and placement of the PEs needs to be specified as well.

n to 1. (Iterative.) Starting with a set of *n* like events (and the PEs which encapsulate these), group the events in each of the PEs into a single PE. Perhaps the most famous example of this more common type of narration is the opening sentence of Proust's *À la recherche du temps perdu*, "Longtemps je me suis couché de bonne heure." (In Scott Moncrieff's translation: "For a long time I used to go to bed early.") Depending upon the set of like PEs and the other PEs in the sequence, the iteration will be either generalizing (external) or synthesizing (internal).

It is also possible to formulate ellipsis as a form of **1 to 0** or **n to 0** frequency, although Genette does not take this step. The "n to 0" formulation has some possible advantages; it could be used to explain, at once, why a whole set of like events are not narrated. However, the same thing can be accomplished by specifying an ellipsis (speed) rule that applies to a category of events and prevents their expression as PEs. For this reason, Genette's singulative, repetitive, and iterative categories will suffice and additional formulations will not be necessary.

7.2.1 A Connection between Order and Frequency

The independent presentation of order, speed, and frequency may have suggested that the three are independent. While they can be discussed individually as discernibly separate variations in the narrative discourse, there are definite ways in which they work together to accomplish particular effects, as is the case with many elements of the narrative discourse.

Mieke Bal, who uses *retroversion* to indicate the same general re-ordering that Genette calls *analepsis*, notes the effect of what she calls *punctual retroversion* and *durative retroversion*. A punctual move to the narrative past is one that skips back only to narrate a single event; a durative one narrates a longer sequence from the "once" of the narrative:

Frequent use of punctual anachrony sometimes makes for a businesslike style; systematic combinations of punctual and durative retroversions can create — or at least add to — the impression that the story is developing according to clear, causative laws ... If durative retroversions are dominant, then the reader quickly receives the impression that nothing particularly spectacular is happening. (Bal 1997:94)

This analysis can be extended to comment on the use of analepsis to a singulative narration and

analepsis to an iterative one. First, note that since a single event cannot be narrated in an iterative way, punctual retroversion or analepsis is incompatible with narrating the past iteratively. On the other hand, iterative narration during the analepsis is a possibility whenever there is a durative move to narrating past events. (Including several events in the range analepsis does not guarantee that the analepsis will be a single iterative narration, of course; it could include several singular narrations or it could include a combination of iterative and singulative narration: “I used to hang out in cafés around here. One day I found my favorite one had closed. I would occasionally run into strange characters in these places. ...”) To narrate the single-event span of the analepsis singulatively, then, is a punctual case, while narrating iteratively is a durative case. But the assertion here is that beyond that, the shift from singulative to iterative frequency leads to a different effect when it is occurs analeptically.

A singulative analepsis (such as “a murder occurred here in February 1975”) provides a businesslike style, as Bal noted, and is more compatible with a heightened sense of urgency than is an iterative analepsis, even if the analepsis includes only one narration that encompasses similar (as in “I used to hang out in cafés around here during the Web boom”). In the two examples that follow, some context for the analepsis is provided, both examples have similar content, and (1) shows the effect of singulative analepsis, as opposed to iterative analepsis in (2):

(1) I kept walking down Geary. A murder occurred here in February 1975. I saw Union Square.

(2) I kept walking down Geary. In the mid-Seventies, there was a lot of really bad violent crime here. I saw Union Square.

The narrator mentions something more focused and specific in (1) than in (2), which influences the style and how businesslike and urgent these narratives seem.

(3) I kept walking down Geary. I drank a cup of coffee here about ten years ago. I saw Union Square.

(4) I kept walking down Geary. I used to hang out in cafés around here during the Web boom. I saw Union Square.

The difference between the style of (3) and that of (4) is similar to the difference between (1) and (2), even the content of both of the latter two narratives is more light-hearted.

Given the effect of using iterative narration during an analepsis, even if there is no

specification to narrate events in the story's "now" iteratively, the desire to narrate past events in this way may still arise. This provides one example of why it would not be adequate to set order, speed, and frequency independently.

7.3 *Narrative Mood*

Genette defines *mood* in narrative by analogy to grammatical mood, quoting specific parts of a definition of grammatical mood: "one can tell *more* or tell *less* what one tells, and can tell it *according to one point of view or another*; and this capacity, and the modalities of its use, are precisely what our category of *narrative mood* aims at" (Genette 1980:161-162). Mood includes both distance and perspective (Genette 1980:162-211, Genette 1988:44-78). Distance refers specifically to the distance between the events and their narration, with richer detail (usually associated with slower speed) and less overt narration leading to more proximity. More cursory narration that makes the narrating itself more obvious leads to a greater distance. Distance includes the extent to which a narrative is either mimetic or diegetic, in Plato's terms — that is, how much of the content is presented by "showing" as opposed to "telling."

The qualities of mood dealt with directly in the current project, however, are those of perspective — specifically, focalization. As described in the section 5.2, Focalizer Worlds have been modeled separately from the base simulation of the IF Actual World, allowing the system to narrate from any actor's perspective while taking into account both the immediate perceptual ability of that actor and the knowledge of the world that the actor has.

The position taken here on distance is that it is better understood as a composite of several lower-level, more fundamental narratological variations. Without attempting a complete model of distance, and without attempting to allow all sorts of variation relevant to distance, it is still possible to affect how distant or how close the narration seems in an interesting way. This is explored in section 9.1 in the context of a creative work called *Lost One*. This piece varies distance interactively, based on whether the user is commanding the player character to wander about or whether the character is staying in the same location.

7.4 *Narrative Voice*

The category of variation that remains is that of narrative voice, which covers the act of narrating

(or narrative instance) and its context. This includes the time of narrating relative to the events of the content plane (as discussed in section 8.4), the narrative levels on which narrating occurs (a topic not considered in this project), and what is loosely called “person” — the relationship between narrator, narratee, and what is told (Genette 1980:212-215). While “person” covers the distinction between first-person, second-person, and third-person narration, the matter of narrative voice also covers how prominently the narrator and narratee are indicated and what signs of them are present in the text (Prince 1982:7-26).

The presence of first-person verbs in a narrative text can therefore refer to two very different situations which grammar renders identical but which narrative analysis must distinguish: the narrator’s own designation of himself as such, as when Virgil writes, “I sing of arms and the man ...,” or else the identity of person between the narrator and one of the characters in the story, as when Crusoe writes, “I was born in the year 1932, in the city of York. ...” (Genette 1980:244)

The former sort of narrator is heterodiegetic (the narrator does not correspond to a character), while the latter is homodiegetic (the narrator is present within the narrative as a character). The narration generated by nn can have either sort of narrator; the “I” can be assigned to any actor or to None, in which case the narrator is heterodiegetic.

As a first means of offering authorial control over the way that heterodiegetic narrations are generated, the level of signs of the narrator can be varied between 0 and 1, corresponding to the minimum and maximum, with 0 as the default. Obviously, there are many other ways in which such signs can vary, but it is worthwhile to first be able to change the level of essentially “neutral” signs of the narrating. A setting of 0 corresponds to as few signs as possible, but, almost certainly, not a complete absence of signs. One position holds that any narrative is itself a sign of the existence of some narrator, so it is not possible to eradicate all signs of these from a generated text. Although it has been argued that some narratives are “non-narrated” and have no narrator, it is extremely difficult to eradicate all signs of a narrator in a verbal narrative. The tense will signal the temporal position of the narrator relative to events, for one thing, and adjectives will usually suggest some evaluation on the narrator’s part.

The “I” that narrates can be indicated implicitly (for instance, with “Fortunately, ...”) or explicitly (for instance, “I narrate that ...”), each more or less often. These are probably not truly independent; an analysis of explicit and implicit signs of the narrator in existing narratives would almost certainly show a strong positive correlation between the two. Generation can sensibly happen with either one set anywhere in its range, however, so these can be provided to IF authors

as independent variables: $0 \leq \text{exp-signs-narrator} \leq 1$ and $0 \leq \text{imp-signs-narrator} \leq 1$.

The narratee can be specified the same way: as any actor or as None. There is no condition that both be actors or both be None. Genette's observation that "[l]ike the narrator, the narratee is one of the elements in the narrating situation, and he is necessarily located at the same diegetic level" (Genette 1980:259) does not really argue against the system permitting this freedom. The "you" to whom the narrative is addressed can be indicated implicitly or explicitly, each more or less often. "The lights blink, indicating that intermission is nearly over" includes an implicit sign of the narratee, for instance; the text assumes that the narratee does not know what it means for lights to blink during intermission and needs to be told. An example of an explicit sign is seen in a sentence beginning with "You may be interested to know that ...". As with signs of the narrator, $0 \leq \text{exp-signs-narratee} \leq 1$ and $0 \leq \text{imp-signs-narratee} \leq 1$.

Signs of the narrating can also be present in a text and can be more or less explicit. "The thief dashes out of the room in less time than it takes to tell it" provides one example; "The preceding recounting of events does not do justice to the strangeness of what happened" is another, more explicit example. The project focuses on signs of the narrator and narratee exclusively, but future work should be able to allow signs of the narrating to be parametrically generated as well.

7.5 *Interactive Mood and Voice?*

The existence of interactive time has been described briefly at the beginning of section 7.1. It is possible to also speak of interactive mood and voice. As described in sections 6.2 and 6.3, some of the output texts generated as part of a reply are not primarily narrative and addressed to a narratee; they are more strongly related to the interactor's working through puzzles and interacting with the world via the character. The text "You notice something unusual" would be in this category, as it cues the interactor to investigate further in a way that is almost certainly more important than its subsidiary narrative purpose. These sorts of text should be distinguished both from ones that describe the state of the program (for example, "Press SPACE to continue", "Do you really want to quit?") and from ones that have a clear narrative function (for example, "A 'lean and hungry' gentleman just wandered through, carrying a large bag."). In the case of interest here, the interactor's ability to command the player character in the simulated world, not to issue directives such as "quit," is what warrants this sort of text, so it can be said to be

addressed to the specific function of the interactor as *commander* as opposed to *user*, a generic role to applies to any computer program. The directive-giver might reply to the program, while the interactive fiction's corresponding function for the command giver would be that of the *suggester*, an entity that offers suggestions about what to do. This function of interactive fiction can still be called *suggester* even if the suggestions given are intentionally misleading.

Interactive mood can be seen as regulating interactive (rather than narrative) information. Distance and the telling versus showing distinction can correspond to direct instruction or hinting (as in "Maybe you should check out the south wall") as opposed to more indirect suggestion (as in "Something seems odd about the south wall"). Perspective, and focalization specifically, can be seen to relate to whether the suggester's information is limited to that of a character within the interactive fiction or is total.

Interactive voice corresponds to the signs of the suggester, which includes whether the suggester speaks as a character (as is the case with the frog in Robert Pinsky's *Mindwheel*) or not. Signs of the suggester, like signs of the narrator, can be more or less explicit: "You should check out the south wall" does not signal the suggester as explicitly as "I really think you should check out the south wall."

Signs of the suggester and commander could be varied in the same way that signs of the narrator and narratee can be, but interactive mood and voice should generally be controlled with at least as much regard for human-computer interface (HCI) principles as for narratological ones, and, no doubt, with reference to a model of the interactor that includes a representation of the interactor's understanding of the simulated world (as commander) not just his or her ability to operate the program (as user). Because a general scheme for generating these signs appropriately would involve using practices from and doing evaluation in terms of yet another discipline, HCI, and because these signs are not as common or as clear in interactive fiction as are signs of the narrated and signs of the narrating, signs of the interaction were not a focus of the initial project.

7.6 *Other Parameters of the Plan for Narrating*

The plan for narrating allows the focalized character, the narrator, and the narratee to be freely set to be any actor — or no actor at all. There is another important function of interactive fiction which derives from its nature as a computer program that accepts input, and which does not

have an analogue in narrative. Not only can one particular character be focalized, a particular character can be commanded. In almost all existing IF, these characters are always the same; they are the entity called the player character. But the two functions of focalized character and commanded character do not have to be locked to the same entity. In the opening sequence of Scott Starkey's *The Beetmonger's Journal*, the interactor commands Victor Lapot by typing input at the prompt, which is initially "What did Monsieur Lapot do next?" But the focalizer is Aubrey Foil, "Monsieur Lapot's chosen chronicler and assistant." Foil always remains in view of Lapot, so the strange situation of commanding an actor who cannot be seen by the focalizer does not arise. Nevertheless, this piece features a very unusual split between these two functions.

In *mn*, the focalized character and commanded character can be set independently and arbitrarily, even when the former cannot see the latter. The commanded character can also be set to None, in which case there is no opportunity for input and the systems runs without input until it reaches a conclusion, if it ever does.

There are a few other miscellaneous parameters of the plan for narrating. *indicate_known_directions* results in the printing of messages such as "South is the slit in streambed" if set to True; these messages are suppressed if it is set to False. Only if *room_name_headings* is True will the name of the current room be printed as a heading when the focalizer looks at the room, so that something like "- End of the road -" will appear before the description of the room.

8 Microplanning and Surface Realization

The previous two chapters discussed the first, highest-level component of the Narrator in the context of Genette’s theory of narrative discourse. This component, the Reply Planner, is the first stage of a three-stage pipeline within the Narrator. After the initial choices have been made — choices about the inclusion or exclusion of events; the ordering of events; the speed with which each should be narrated; and the time of speech, reference, and event — the representation of an event must still be converted first into an abstract syntactical representation and then into a string. These functions are carried out by two additional components which access the Discourse Model. These two do not need to refer to the plan for narrating, since the Reply Planner has already used this plan to formulate its output. The two additional components are the Microplanner — which accepts a reply structure as input and provides a proposed section of text — and the Realizer, which takes the proposed section and converts it to a string, ready to be output.

8.1 *Microplanner*

The idea for a microplanning stage in the Narrator comes from work in the last decade that has revealed that “there are a number of NLG issues that do not seem to be centrally concerned with text content and structure, nor with sentence-level syntax and morphology, but which are nonetheless important in building systems that produce high-quality texts” (Reiter and Dale 2000:115). There is not complete agreement about which tasks are best handled between the better-established first “strategic” stage and the better-established last “tactical” one (Thompson 1977), but they include aggregation, done in a rudimentary form in nn, and also referring expression generation, which is only done only in one way, by indicating which NPs should be pronominalized. The Microplanner that is part of nn’s Narrator mainly handles a task specific to a narrating system — the task of mapping between a specification for narrating and the detailed grammatical information that the Realizer needs to produce textual output. The Microplanner has some additional tasks related to other parts of the plan for narrating. For instance, it uses the length specified for each proposed expression, which has been determined by the Reply Planner

based on the setting for speed in the plan for narrating, to generate more or less text.

8.1.1 Mapping the Narrative to the Grammatical

The Microplanner takes as input a detailed representation of what is to be narrated that specifies the order of events and their relationship to one another and to the narrator in time. This is the Reply Structure (RS). It outputs an abstract representation of a section, which can contain a sequence of abstract representations of paragraphs and headings. The paragraphs can contain abstract representation of sentences, which themselves contain sequences of parts. These parts can be strings which are to be printed out essentially as they are (with some orthographic changes such as punctuation and sentence-initial capitalization) or they can be special symbols that the Realizer will convert into strings using grammatical information that the Microplanner specifies for each string.

To determine this grammatical information, the Microplanner refers to the plan for narrating, the proposed expression (PE) itself, and the way that PEs are arranged in the reply structure. For instance, if the time of narrating is simultaneous and an event is being related in an analepsis, the Microplanner may use this information to determine that the sentence narrating this event will be in the past tense. Events in the main sequence would be related in the present tense in this case, but within an analepsis, reference time and speech time would be fixed to the “now” of the telling while an event from the past is related, so the tense would shift. If the same event were in the future relative to what was being told (event time changes while everything else remains the same), the Microplanner would indicate that the sentence should be in the future tense.

Specifically, the Microplanner considers the relationship between the three types of time that are associated with each PE: reference (R), speech (S), and event (E) time. The value for E is always initially specified as numerical, when the event is registered during simulation. When a reply structure has been set up by the Reply Planner and sent to the Microplanner, each internal node is supplied with rules for both R and S. As described in section 6.4, these rules are either numerical values or one of (FOLLOW, MIN, MAX, HOLD).¹⁸ The Microplanner determines the specific R and S values for each PE by using the rules on its parent; it then determines the tense

¹⁸ There are also rules to set reference time to one greater than or one less than event time. These are convenient because they allow aspect to be easily changed. To simplify the discussion of the basic cases, these two rules are not described in any more detail.

that will be used to express this event by considering the relationship between E, R, and S:

$E < R \rightarrow$ anterior	$E = R \rightarrow$ simple	$E > R \rightarrow$ posterior
$R < S \rightarrow$ past	$R = S \rightarrow$ present	$R > S \rightarrow$ future

For each action indicated by a reply structure, the Microplanner generates proposed sentences and at least one proposed paragraph. If the speed is slow enough, the Microplanner will generate a proposed sentence introducing the action. (This will end up being output as a text such as “The adventurer will start to examine the area below the grate.”) Whether or not this happens, though, some text associated with the action will eventually be produced, since an action included in a PE cannot be entirely elided. The Microplanner goes through each event in the action, producing proposed sentences for each one that is salient enough, given the speed with which this action is to be narrated. When the event is SENSE, information will be included in the text describing the appearance (in the case of sight) of whatever was sensed. Time words are added at this stage, too. When the Microplanner finishes its work and outputs a proposed section, there are no simulated events and simulated times in the output, only representations of sentences. These representations do have the tags of existents embedded in them, however; their names are generated in the final step, as verbs are conjugated.

8.1.2 A String-with-Slots Representation for Building Sentences

The string-with-slots representation was developed as a compromise between a simple string representation, which would not have been manipulable in any way, and a richer abstract sentence representation, which would have been prohibitively difficult for IF authors to use. There are two main places where strings-with-slots are used: in representations of events and in the *appearance* properties of existents. From the IF author’s perspective, either one can be custom-written for a particular game, although standard strings-with-slots are provided by the nn system for representing events. Existents are always specific to works of IF (although they can be re-used in various works) and so are always custom and author-provided. Representations of events are, from a literary standpoint, also very useful for authors to specify. Currently, the particular strings-with-slots that are used to generate event representation are hard-coded in the Microplanner. In the release version of nn, they should be moved to somewhere where they can be more easily seen and in a way that allows them to be easily replaced and overridden. Placing

them in the Discourse Model, with allowance for defining event-related strings-with-slots in a game file, is probably most appropriate.

As an example of how the formalism is used, consider the problem of generating the text that begins this author's *Winchester's Nightmare*, for instance:

Sarah Winchester has forgotten being awake. It is night, or predawn morning, and moonless. She is on a sandy strand extending north and south from here. The sea is before her to the east.

The following strings-with-slots, along with the appropriate existents, will work:

'S_FC V_forget_PERF being awake',
'it V_be_S night, or predawn morning, and moonless',
'S_FC_PN V_be on a sandy strand extending north and south from D_HERE',
'the sea V_be_S before O_FC_PN to the east'

The subjects are indicated with "S_", the verbs with "V_", and the objects with "O_". The addition of "_PN" at the end of a noun forces it to be realized as a pronoun. The modifiers for verbs that are seen here are "_PERF", forcing the perfect aspect, and "_S", which causes verb formation to be done assuming a singular subject. "_S" is used in the second line, when there is no subject that can be easily simulated. In these strings-with-slots, there are no existents mentioned by tag. To generate this particular text, the tag "SARAH" could have been used wherever "FC" appears. As written, however, these strings-with-slots apply very generally to any focalized character. So if one removes the SARAH existent, replaces it with the BARNEY existent, and sets BARNEY to be the focalized and commanded character, this text results:

Barney has forgotten being awake. It is night, or predawn morning, and moonless. He is on a sandy strand extending north and south from here. The sea is before him to the east.

Or, a conventional "first-person" IF style can be generated by setting SARAH as the narratee:

```
>pfm narratee sarah  
  
<<< The narratee has been set to sarah. >>>  
  
>look  
  
You look at the shore.  
  
- Shore -
```

You have forgotten being awake. It is night, or predawn morning, and moonless. You are on a sandy strand extending north and south from here. The sea is before you to the east.

The representation is adequate to allow changes in tense, so that the time of narrating can be changed from simultaneous to subsequent, or, less conventionally, to a time previous to the events:

>pfn time previous

<<< The time of narration has been set as previous to the simulated events. >>>

>pfn narrator sarah

<<< The narrator has been set to sarah. >>>

>look

I will look at the shore.

- Shore -

I will have forgotten being awake. It will be night, or predawn morning, and moonless. I will be on a sandy strand extending north and south from there. The sea will be before me to the east.

When the appearance of an existent needs to be described, one can either write a string-with-slots that will be used to complete a sentence along the lines of "You see ...", or one can supply the whole sentence, or sentences, that are to be used. In this case it is necessary to refer to the actor doing the sensing. This is done in appearance strings with "*S_?*" and "*O_?*".

The number of verbs is determined automatically during realization whenever possible. If there is exactly one singular subject that comes before the verb in the sentence, the singular form of the verb is generated; otherwise, the plural form is used. "*_S*" and "*_P*" can be used to force the verb to be of a particular number, as in the example above.

This flattened representation of a sentence does not, of course, work to model complex syntax, which is a disadvantage for a literary system. But it is powerful enough to generate many, almost certainly most, sentences that occur in existing IF, and it is simple enough that the process of developing a string-with-slots is linear and is at least somewhat like writing rather than resembling a detailed data modeling or syntax representation process.

Strings-with-slots are converted to sentence representations in the Microplanner with little effort. The additional information that must be added to these strings includes grammatical information specifying tense and number along with some information about deixis. With this specified, a string-with-slots is a complete sentence representation ready for realization.

8.1.3 Section, Paragraph, and Sentence Representations

Although the representation provided by the Microplanner has more information than is shown here, this printout of a section indicates its basic structure:

---SECTION---

P-((S-(N.VISITOR.0 V.look.0 at N.PLAZA-CENTER.1)))

P-((S-(P.VISITOR.2 senses V.hum.2 as N.VISITOR.0 V.view.0 the broad , circular , encircling Plaza of the Americas) S-(the morning V.conclude.1) S-(it V.be.1 midday D.NOW)))

P-((S-(N.VISITOR.0 V.see.0 that N.PLAZA-CENTER.0 V.contain.0 N.FLANEUR.1)))

P-((S-(from D.HERE , N.VISITOR.0 V.see.0 in the distance N.STATUE.1 to the north , N.FOUNTAIN.1 to the east , N.COLLECTOR.1 to the southwest , N.BOY.1 to the northeast , N.OBELISK.1 to the south , and N.PUNK.1 and N.TREE.1 to the west)))

P-((S-(N.FLANEUR.0 V.go.0 to N.PLAZA-N-INNER.1)))

P-((S-(N.PUNK.0 V.kick.0 N.TREE.1)))

P-((S-(N.BOY.0 V.throw.0 N.BALL.1)))

P-((S-(N.COLLECTOR.0 V.pick.0 up something)))

P-((S-(N.BALL.0 V.fall.0 to the ground)))

The nouns (beginning with “N.”) contain words in all uppercase, such as “VISITOR,” which are tags referring to existents. It is necessary to see what the focalizer thinks the corresponding existent is before an NP can be realized from these. The number at the end of the noun indicates the form; 0 is the subject, 1 the object, 2 the possessive, and 3 the reflexive. Not shown, but present in the representation, is a flag which indicates whether the existent should be realized as a pronoun. In English, the distinction between the subject and object form is only important in this particular case.

The verbs are indicated in the style “V.look.0”, with the base form of the verb in lowercase and the default number of the verb at the end. In this case, “0” means there is no default and number of the verb must be determined by the context. The number of the verb will be determined by the subject — in this case, VISITOR being the subject means that look is singular. If there is no representation for the subject in the IF world, number needs to be indicated as a parameter, which is represented with “1” for singular and “2” for plural.

This section is realized as follows:

You look at the center of the plaza.

Your senses are humming as you view the broad, circular, encircling Plaza of the Americas. The morning has concluded. It is midday now.

You see that the center of the plaza contains a flaneur.

From here, you see in the distance a statue to the north, a fountain to the east, a trash collector to the southwest, a boy to the northeast, an obelisk to the south, and some punk and a tree to the west.

The flaneur goes to the northern area.

The punk kicks the tree.

The boy throws a ball.

The trash collector picks up something.

The ball falls to the ground.

8.2 Realizer

The Realizer's task is straightforward: Convert the section representation into a string of English text. Realization code is included on each Section, Heading, Paragraph, Sentence, and part of speech class in realize methods. So, realizing a section is a simple matter of calling the realize method on the Section instance of interest; this results in realize being called on each of the items (which may be Paragraphs or Headings) in the section. Within each Paragraph, realize is called on each Sentence. Finally, each Sentence calls realize on each of its parts. There are no dependencies on other parts of the sentence or across wider spans of text, although a world instance must be passed to the realize methods as a parameter along with the focalized character (so that the noun or pronoun with the tag "FC" can be realized properly) and the Realizer must have access to the list of givens (which will grow as new existents are mentioned). The world used to provide information about existents to the Realizer is the Focalized World corresponding to the current focalized character.

For adjectives, realization is a simple, two-step process. It first means looking up the appropriate state to determine, for instance, whether a lamp's LIT state is False or True. Then, the English word or phrase associated with that state is looked up — it is either "unlit" or "illuminated," in this case.

Noun realization is a bit more complex. Pronominalization can happen at this stage, if the existent that is to be realized refers to the narrator ("I") or narratee ("you"). Or, pronominalization may have been indicated explicitly already. If an existent is not going to be pronominalized, though, the process involves looking up the current article and name, making the noun possessive if necessary, and updating the given list if necessary.

The pronoun class contains a list of English pronouns and indexes into it based on number, gender, person, and form. The appropriate pronoun is selected from this list.

Deictic expressions are realized based on whether the narrator is in some way there at the events, which corresponds to narrating in the present tense. If the narration is present-tense, "D.NOW" is realized as "now"; if not, it is realized as "then." The other words that can be shifted in this way are here/there, this/that, and these/those. The description of the center of the Plaza in *Lost One* uses D_NOW to produce "The morning has concluded. It is midday now." in the present tense and "The morning had concluded. It was midday then." in the past tense.

Verb realization is by far the most involved. No unencumbered system (free for academic and IF author use) could be located, so verb realization was implemented from scratch. All the standard tenses are supported by nn's Realizer in any combination of progressive and perfect aspects. The system also supports generation in Reichenbach's posterior future ("you will be about to go east"), the use of different future styles ("will," "shall," and "going to") and the generation of intensive and negated verbs. Still, the system used is rather simplified and relies on a dictionary of more than 1000 "irregular" verbs which includes many consonant-doubling cases. With this and the built-in support for conjugation, at least ten thousand of the most common verbs can be correctly realized. There is no support yet for some sorts of generation, such as adding modal auxiliaries, but since the system does not have the ability to produce Reply Structures representing different modalities this does not present an immediate difficulty.

Most typesetting details (all the ones that are not managed by the Presenter) are handled by the Realizer. Headings are printed with a certain amount of space before and after and paragraphs are indented by a certain amount. Indentation can be suppressed for the first paragraph in a section or after a heading. Sentences are capitalized and given final punctuation if they do not already have it. If there is an output filter (such as the distance filter with distancing expressions in *Lost One*, described in the next chapter, or the filters in the Hesitant and Surprise plans for narrating in appendix E), it is applied in the Realizer. Output filters can operate individually and independently on parts, Sentences, and Paragraphs. More sophisticated filters that are sensitive to syntax as well as text structure could be developed with little effort and with no major changes to the framework.

9 Creative Work Using Narrative Variation

Two short interactive fiction pieces, intended as the first two in a collection to be called *The Ground Set*, were developed along with the system to demonstrate the capabilities of nn and to guide its development. Implementing existing interactive fiction in nn is useful to some extent, but has serious limitations. This is the case not only for the venerable but simple *Adventure* and the minimal *Cloak of Darkness*, but even for more sophisticated recent works of IF. Since all interactive fiction created so far has been developed without facilities for narrative variation, none of it, when ported unmodified to nn, will make use of the system's main capability. The *Ground Set* was created to show some interesting examples of narrative variation in an original setting and in closely related original IF worlds.

These two pieces are probably better characterized as demos rather than games, but, nevertheless, they show some ways that changes in the plan for narrating can be done interactively during a session and some ways that these changes can have literary effects. These techniques can clearly be used in larger-scale works of interactive fiction that have other literary aspects not directly related to narrative variation, such as riddle-like, figurative systems for the interactor to explore and solve; powerful descriptive writing; characters who have compelling behavior; and locations that refer to and play with various real and fiction spaces.

The pieces in *The Ground Set* share a significant subset of their existents, including all of their rooms. They take place in the Plaza of the Americas, a setting that was created by drawing upon North American and European public spaces but which does not correspond to any real urban location. The monuments and large features of the Plaza are present in all three pieces, as are various scraps of trash. Randomness can be easily used in an nn interactive fiction, but in these three pieces, what the actors do and everything else in the simulated world, during an interaction, is deterministic. This choice was made to allow these pieces to be used more effectively as demos and examples, and so that one session can be more easily compared to another. Some aspects of the generation of text are done at random, however, so the system does not always present the identical strings each time the same inputs are given.

9.1 *Distance in Lost One*

In *Lost One*,¹⁹ the player character is a woman about whom very little is specified. She has returned to the city where this interactive fiction is set, one where she had previously resided. She is now supposed to meet an old friend here in the Plaza. The two basic strategies for locating someone in a large public space are, first, to stay put hoping you will be found, and, second, to walk around trying to find the person you are looking for. *Lost One* recognizes which of these strategies the interactor is pursuing and either leaves the plan for narrating alone (in the first case) or changes it to increase the narrative distance (in the second).

A transcript of a complete, successful traversal of *Lost One*, using the first strategy, follows. As the transcript shows, the interactor can “win” — a rather hollow victory — simply by staying put while looking around, waiting, and engaging in other unremarkable activities. The plan for narrating is never changed in this case, and all actions are narrated in chronological order.

Lost One

from *The Ground Set*, 2007
by Nick Montfort

The sun hits the plaza. Now they drive cars, seeking flatpacks across the sprawl. Once they were supposed to cluster here. You have arrived, visitor to this place where you briefly lived years ago, where you knew spaces and faces now almost forgotten. There is one here less lost to you than the others, though, and it is right here in this plaza, about now, that you are to meet him. Somewhere right around here.

>look

You look at the center of the plaza.

Your senses are humming as you view the broad, circular, encircling Plaza of the Americas. The morning has concluded. It is midday now.

From here, you see in the distance a statue and a flaneur to the north, a fountain to the east, a trash collector to the southwest, a ball and a boy to the northeast, an obelisk to the south, and some punk and a tree to the west.

The flaneur goes to the northern area.

The punk kicks the tree.

The boy throws the ball.

The trash collector picks up something.

The ball falls to the ground.

>examine the trash collector

You look at the trash collector.

You see a nondescript man who is, to all appearances, a bona fide trash collector.

¹⁹ Neither of these pieces have been made available to the public, but *Lost One* was presented at the Electronic Literature Organization Open Mic/Open Mouse at the University of Maryland on May 2, 2007.

The boy takes the ball.
The trash collector conveys himself to the western area.

>examine myself

You look at yourself.
You see someone who is out of place, although not attired like a stereotypical tourist.
You see that you possess a tasty tortilla.
The trash collector takes something.

>eat the tortilla

You consume the tasty corn tortilla.
The flaneur wanders around briefly.
The punk conveys herself to the southern area.
The boy throws the ball.
The trash collector goes to the northwestern area.
The ball falls to the ground.

>examine the punk

You look at the punk.
You see a girl who is participating in the punk subculture.
The punk kicks the obelisk.
The boy picks up the ball.
The trash collector conveys himself to the northern area.

>examine the obelisk

You look at the obelisk.
You see the stone pointing the way it has for centuries.
The flaneur heads over to the northeastern area.
The trash collector heads over to the center of the plaza.

>wait

You wait.
The punk goes to the western area.
The boy throws the ball.
The trash collector walks to the southern area.
The ball falls to the ground.

>wait

You wait.
The punk kicks the tree.
The boy takes the ball.
Your friend steps up quickly from behind you with a greeting, and the two of you walk off.

- the end -

What happens when the interactor chooses to have the visitor walk around, in a more typical adventure-gaming mode of exploration, is more interesting. The more rapidly the visitor moves

from place to place, the more *Lost One* takes steps to increase the narrative distance — that is, to narrate in a way that is less immediate. These two cycles of command and reply show the first shift happening as the visitor moves from the northeastern to the northern area:

>examine the ball

You look at the ball.
You see a baseball.
The boy takes the ball.
The trash collector goes to the northern area.

>go west

You headed over to the northern area.
It appeared that the flaneur walked to the northeastern area.
You looked at the northern area.
It was as if you saw the space north of the plaza's center, which was particularly barren of vegetation and ornament.
You saw that the northern area contained the statue and a rock.
From there, you saw in the distance the trash collector to the south, the fountain to the southeast, the flaneur and the boy to the east, the obelisk across the plaza, and the trash collector and the tree to the southwest.
The trash collector conveyed himself to the center of the plaza.

Two changes happen at once after the “go west” command: First, some distancing phrases are inserted at random at the beginning and end of sentences. In this example, “It appeared that” is prefixed to the second sentence and “It was as if” is prefixed to the fourth. Such phrases continue to crop up in the rest of the transcript from this session. The probability of adding some distancing expression becomes 1/6 at this step; a sentence-initial one will be added half the time and a sentence-final one the other half. It is useful to be able to mix in fixed phrases like this, but a text generation system is hardly necessary for such a transformation. The second change relies more heavily on the Narrator's text-generation capabilities. The time of narrating is shifted to be subsequent to, rather than simultaneous with, the events. The effect is that the text is generated in the past tense. By moving the narrating apart from the events temporally, narrative distance is increased.

The next shift can occur in as few as two more turns, if the interactor keeps the visitor moving:

>go south

It appeared that you went to the center of the plaza.
The punk conveyed herself to the western area, evidently.
It seemed that the boy threw the ball.
You looked at the center of the plaza.

Your senses were humming as you viewed the broad, circular, encircling Plaza of the Americas. The morning had concluded. It was midday then.

From there, you saw in the distance the statue to the north, the fountain to the east, the flaneur, the ball, and the boy to the northeast, the obelisk and the punk to the south, and the punk and the tree to the west.

The trash collector walked to the southern area.

The ball fell to the ground.

The trash collector looked at the southern area.

>s

The visitor headed over to the southern area.

The punk kicked the tree.

The boy picked up something.

The visitor looked at the southern area.

The visitor saw the space south of the plaza's center.

The visitor saw that the southern area contained the obelisk.

From there, the visitor saw in the distance the statue across the plaza, the fountain to the northeast, the trash collector to the east, and the punk and the tree to the northwest.

The trash collector walked to the southeastern area.

Now the narratee, the one addressed as "you," is set to be None — that is, there is no actor within the IF world who is addressed as "you." The result is that the visitor, previously referred to in the second person, is now referred to in the third person. The idea here is not mainly to increase conventional narrative distance, although this is probably done to some extent, but to weaken the tie between interactor and player character by no longer stating that this particular actor is "you."

At any point, the interactor can command the visitor to return to the center of the plaza and wait there. Distance will slowly diminish and the "winning" ending can be attained by staying put. It will just take longer.

If the interactor continues to command the visitor to move about, the next transformation is a change in order:

The trash collector went to the southeastern area.

Before that, it seemed that the visitor looked at the southern area.

It was as if the visitor saw the space south of the plaza's center.

The visitor saw that the southern area contained the obelisk, or so it seemed.

From there, the visitor saw in the distance the statue across the plaza, the fountain to the northeast, the trash collector to the west, the trash collector to the east, and the punk and the tree to the northwest.

Before that, the boy picked up something, or so it seemed.

Before that, it appeared that the punk kicked the tree.

Before that, the visitor headed over to the southern area.

The order switches to retrograde, which is an unusual and artificial way of narrating and one that may suggest searching back in one's memory for events that are further and further removed. To signal that this unusual order is being used, time words are turned on.

One final change in the narrating is made if the visitor continues wandering: the order of

narrating is switched to random, an even more artificial and difficult order. The final change in the plan for narrating is actually not a variation in narrating, but in interactive control. To make the visitor as distant from the interactor as possible, and to shift the experience away from the interactive fiction framework entirely, the commanded character is set to None. This has the effect of not allowing the interactor to control anyone, or indeed to type anything in to the system: *Lost One* runs, instead, like a particularly pointless story-generating system, with the visitor wandering about at random under the system's control.

The flaneur went to the eastern area, perhaps.

Then, the punk conveyed herself to the southern area, if memory serves.

Then, the visitor looked at the center of the plaza.

Her senses were humming as the visitor viewed the broad, circular, encircling Plaza of the Americas. No doubt, the morning had concluded. It was midday then.

From there, the visitor saw in the distance the statue to the north, the flaneur, the trash collector, and the fountain to the east, the trash collector to the southwest, the boy to the northeast, the obelisk and the punk to the south, and the punk and the tree to the west.

Then, the trash collector took something.

Before that, the boy took the ball, or so it seemed.

Before that, the visitor walked to the center of the plaza.

The visitor waited, apparently.

The punk kicked the obelisk.

It was as if the visitor waited.

The ball fell to the ground, if memory serves.

Before that, the boy threw the ball.

Before that, the flaneur wandered around briefly, or so it seemed.

The punk headed over to the western area.

Then, the visitor looked at the southern area.

It looked like the visitor saw the space south of the plaza's center.

The visitor saw that the southern area contained the obelisk, or so it seemed.

From there, the visitor saw in the distance the statue across the plaza, the flaneur, the trash collector, and the fountain to the northeast, the trash collector to the west, and the punk and the tree to the northwest.

Before that, no doubt, the boy picked up the ball.

Meanwhile, the visitor walked to the southern area.

The visitor waited.

The punk kicked the tree.

Before that, the flaneur wandered around briefly, perhaps.

....

9.2 *Baseball Narration in Oddly Angled*

Narration in interactive fiction has some unusual qualities. One of these, often remarked upon, is the conventional use of “you” to refer to the player character. While there are some novels that have a character who is spoken of as “you,” such as Italo Calvino’s *Se una notte d’inverno un viaggiatore* (*If on a Winter’s Night a Traveler*), and while there are interactive fiction pieces that do not use the “you” (this author’s *Winchester’s Nightmare* is one example) this is nevertheless a prominent convention in interactive fiction that distinguishes it from the novel. A similarly well-established convention, and one which is also almost certainly inherited from the role-playing exchange of *Dungeons & Dragons*, is the use of simultaneous narration. This is rather unusual in non-digital prose fiction and in prose of any sort. Narrating in the present tense is usually called “the historical present” to emphasize that the use of this tense is rhetorical and, perhaps, to reassure us that the events did in fact happen in the past. In interactive fiction, however, events are being narrated as they are being simulated, so there is a particular type of presence to the narrating and it makes sense to think of it as happening simultaneously. To find an analogous situation in which events are truly narrated as they are told, it is worthwhile to consider the rich cases seen in sports commentary — specifically, the narration of baseball games by commentators as they happen.

The main aspect of baseball narration that is played upon in *Oddly Angled* is the tendency to retell an interesting play more than once (Ryan 2006:81-93), even as other events are passed over with a single cursory narration. A simultaneous, present-tense telling of events can be followed by multiple retellings of those same events which are more reflective and evaluative and are in the past tense. While this is not the only interesting feature of baseball narration, it is the one that inspired the heightened plan for narrating that has been implemented in *Oddly Angled*.

Oddly Angled shares the plaza setting of *Lost One*, and initially the narration is not very interesting. Two categories of actors are notable — the police and the performers — and their actions, as they do something that is vaguely reminiscent of a highly stylized game of baseball, are always narrated at least briefly. The other people on the plaza are not mentioned in the ordinary narration scheme except in the aggregate. If only one of the non-police, non-performer actors is doing something, that is narrated; otherwise their activities are all glossed in the

statement "People are doing some things."

Oddly Angled

from The Ground Set, 2007
by Nick Montfort

It's looking to be a beautiful day at the plaza. The skies are clear. There's a gentle breeze. The police have filtered out now and have taken their places. The performers are limbering up ...

>look

You look at the center of the plaza.

Your senses are humming as you view the broad, circular, encircling Plaza of the Americas. The morning has concluded. It is midday now.

You see that the center of the plaza contains a flaneur and a police officer.

A ball falls to the ground.

People are doing some things.

>wait

People are doing some things.

>wait

A crackling issues from the police officer.

A mime freezes.

A trash collector picks up something.

A crackling issues from a woman in a suit.

>wait

The ball falls to the ground.

People are doing some things.

>wait

People are doing some things.

>wait

A crackling issues from the police officer.

The mime waves.

A crackling issues from the woman in a suit.

People are doing some things.

>wait

The ball falls to the ground.

People are doing some things.

>wait

People are doing some things.

“A ball falls to the ground” is actually not part of the abstracted baseball-like activity; the boy is here throwing his ball into the air, an activity which is aggregated together with all the other things that people are doing. The ball’s descent to the ground is not being aggregated. The trash collector’s action is narrated on the second turn because he is the only “bystander” doing anything at that point.

The narration changes significantly after the next command because some interesting events have transpired:

>wait

A crackling issues from the police officer.
The mime walks to the eastern area.
A crackling issues from a park ranger.

The flaneur started to wander. The flaneur wandered around briefly.
Then, the police officer started to crackle. A really substantial burst of static came from her walkie-talkie there.
Then, the mime started to go. The mime got himself moving. The mime booked it over to the eastern area over there like there was no tomorrow, and he made it there with no problem.
Then, the trash collector started to take a plastic scrap. The trash collector lifted the plastic scrap. The trash collector took something.
Then, the park ranger started to crackle. A really substantial burst of static came from her walkie-talkie there.

The flaneur wandered around briefly.
A crackling issued from the police officer.
The mime went to the eastern area.
The trash collector picked up something.
A crackling issued from the park ranger.

For this turn only, the system uses a list of three plans for narrating, one after the other. The first one narrates a small amount of information in the present tense. The speed is slowed down substantially in the second narration, which is done from a standpoint subsequent to events, and everything that has happened is narrated. Finally, everything is narrated again, still in the past tense but much more tersely, as if in summary.

While this scheme does not capture the richness of real baseball narration — in which there are, for instance, references to events earlier in the game, in the season, in a player’s career, and in a team’s history, along with analysis and many more nuanced forms of variation — it shows how the narration can change in response to underlying events in the world. In this example, the interesting events are not due to the player character’s actions, but narration could be similarly

changed as the exploration of an area is being started or finished, as interactors hone in on the solution to a series of puzzles, or as some new secret is being uncovered through detection or conversation.

10 Pilot Evaluation

There were too far many open questions about the nature of a useful evaluation to begin immediately by having annotators trained at length and setting them to work on a large corpus. Certainly, it would be too much to expect statistical significance for results about the quality of generation and ability of different texts to convey information about the chronological order of events. The useful measures for this sort of system and domain were not known. It was not obvious what could serve as a baseline and an upper bound or goal, or if such were necessary. The useful ways of excerpting and displaying text for evaluation were also not known. While human-authored texts are readily available for many other domains (patient information, weather reports, sales copy), it is not obvious that there are human-authored texts representing interactive fiction events which are useful to compare to the system's output. IF outputs are not really human-authored, in the usual sense, to begin with. Although more or less of a reply may be pre-written as a string and the game has an overall programmer or author, an interactive fiction system is a new sort of dialogue. This dialogue borrows from role-playing exchanges but does not simply implement them.

An extensive evaluation of output texts might use the system to generate various permutations of tellings of events; the outputs could then be evaluated for naturalness and something could be learned about whether different orderings in the narrative were felicitous. This result would be valuable, but an experiment of this sort could not be undertaken initially for several reasons. Annotators would be biased depending upon how soon or late they read texts, or they would simply not be able to read and comprehend the texts very well, if all the texts were re-shuffled tellings of the same events. If they were not, the difference in events would introduce another variable. Other factors such as causality and the connotations of certain events would probably plague a small-scale study. If events told were not causally connected, there would be nothing to indicate their underlying order in time. If they were, the strength of their causal relationships might vary — some events would obviously cause others, while it would be less clear in other cases — leading to additional problems for the evaluation. The strength of causal ties could have been empirically tested, too, but that would involve an extensive additional study.

Instead of an extensive evaluation, a small pilot session was arranged in which two annotators rated 14 texts for their naturalness on a scale of 1 to 10, identified events by underlining portions of the text, and tried to determine the chronological order of events. Importantly, annotators were also asked afterwards to describe what made some of the texts seem natural or unnatural. While the pilot results will help in designing a more useful future evaluation, this qualitative advice from the annotators will be more useful than the quantitative results in helping to direct short-term system development.

To reduce bias, annotators were given the texts in the same randomly-determined order and were not told their sources. The annotators were both computer science Ph.D. students working in natural language processing. The first annotator was familiar with interactive fiction; the second annotator was not.

The texts evaluated are included, along with the instructions to annotators, as Appendix F. All of the texts are excerpts from interactive fiction. They include unmodified command-and-reply exchanges (1, 2, 5, 11,

12); a cut-up command-and-reply exchange (3); recountings generated by nn (4, 6, 8, 10, 14); a passage from a single interactive fiction reply that was rearranged to reverse order, paragraph-by-paragraph (7); and two lists that come from “full score” reports in pre-existing interactive fiction works (9, 13). Two of the generated recountings were produced in random order:

(4) was randomized but used time words, while (6)

	<i>Rank</i>	<i>Order</i>	<i>Ann1</i>	<i>Ann2</i>
1	3	Chronological	0.83	0.33
2	1	Chronological	0.83	1.00
3	4+	Random (no cues)	0.50	0.33
4	8	Random (cue words)	0.17	0.00
5	4+	Chronological	0.83	0.00
6	9	Random	0.00	0.00
7	6	Retrograde	0.50	0.00
8	7+	Retrograde	0.33	0.00
9	2+	None	1.00	0.33
10	7+	Retrograde	0.33	0.00
11	5+	Chronological	0.67	0.00
12	2+	Chronological	1.00	0.33
13	4+	None	0.83	0.00
14	5+	Chronological	0.67	0.00

Table 2: Ratings of naturalness from the two annotators, each normalized over the interval that annotator used in assigning scores. Order is not the only thing that varies, but it is shown here to characterize the texts.

used none, making it essentially a nn-generated cut-up. Two were generated in retrograde order with time words to indicate this order (8, 10).

10.1.1 Naturalness

There was not strong agreement between the annotators about the naturalness of texts. There was a moderate positive correlation (0.46) that, with these instructions and given this sample size, is significant only at $p = 0.1$. A more specific set of criteria rather than the briefly-defined “naturalness” might be needed to evaluate texts of this size. Alternately, annotators might be able to assess the overall naturalness better if they are given multiple questions about the text or the ability to assess different sections of it. Also, familiarity with interactive fiction conventions might be important in understanding whether or not a text is natural. Nevertheless, these rankings of naturalness suggest a few possibilities for future, more tightly-focused investigation with annotators, and their comments about what made the texts seem more or less natural will be helpful in further developing the system.

Normalizing the scores to the intervals used by the annotators, the highest average score went to (2), the excerpt from commands and replies in Ian Finley’s *Babel*. This excerpt was one of only two that included dialogue in direct discourse. The other excerpt with one brief (and odd) phrase quoted in direct discourse was (12), an excerpt of commands and replies from *Book and Volume* which tied for second. The presence of direct discourse may, by itself, make a text seem more natural; this could be controlled for or investigated in a future evaluation. The other text tied for second was (9), a list of tasks accomplished from *Book and Volume*. Lists of accomplishments did not seem to be generally rated as unnatural simply because they were lists; a full score report from Emily Short’s *Savoir Faire* (13) also ended up in the top half of the texts.

Ranked third was (1), the excerpt from commands and replies in Andrew Plotkin’s *Shade*. After this, three texts tied as fourth most natural. One of these, interestingly, was (3), a cut-up rearrangement of commands and replies from Michael Gentry’s *Anchorhead*. This ranking suggests that a rearranged text without any cues about its reordering (except semantic ones related to causality) can still be judged highly natural. An excerpt of commands and replies from the nn implementation of *Adventure* (5) was ranked alongside this, as was (13), the full score report from *Savoir-Faire*. Two texts tied for fifth. Both were generated by nn in chronological

order; one of them (11) in the course of an interaction with *Lost One* (with commands and replies) and one of them (14) from recounting part of a session of *Adventure*.

Of the texts that were least natural, those ranked sixth through ninth, none featured events presented in chronological order. Ranked eighth, and next to last, was (4), a randomly-ordered recounting of events from *Lost One* that used time words (such as “before that” and “then”) to show the relationship between adjacent sentences. This would suggest that using time words, by itself, does not help much in making a randomly-ordered list of events sound natural. Finally, the text that was marked as least natural was (6). The annotators both gave their lowest scores to this randomly-ordered list of sentences, without time words, narrating events from a session of *Adventure*.

	<i>Ann</i>	<i>Score</i>	<i>Events</i>	<i>Essential</i>
1	1st	0.83	123__45	Chron.
	2nd	0.33	????????	?
2	1st	0.83	12344_55__667	Chron.
	2nd	1.00	2__44566?1__8	Chron.
3	1st	0.50	_8____7_7345_621	Achron.
	2nd	0.33	4?443184_55567_?	Achron.
4	1st	0.17	785634412	Retro.
	2nd	0.00	231???????	Retro.
5	1st	0.83	12345__6	Chron.
	2nd	0.00	234561?_	Chron.
6	1st	0.00	????????_???	?
	2nd	0.00	??????????????	?
7	1st	0.50	77_888_92222_334_5_61222222	Retro.
	2nd	0.00	ABBCCC_?23_?4556788911_1111	Retro.
8	1st	0.33	8677455321_	Retro.
	2nd	0.00	5444333211?	Retro.
9	1st	1.00	????????	?
	2nd	0.33	????????	?
10	1st	0.33	7654231?	Retro.
	2nd	0.00	76543211	Retro.
11	1st	0.67	1234456789ABCDEFGG?H	Chrono.
	2nd	0.00	111223456789ABCDDDE	Chrono.
12	1st	1.00	12333?????_??_?????	Achron.
	2nd	0.33	?????1234567_89AAABC	Achron.
13	1st	0.83	????????????????	?
	2nd	0.00	????????????????	?
14	1st	0.67	12_345678__9AB	Chron.
	2nd	0.00	????????_???????	?

Table 3: Whether events identified by annotators matched up and what chronological ordering was assigned to events. “_” indicates that one annotator did not mark something as an event when the other did. Numerals represent the chronological rank assigned, with ranks 10, 11, 12 ... indicated with A, B, C... The essential identification by each annotator of the prevailing order in each passage is indicated in the last column.

10.1.2 Agreement about Events

Interannotator agreement about what events were indicated in texts was reasonably high (89%, using the F-measure of pairwise comparisons), considering that very little instruction was given about what should be considered an event. Annotators agreed about 150 events (300 markings). They did not agree about 40 markings of events. The texts ranged fairly widely, from lists of accomplished tasks to straightforward recountings to more complex excerpts from IF discourses. It is not clear that all of the cases where annotators disagreed would have been easy to resolve. Nevertheless, with some refinement of instructions and with provision for discussion and consensus judgments, it should be possible to have annotators identify when events are narrated in more than 90% of the cases.

10.1.3 Ability to Communicate Events and Order

The numbers assigned to events were seldom identical and did not always even look similar. There were only two cases, however, where annotators disagreed about the essential overall ordering of events. In both of these cases, familiarity with interactive fiction led the first annotator to correctly identify the texts as narrating chronologically, while the second annotator could not tell that they were in any particular order. The second annotator also marked more events in both of these cases, making the process of determining the order more difficult.

The annotators made different assumptions about a list of randomly-ordered sentences with the time words “before that” and “then.” The first assumed that these words were relative to the entire discourse so far (that is, that “before that” meant “before everything that has been narrated so far”) while the second thought they referred only to whatever single event had been narrated just previously.

10.1.4 Comments from Annotators

The annotators were first instructed to evaluate naturalness and were told that “naturalness means being written in ordinary, fluent English. The more awkward and disfluent the text is, the less natural it is.” They indicated that they appreciated less repetitive prose. The second

annotator said that the phrase “As your fingers touch ...” in (2) offered a welcome relief from a formulaic way of reporting that the “take vase” command was successful, such as “The vase was touched.” Remarkably on (1), that annotator said that “words like ‘Taken’ are extremely irritating. Regarding the *Adventure* transcript, (5), that annotator said of the repeated statements about walking and looking that “there’s so much redundancy here” and “explaining the perceptual act of the agent seems a little odd.” The first annotator found the presence of causally connected events important, noting about the boy throwing the ball and the ball falling to the ground in (11), “it helps you see that ... the whole thing is in chronological order.” Both annotators agreed that recountings of events would benefit from variety and from better selection of salient events. Both also found the retrograde narratives particularly difficult and unpleasant to read.

10.2 *Lessons for IF Output Evaluation*

There are no real conclusions to be reached about the quality of the system’s text generation ability as compared to anything else based on this pilot evaluation, which was not mainly undertaken to assess the system. The goal of this session was to learn more about how people are able to annotate, initially with minimal instructions and in a short time, so that a more focused evaluation using better instructions, an opportunity to adjudicate among the annotators, and a more appropriate corpus of texts can be done in the future. There are several conclusions about how to have texts annotated in the future:

Interactive IF sessions are not necessary. Annotators can judge the quality of generated text and can determine what events are narrated without typing input to an IF system. Whether the full interface, the full interactive situation, and the simulated world works well are different questions, and question for a different sort of evaluation.

Context should not be omitted. Annotators should be given a complete text: a transcript or recounting that can stand alone. A very short transcript is ideal, but all transcripts should begin with the system’s initial output, or at the beginning of a hypodiegetic simulation or some other major discourse boundary, not in the middle of a transcript. To facilitate this pilot evaluation, short excerpts were taken out of a larger context of interaction and printed out. This proved problematic in a few cases because important context or even something that would explain an offhand reference (and thus make the text seem more cohesive and natural) was missing.

Jumbled and “real” IF texts are not necessary. For a complete evaluation, it seemed that a baseline of some sort and something like human-authored “gold standard” output would be ideal. However, cut-up texts were not judged as being much less natural than were the unjumbled texts, and in one case, both annotators mistook an unaltered transcript for one that had been cut up. Stylistic differences and the presence of features that were not being evaluated (such as the presence or absence of direct discourse) seemed like they may have overwhelmed the questions being asked about the quality of the system’s output. The human-authored IF outputs were arranged in fairly large paragraphs with all sorts of stative information, some use of modal verbs, and other grammatical complexities. Instead of adding such “control” texts to the pool being assessed, it seems more useful to vary a single parameter (such as the way that anachronies are indicated: using tense, the same time words, a mix of different synonymous time words, or a mix of all of these) while presenting a set of system-generated texts that are otherwise the same. A large group of annotators could be used, with each given a small subset of texts. This will not reveal how close such output is to human-authored text. Comparing different system-generated texts could still help to determine whether, for instance, using a random mix of synonymous time words in the appropriate places is better than repeating the same ones, and whether mixing in some changes in tense is even better. Then the actual full interactive system, implemented using the best of these techniques that were found, could possibly be evaluated against works done in state-of-the-art IF systems.

Knowledge of IF conventions matters. The annotator familiar with IF was able to resolve the chronology of more of the excerpts, and in discussion at the end of the session, it was clear that there was an important difference in the two annotators’ ability to make sense of the excerpts. If the target readership for work in the system includes both the IF-savvy and those not familiar with IF, it makes sense to have annotators of both sorts judge the texts. However, it will also be important to record annotators’ IF experience, probably with more granularity than “familiar/unfamiliar.” The differences between the assessments of the two groups might represent sticking points for newcomers that should be carefully attended to; alternatively, the less-familiar group might simply have a harder time understanding the interaction that produced the transcript.

11 Discussion

11.1 Consideration of the Project and Early nn Development

The project described here has been successful in contributing several narratological advances and making progress in text generation. It has served as the first stage of a broader aesthetic and educational project. The contributions that have been made are reviewed in the next section. In this section, some of the challenges faced during the project and some of the limitations of the chosen approach are discussed.

The major limitation of the work done so far is that it does not include the development of more original creative work, the porting and modification of work done in previous systems, and the use and evaluation of the system by IF authors. These deficiencies are not just problems from an aesthetic, literary perspective or from the practical standpoint of the IF community. The development of quality text generation in the Narrator was almost certainly hindered because there was less to guide the development of the theory and of the system. A few IF authors were given access to a preliminary version of nn during the project, but the state of the code, the lack of documentation, and the limited amount of time available did not allow for any original work to be done by these IF authors. In this first preview, IF authors were only able to discuss the basic concepts of the system. This lack of creative development is not a complete surprise, given that this was a research project in computer and information science (a discipline which favors quantitative evaluation) with relevance to narratology (a discipline favoring the development of theoretical models). However, the absence of work does make it difficult to do more than speculate about the difficulties that IF authors may have in implementing plans for narrating, code to update the plans for narrating, and actors, and in writing strings-with-slots to define appearances and the representation of events. The strings-with-slots formalism may be too difficult to use, or too limited in its ability to represent syntax, or both; only when IF authors attempt to create games will it be possible to figure out more about this.

As development of nn continues, the next project should embrace creative, literary development and use of the system by IF authors. The potential of this computationally

inexpensive interactive fiction system, one which is unencumbered by modules that require license payments, should be used to bring this line of research out of the lab and into the broader creative community of gaming and literary work. Interested IF authors should be invited to join as collaborators in creating new versions of the system that respond to the needs of developers and that incorporate new features that are desired.

Because of the project's focus on the Narrator and text generation, the overall maintenance of the discourse — done by the Recognizer as well as the Narrator, using the Discourse Model — was not considered in adequate depth. Referring expression generation and resolution should be developed further throughout all the relevant modules and models. Also, the system should be enhanced so that all words generated by the system in descriptive text are automatically understood by the Recognizer.

There are also limits to the basic approach of specifying behaviors and information about the telling on objects (things, actors, and rooms). On the behavior side, this distribution of code inclines IF authors toward autonomous, bot-like actors rather than a tight cast of characters that is centrally directed. Drama management would not be impossible in nn, and could be accomplished by adding code to COSMOS, but the framework does not encourage it. On the side of descriptive text, that this text resides on objects rather than on the perceiving actors also biases the system a certain way. As a practical matter, the things being perceived tend to be modified more than the perceivers, so it is less trouble to store the string-with-slots representing appearance on the existent being seen. However, while this also does not prohibit variation in appearance, it does encourage IF authors to write games in which all narrators "see" existents in the same way. It would be worthwhile to explore other paradigms besides the object-oriented programming one that is typical in state-of-the-art IF development systems.

It is hoped that the work on narrative variation begun in this project will be significant for the future of interactive fiction. This project will almost certainly not be the only important type of advance made in the next few decades, though. Those working with interactive fiction could find an effective way to use cell phones as a platform (via SMS or voice interfaces), could integrate the textual exchange with graphics in new ways, could craft characters and other aspects of the world so that they are much more fine-grained, or could simply take on new literary projects that had previously been unimagined — all without making use of narrative

variation. The development of general facilities for varying the telling, however, will provide new opportunities for work that draws on literary traditions. This has the potential to refresh interactive fiction and perhaps computer gaming more generally, and to connect this work to literary art in a unique way.

11.2 Summary of Advances

Significant progress has been made in this project toward the goal of automatic narration. This includes theoretical work in relating narratology to IF and in formalizing narratology, architectural advances and the implementation of an IF system that uses this framework, and the composition of complex ways of narrating in creative pieces. Summaries of each of the major contributions are provided below.

11.2.1 A Theory of IF Based on Narratology

A detailed theoretical framework for understanding interactive fiction using an approach based on narratology has been developed (chapters 2 and 4). This includes a mapping of IF to narrative concepts at a high level and element by element. The analysis takes into account what is different about interactive fiction and that an IF work is an interactive computer program and a potential narrative rather than a narrative. This work also includes an analysis of IF outputs which has resulted in a detailed typology of these outputs.

11.2.2 A New IF Architecture Abstracting Content from Expression

A new architecture for interactive fiction has been developed, using insights from narratology about the distinction between content and expression in narratives (chapter 5).

11.2.3 A Working IF System, nn

The proposed architecture has been implemented in Python in an IF development system called *nn*. Adequate Recognizer, Simulator, Joker, Preparer, Clarifier, and Presenter modules have been written to make *nn* a complete and usable interactive system which can be accessed locally or remotely by interactors. As a result, research focusing on the Narrator can be conducted in the full context of a working IF system.

11.2.4 Formalization of Narrative Variation

Possible variations in the narrative discourse, from the perspective of an author or a text-generating system, have been formalized so that they can be used in a natural language generation system. Formal models have been provided to represent types of variation in each of the five major categories of Genette's *Narrative Discourse*: order, speed, frequency, mood, and voice (chapters 6 and 7).

Variations in order have been formally described by defining algorithms to sort events into chronicles and retrograde narratives, to disarrange them in achrony, and to categorize events using zigzag and sylleptic schemes. The inputs necessary to specify analepsis and prolepsis are also defined, with algorithms to add those types of anachronies also provided. These models of order are combined with the concept of time of narrating, which Genette discusses as part of narrative voice, in a reply plan that is an ordered tree and that encapsulates the event time, reference time, and speech time of everything that is to be expressed, the order of these proposed expressions relative to one another, and the way that they are embedded within one another. From Genette's description of how one sequence (the order in which events are narrated) relates to another sequence (the chronological order in which these events transpired), this enriched model using ordered trees and temporal information has been developed and the algorithms needed to produce different types of order have been specified.

Genette's other categories of narrative discourse have also been explored, and some elements from each have been further formalized and implemented. Speed has been implemented as a straightforward length parameter, with the requirement that any event that is to go unmentioned (that is, to be related at infinite speed or with a text of length 0) must be dropped during content selection. The different frequencies identified by Genette (singulative, repetitive, and iterative) have been mapped to different sorts of proposed expressions (ordinary, duplicated, and grouped for later aggregation). A rather extensive model of actors' perceptions and knowledge has been developed and implemented to allow variation in focalization, an aspect of narrative mood. This model includes a history so that it can roll back to a point in the past when it is necessary to narrate a past event from a particular actor's perspective.

11.2.5 An Automatic Narrator Based on Genette's Theories

A Narrator has been developed to perform some sorts of narrative variation (chapters 6 and 7). The focus has been on variations in order and how these relate to the time of the narrating. The Narrator is a three-stage pipelined text generation system which maps the high-level IF world and plan for narrating through Reply Structures and syntactical representations to output text. The appropriate tasks for each stage, the interfaces between each stage, and the workings of each stage were all determined during the process of implementation.

11.2.6 Implementation of Standard IF

Two “benchmark” interactive fiction pieces, *Cloak of Darkness* and *Adventure*, have been coded in nn (appendices C and D) and plans for narrating have been developed which are compatible with both (appendix E). These serve to show that the system can implement standard IF in a reasonable way. The pieces also demonstrate what different plans for narrating can do when applied to the same known IF world.

11.2.7 Implementation of New Pieces with Complex Plans for Narrating

Early versions of two original works of interactive fiction — *Lost One* and *Oddly Angled* — have been developed in nn to guide system development and for use in evaluation (chapter 9). Although not yet interesting as literary works and games in their own right, they show how more complex narrative effects can be built up from the fundamental elements of the plan for narrating.

11.2.8 Pilot Evaluation

A pilot evaluation has been conducted (chapter 10). The evaluation can inform a more focused annotation involving larger numbers of IF output texts. It indicates that transcripts will work as texts to be annotated, but that they should start from the beginning; that slightly differing outputs from the system will be best to compare against one another; and that the annotator's awareness of IF conventions should be noted.

11.3 *Future Work*

The research done here and the development of the IF system *nn* can contribute to a great deal of further work in digital media, narratology, artificial intelligence, and natural language generation. This final section describes several additional projects in these disciplines that could be achieved with the right support and with efforts from the right team of collaborators. They have been roughly classified as small (S), semester-long projects that could be done mainly by one person; medium (M), projects that would involve a team on a shorter timeline or an individual working over a period of a year or so; large (L), multi-year projects that would involve collaboration; and extra-large (XL), projects that would require significant effort from collaborators in different disciplines over more than a year. The projects are organized based on what seems to be the primary field in which they lie, although several projects are situated in more than one.

11.3.1 *Digital Media / Digital Writing*

- **A public release of *nn* for IF authors (M-L).** This will involve completing system development and documentation to the point where *nn* can be publicly released as free, open source software, so that anyone who has Python installed can develop and run interactive fiction using this system. The modules of the system other than the Narrator will particularly have to be further specified and developed — this project requires that at least some significant work has been done on the next two projects. The system will also have to be tested for usability and functionality. This is a software development project in digital media rather than a research project in computer science, but one which would be very useful in broadening the use of the system and facilitating further creative development and more widespread research. Continuing the development of *nn* and supporting it will be an ongoing effort, but a public release could be accomplished by a team in a fairly short time, perhaps in time for next year's IF Competition and for use in teaching the following academic year.
- **New IF collaborations between world-builders and narrator-builders (M).** Such collaborations, where one person creates the IF world and the other determines how it will be narrated, would have results that are mainly interesting on aesthetic grounds. They

would have a side benefit of demonstrating the generality of the Narrator, and could supply examples for narratological study.

- **Bringing other aesthetic text generation work into nn (S).** This involves incorporating the results of ongoing research on narrating and aesthetic text generation, when they are available, into the system. For instance, support could be added for multiple narrative/simulative levels and deeper schemes for referring to the narrating itself, or capability for metaphorical expression could be incorporated into the Narrator. Such work can be used to test the applicability of these techniques to IF; in adapting them for the IF situation, more will certainly be learned about how to use and improve these techniques. There are several generation techniques that could be integrated into nn; each of these considered individually might be a small-scale project.

11.3.2 Narratology

- **Adding an awareness of and ability to discuss the narrating (M).** Currently, nn can generate many different sorts of narrative text referring to the same events. But it does not have the ability to model and refer to the narrating itself by producing such texts as *this happened more quickly than it can be told* and *the exciting way in which I am relating these occurrences may conceal that fact that they were very tedious*. (Producing the second of these examples, of course, would require some sort of emotional model as well.) Being able to refer to the narrating in a wide variety of ways will be difficult, but a rudimentary ability to do this could probably be added in a medium-scale project.
- **Developing nn as a tool for teaching about narrative (L).** Making the system suitable for IF authors is one important direction, but students of narrative theory can also make use of the system. They will have different needs; the ability to see the same set of events played out again and again using different plans for narrating is the most important of these. A rich and versatile system for interacting with a simulated world may be unnecessary for this use. Even a non-interactive program that takes existents, events, and a plan for narrating as input and that presents a narrative could be helpful. Determining the best interface and the appropriate range of capabilities would be the first step in this project.
- **Adapting automatic narration for other applications (L).** This entails moving the ability to

narrate from this particular aesthetic application, interactive fiction, into other systems, including speech systems. Mobile robots and speech interfaces to personal computers, both of which should model and retell events, provide two examples of possible new application areas. Implementing a modified narrating capability in other systems should be undertaken in collaboration with those who have expertise in such system.

11.3.3 *Dialog Systems / Text Generation / HCI*

- **Improving the Recognizer (S-M).** A pronoun resolution capability is particularly important for allowing more English-like, natural input. Other improvements in parsing (such as the ability to use multiple direct objects with verbs) would also be of use. An ideal project would incorporate these improvements in the context of computational linguistics research. New recognition techniques could be tested in nn, particularly if it is useful to investigate how such techniques work in the semantically rich but usefully bounded framework of an IF world. As a smaller-scale project, multiple existing pronoun resolution techniques could be implemented and could be compared against one another in IF sessions.
- **Improving the Clarifier (S-M).** In a publicly-available system, this module should at least ask questions to disambiguate commands. It could be improved in several other ways. A Clarifier based on HCI and dialog system principles would be ideal as part of the system. nn also could allow researchers the opportunity to implement and compare different clarification techniques in otherwise-similar IF systems.
- **Adding generation from abstract syntax representations (M).** The string-with-slots representation is easier for IF authors to understand, but a more abstract representation for texts in the generation process would allow for finer-grained manipulations. It would facilitate multi-lingual generation as well. While the outcome of this project would probably not be useful for IF authors or narratology students, it could contribute to progress in computational linguistics.
- **Tools for building strings-with-slots (S).** Some software tools to help IF authors generate the Narrator's strings-with-slots representation would be helpful. It may be possible to semi-automatically create these representations from ordinary texts under user

supervision.

- **Providing multi-lingual text generation and corresponding recognizers (L).** This would allow the same interactive fiction to accept input and produce output in multiple languages. If successful, this would demonstrate the generality of the Narrator to different languages. This would require extensive development of grammars, new development in the Recognizer module, and research staff with expertise in different languages, and it is not a core investigation of the issues explored by this project.
- **Modeling the user (M).** This would involve predicting the interactor's current activity or goal and employing a model of the user to improve recognition and narration. Hand-crafted, limited sorts of user models have been implemented in existing IF, but the general principles for how a user model applies to the IF situation have not been well-developed. Research in this area might determine better ways to narrate for different users, and could, of course, have benefits that extend beyond IF.

11.3.4 *Artificial Intelligence*

- **Adding planning for actors (L).** The focalizer words have the basics of what is needed for planning, but the modality of planning an action is not the same as the modality of narrating, and more would have to be done to allow actors to plan based on their own knowledge. A capable system could be developed based on previous work on the relationship between knowledge and action (Moore 1980). While entailed actions would need to be narrated in a slightly different way — specifically, the interactor would need to be informed of what the player character is doing — almost all of the work on this planning capability would be in extending the Simulator and the world models. Questions related to believability and lifelike behavior could be tested in the context of this enriched planning model.
- **Adding drama management (M).** nn excels at providing actors with their own individualized views of the IF world, but there is not much of a facility provided for high-level orchestration of these actors. A drama manager (almost certainly located on or connected to the COSMOS existent) might be a particularly interesting element of an nn interactive fiction. Drama management could be interesting to experiment with given the

support for focalization and first-order representation of events, and of course given the ability to change the plan for narrating.

- **Combining narrative variation and story generation (L).** Systems such as MEXICA are advancing the state of the art of story generation by automatically determining what events are interesting to tell. It would be ideal to combine this event-selection and event-sequencing ability with an ability to arrange these events in the telling in a variety of interesting ways.
- **Developing subjectivity (XL).** This includes developing both characters and narrators with some sort of subjectivity, providing for variations in character behavior and narrator mood and voice. This extension would improve the Narrator, but it would involve the development or use of very rich, elaborate mental models of character and narrator. To accomplish this more elaborate subjectivity, the Simulator, the world models, and the Discourse Model would have to be extensively redeveloped to broaden the simulation beyond the physical world to include characters' mental states. Of course, *ad hoc* forms of subjectivity can be developed as games are written, and authors who write in nn will inevitably impart some sorts of subjectivity to characters and narrators. But establishing a general framework in which different sorts of subjectivity can be obtained parametrically will not be easy.

These are directions that seem the most interesting; there are still others that developers and researchers may choose to take. nn could be made into a multiplayer system or support for multimedia elements could be added, for instance. Once the system is released, people will be free to develop or even fork it in whatever ways they wish and these sorts of possibilities can be explored as well. The potential for practical and research use should allow for many sorts of new creative work, new directions of IF system development, and new research that builds on what has been accomplished here.

12 Glossary

Definitions relating to interactive fiction specifically are mainly based on those given in chapter 4, while definitions from narratology are based on those in A Dictionary of Narratology, 2nd edition (Prince 2003), which can be consulted for more complete definitions and for references to the original works in which these terms were defined. Only the most important sense in which terms are used in this document is indicated.

achrony. An order in which events are narrated without the obvious indication of their relationship to one another in time.

action. A sequence of events and actions which is connected by a common intention.

actor. In interactive fiction, an existent in the simulated world capable of undertaking actions, as distinct from things (which are inert) and rooms (which represent locations).

anachrony. The telling of an event out of chronological order.

analepsis. A flashback or retroversion in which events that are previous to the main sequence are related.

baseball narration. The style of narrative in which baseball commentators tell what is happening in a game. This style can involve a simultaneous, present-tense narration as the play is happening followed by one or more retellings of the same events in the past tense.

character. In interactive fiction, an actor who is anthropomorphic.

chronicle. A narrative in which events are related in chronological order.

clarification. An output from an IF system that serves to help disambiguate an input or that simply explains that it cannot be understood. For example, "Are you referring to the first Starbucks or the second Starbucks?"

command. An input to an IF work that is diegetic and instructs the commanded character, within the simulation, to undertake a particular action. For example, "GET THE LAMP" or "HITCHHIKE".

commanded character (CC). The character who is directly commanded by input from the interactor. This function is almost always combined with that of focalized character;

together the two functions are referred to as player character.

content. The plane of the events and existents of a narrative or of a simulation; what happens as opposed to how it is told.

course. A series of interactive fiction exchanges (commands and replies) that are consecutive in terms of the simulation.

cycle. In interactive fiction, any input along with the output that follows it. If the input is a command, that command and the reply that follows are an exchange.

diegesis. The level of narration or simulation, as opposed to the level dealing with the text or program itself.

directive. An input to an IF work that is extradiegetic and refers to the program or the "game" itself, not the simulation. "RESTART" and "SAVE GAME" are examples.

discourse. In narratology, the expression plane. Because discourse is used in computational linguistics to emphasize other things, "expression" is used in this document for this sense of "discourse."

distance. The metaphorical space between the narrator, the events and existents (including characters), and the narratee.

event. A change in the state of the fictional or, in this case, IF world, as opposed to something that exists in that world and is part of the state. Some events result in no underlying change but are registered because they need to be narrated. Actions are used to group events together based on intention.

exchange. In interactive fiction, a command followed by the reply that is output after it. The exchange is a specific type of cycle. The text of an exchange corresponds to what is often called a "turn."

existent. Something that exists in a fictional or IF world, as opposed to an event, which is a change in the state of the world. In the IF system described in this document, an existent is either a thing, actor, or room.

expression. The plane of telling or narration; how events and existents are represented rather than what they are.

extradiegesis. The level above the narration or simulation, which deals with the text or program rather than the narrated or simulated events.

final reply. The last diegetic output from an IF work, ending a traversal and often indicating that the interactor has won or lost.

final situation. The state of the simulated IF world at the point of the final reply.

focalization. The perspective from which a narrative is told. When a particular character focalizes the narrative, there is internal focalization; when the narrator does not restrict the telling to what one character knows, there can be zero focalization.

focalized character (FC). The actor who is being focalized. This role is almost always combined with the function of commanded character to constitute the function of player character.

Focalizer World. A model representing the perceptions and knowledge of a particular actor, relative to an IF Actual World.

frequency (narrative tense). The relationship between the number of times an event happens and the number of times it is narrated.

game. A term often used for any work of interactive fiction, whether or not it is a game according to definitions advanced in game studies.

heterodiegetic narrator. A narrator who is not a character in the story he or she tells, and who therefore stands apart from the diegesis.

homodiegetic narrator. A narrator who is a character in the story he or she tells, and who therefore is part of the diegesis.

hypodiegesis. A level of narration or simulation within the “frame story,” as when a character in the narrative tells his or her own nested narrative.

IF Actual World. The re-centered actual world of an IF work. It is a fictional simulation, but it is “actual” to the characters within it; their actions take place within this world and their own perceptions and knowledge are represented in other models of this world, called Focalizer Worlds.

IF world. The model world of an IF work, including existents, rules for action, and physical laws. In traditional IF systems, there is a single IF world modeled; the system described in this document divides the model into an IF Actual World and a Focalizer World for each actor.

initial situation. The state of the IF world at the beginning of a session.

input. In text-based IF, the text provided by the interactor to the system, which may be a command, a directive, or some unrecognized input.

intended action. The system's interpretation of a command as something particular which the interactor wishes the commanded character to undertake.

interaction. A single "playing" of an IF work that is continuous from the interactor's perspective.

interactive fiction. A form of text-accepting, text-generating computer program that narrates what is happening in a simulated world in reply to input from a user, or interactor. Interactive fiction can have literary qualities and qualities of a game.

interactor. The user of an IF work, who provides input and reads the output.

iterative (frequency). "n to 1" narration, in which a group of similar events are told in a single narration. For example, if a character drove to work on Monday, Tuesday, and Thursday, "She drove to work a few times that week."

mood (narrative). The set of modalities regulating narrative information. Specifically, these modalities are distance and focalization.

narratee. The one to whom the telling is addressed; the "you." The narratee is a function of the text, not the same entity as the reader and not on the same level.

narrated. The set of situations and events recounted in the narrative; the story or content as opposed to the discourse or expression.

narrating. The telling, as opposed to the narrated, which is what is told.

narrative. The representation of one or more events by one or more narrators to one or more narratees. The narrators and narratees may be more or less overt. Narratives are distinguished from texts that contain only propositions and description. Discussion continues about whether the definition of narrative should include such cases as dramatic presentations (which seem to lack a narrator) and instruction booklets (which indicate what to do but may not represent real or fictive events). These questions do not bear very directly on the project of narrative variation in interactive fiction, however.

narrative variation. In this document, the term is used to indicate a capability to change the narrative discourse independently of underlying events and existents. That is, the term indicates variation in expression rather than at the underlying content level.

narratology. The study of the nature, form, and functioning of narrative. Narratology studies fictional and non-fictional narrative, literary and non-literary narrative, narratives in all media, and the distinction between narratives and non-narratives.

narrator. The one who does the telling; the "I." The narrator is a function of the text, not the same entity as the author and not on the same level.

non-player character. A character in interactive fiction who is simulated but is not commanded by the interactor and is not focalized in the narration.

omniscient narrator. A narrator who, rather than being restricted to the perspective of a single character, knows essentially everything about the events being narrated.

order (narrative tense). The relationship between the chronological sequence of events as they occurred and the sequence in which they are told.

other person. In interactive fiction, some person mentioned in the output text is not simulated, that is, who does not have an underlying representation in the IF world.

output. In text-based IF, the text provided by the system to the interactor, which may be a reply about the IF world, a report at the program level, or a clarification.

player character (PC). The character in an IF work who focalizes the narrating and is commanded within the simulation.

potential narrative. Refers to a system (such as a computer program) which is not itself a narrative text but which can generate narrative.

prolepsis. A flashforward or anticipation in which events that happen after main sequence are related.

prologue. A representation of the IF world, possibly including description of states and narration of events, which is provided before the first opportunity for input.

puzzle. A challenge in an IF work that can only be met with a non-obvious set of commands.

repetitive (frequency). "1 to n" narration, in which an event that happens once is narrated multiple times. For example, if a character dropped the ball once: "He dropped the ball. He dropped the ball! He dropped the ball!"

reply. An output that follows input from the interactor and represents anything about the IF world and events in it. For example, "It's too dark to see, like inside of a dog."

report. Outputs that are at the game or program level and do not deal with the IF world. For example, "Abbreviations on" or "Save failed."

retrograde. The telling of events in reverse chronological order.

room. In interactive fiction, an existent in the simulated world which represents a location, as

distinct from things (which are inert) and actors (which can undertake actions). Both outdoor and indoor locations are called “rooms.”

session. A complete execution of an IF program from when the program first begins executing to when it terminates.

singulative (frequency). “1 to 1” or “n to n” narration, in which each event is narrated once. This is the standard case.

simulation. A computer model which describes the state of a system and the system’s dynamics.

It may be an abstraction of something in the real world or not. An IF world is a simulation.

situation. A state of the IF world.

speed (narrative tense). The relationship between the time it takes for an event to occur and the time it takes to narrate it. In a textual narrative, the time it takes to narrate is measured by considering the amount of text devoted to that event.

story. In narratology, the content plane. Because this term has a variety of meanings, “content” is used in this document for this sense of “story.”

successful traversal. A traversal in which the interactor wins, reaching the best final reply.

syllipsis. The organization of events in the narrating by category rather than by chronology.

tense (narrative). The temporal relations between the events narrated and the narrating of them.

This includes order, speed, and frequency.

thing. In interactive fiction, an inert existent in the simulated world, as distinct from actors (which can undertake actions) and rooms (which represent locations).

time of narrating. The temporal position of the narrator relative to the events. Usually the narrator’s position is subsequent to the events, corresponding to a mainly past-tense narrative.

traversal. Some “completion” of an IF work, not necessarily successful, beginning with the initial situation and prologue and ending with some final situation and final reply.

unrecognized input. An input that is not recognized by the IF work as either a command or a directive.

voice (narrative). The signs characterizing the narrator, the relationship between the narrating and the narrative (that is, the text) and the relationship between the narrating and the narrated events.

work. A generic term for any piece of interactive fiction. "Game" is often used instead, although that term implies that the work can be won, which is not always the case.

zero focalization. Focalization using an omniscient narrator rather than via a single character.

zigzag. An ordering of events in which the movement of time is back to the the past or "then," up to the present or "now," back to the past, up to the present, and so on. This can be used when a character is in a situation similar to one in the past, for instance.

13 Appendix A

Example Interactive Fiction Transcript

From Adam Cadre' *Varicella* (1999)

Lines beginning with ">" were typed by the interactor. Boldface in original.

>go southeast

Palace Asylum

The walls and floor are padded here, a cylindrical pattern of white squares spiralling upward to a skylight far overhead. The room is exactly half as wide as it is tall, which at this latitude means that the sun appears in the skylight on one day and one day only: the first day of summer. The door, being padded exactly like the walls, is scarcely discernible, but you can sense its outline to the northwest.

Princess Charlotte, tightly bound in a straitjacket bejeweled with a mad assortment of silver buckles, sits propped up against the wall opposite the door. Youngest daughter of the King of Paris and adoptive sister of Queen Sarah of Piedmont, Princess Charlotte was betrothed to your younger brother Terzio, with the wedding scheduled to take place upon her sixteenth birthday. But as they exchanged their vows in the chapel just downstairs from this very tower, a pair of assassins shot Terzio dead, splashing the young princess's wedding dress with the blood of her almost-husband. She let out a great wail as the assassins fled, and her keening continued unabated for close to an hour; finally she was sedated, placed in a straitjacket, and taken to the tower. That was more than four years ago. She's still here.

She looks up. "hello, varicella," she says. "face it, tiger, you just hit the jackpot!"

>examine charlotte

As the nominal head of the House of Varicella, you accompanied Terzio to Paris when he went to ask the Parisian king for his daughter Charlotte's hand in marriage. This was the first time you ever saw the object of his affection, as she emerged from the royal woodlands, flowers strewn through her flame-colored hair, her skirts tattered by her unaccountably joyous dash through the forest -- a daily ritual of hers from the time she was barely old enough to walk, you were later to learn. The picture of health, she, as she bounded out of the woods and across the lawn, cheeks flushed with the exertion of a breakneck run and with young love, merrily racing toward the storybook life that awaited her in Piedmont...

Now she is frail as a great-grandmother, her skin milk-white from her years locked in this tower... but more than anything, you are struck by her eyes. Dazzling green like Miss Sierra's, they are far from hollow or dead as one might expect -- but neither are they anything close to lucid. Rather, their intensity is cranked up beyond anything you could possibly construe as human; to Princess Charlotte, you reflect, the world must look as if it's all been shot on a roll of overexposed film.

>ask charlotte about the asylum

"Are you displeased with your current accommodations?" you ask.

"well, it's not like i'm asking for a nice room with a comfy bed like i had back home," Princess Charlotte says. "all i ask is that you take away some of the padding so that when it gets intolerable i can bang my head on the wall till i pass out. is that so much to ask?"

>ask charlotte about prince charles
"Can you tell me a bit about Prince Charles?" you ask.

"he comes to visit me sometimes," Princess Charlotte says. "sometimes he throws things at me and sometimes he calls me names and sometimes he talks about how his teacher does things to him in the day like louis and rico do to me at night and sometimes he brings his car and hits me with it and sometimes he just looks at me and then leaves."

>ask charlotte about night
"Can you tell me a bit about night watchman?" you ask.

"it's the same thing pretty much every night," Princess Charlotte says. "the sun goes down, the guards change shifts, and some time later, the night watchman lets the bad men in. sometimes prince louis and sometimes argento rico and sometimes both. and they do things to me that are not as fun as i had been led to believe, let me tell you. but the worst thing is that after staring at a blank wall for twenty-three hours with half of it in the dark sometimes the pretty colors come and there are other times and other places and then the door opens and the bad men come in and they drag me back to the bad room where all i can see are white squares. well, some red squares when they're done."

>unbuckle princess charlotte
You unfasten the straps securing the Princess's straitjacket, and she gleefully hops to her feet. "mein fuhrer, i can walk!" she cries, shaking her arms to bring back the circulation. "so what are we going to do today, varicella?" she asks. "take over the world? won't that be fun!"

>leave
You'll have to say which compass direction to go in.

>northwest

Southeast Tower, Top Level

You are standing at the top of the massive marble spiral staircase that comprises the southeast tower. Before you, to the southeast, is the door to the palace asylum.

A guard stands watch over the entrance.

Princess Charlotte follows in your wake.

"What on earth?" the guard says. "The patient is not to leave the asylum under any circum--"

"so long, sucker!" Princess Charlotte interrupts. "i'm taking this thing to mexico!" She grabs your sleeve and pulls you down the stairs.

14 Appendix B

Example Interactive Fiction Transcript

From Emily Short's *Bronze* (2006)

Lines beginning with ">" were typed by the interactor. Boldface in original.

Law Library

Many books of precedent line these walls, containing every kind of contract that can be made to bind every kind of soul.

A hole in the floor descends to the other, less savory portion of this place.

You can see a great contract book here.

>examine book

The runes are unfamiliar to you, but you know what the book is: a record of all the contracts of all the souls enslaved to the king of this castle.

You caught him staring at you once. "Your clothing is wearing out. I'll look up a seamstress in the contract book for you."

You plucked the erring sleeve back into place. "You needn't," you said. "I don't mind."

"Yes, but I do," he snapped. "I was once a -- the polite term would be a connoisseur of ladies -- and it is not a taste that goes away. So for everyone's sanity it would be best if you went about fully clothed."

You avoided him for three days, after that incident. But your gowns were all replaced.

>north

Lower Bulb

In this very tall room, like a silo, is a glass of running sand: not an hourglass, or even a dayglass, but a timer whose duration you do not know. A whole Sahara has poured into its lower chamber, but the trickle from above continues, very fine.

Around the outside of this contraption ascends a wooden staircase.

>west

Translation Room

Endowed with a deep power that dissolves into ready vernacular even the most ancient and secret scripts of the Law Library.

Usually blazing with the captured light of many thousands of candles, but it has been let to go out, and everything is dim.

And here Beast lies, sprawled on the ground as if he'd fallen.

You can also see an ordinary quill pen here.

"Nothing bad will happen for the first seven days," he said, when you left.

And yet here he is, looking very nearly dead.

>wake beast

Though you shake him vigorously, he does not stir.

>take pen

It is too heavy for you to lift.

>east

Lower Bulb

In this very tall room, like a silo, is a glass of running sand: not an hourglass, or even a dayglass, but a timer whose duration you do not know. A whole Sahara has poured into its lower chamber, but the trickle from above continues, very fine.

Around the outside of this contraption ascends a wooden staircase.

>south

Law Library

Many books of precedent line these walls, containing every kind of contract that can be made to bind every kind of soul.

A hole in the floor descends to the other, less savory portion of this place.

You can see a great contract book here.

>down

Darkness

Though a dim light filters down from the law library, you can see almost nothing of the contents of your current location. You find yourself concentrating all the more alertly on your hearing, as though the slightest echo might offer a clue.

You hear some dry sifting from the north.

15 Appendix C

Example nn Code

Cloak of Darkness

```
"""Cloak of Darkness

An implementation of Roger Firth's Cloak of Darkness (1999)
in nn, Nick Montfort's research IF development system.
"""
__author__ = 'Nick Montfort <nickm@nickm.com>'
__version__ = '0.1'

from Models.discourse import Message, command_lists, compass, main_directions,
directives, relation_names, text_to_relation, state_names
from Models.world import Actor, Thing, Room
from Simulator.events import *
from Narrator.microplanner import Paragraph
from copy import deepcopy

#- DISCOURSE -#
#-----#

frontmatter = (
    'Cloak of Darkness',
    'A basic IF demonstration',
    'Roger Firth - implementation in nn by Nick Montfort',
    "Hurrying through the rainswept November night, you're glad to see the
    bright lights of the Opera House. It's surprising that there aren't more people
    about but, hey, what do you expect in a cheap demo game...?"
)

message = Message()
```

```

#- WORLD -#
#-----#

class Scrawled_Message(Thing):
    def __init__(self, tag, **keywords):
        Thing.__init__(self, tag, **keywords)
    def react(self, world, e):
        new_actions = []
        if e.get(0) == 'IMPEL' and not e.get(3) == 'NORTH' and \
            world.ex[self.room(world)].light(world.ex) < 0.4:
            modify1 = Modify_Event('OPERAGOER', str(self), 'INTACT', \
                True, False, e.timestamp, e.id)
            modify2 = Modify_Event('OPERAGOER', str(self), 'APPEARANCE', \
                self.state['APPEARANCE'], \
                'the message, now little but a cipher of trampled sawdust, which
V_seem_S to read ... *** S_? V_lose_PERF ***', \
                e.timestamp, e.id)
            new_actions.append(Action([modify1, modify2], e.id))
        if e.get(0) == 'SENSE' and e.get(2) == str(self) and \
            e.get(3) == 'SIGHT':
            conclude = Conclude_Event(['OPERAGOER'], e.timestamp + 1, e.id)
            new_actions.append(Action([conclude], e.id))
        return new_actions

existents = [

    Actor( 'OPERAGOER',
        called=('the', 'operagoer'),
        parent=('IN', 'FOYER'),
        adjs=['my'],
        nouns=['self', 'myself', 'me', 'operagoer'],
        allowed=[('OF', '*'), ('ON', 'CLOAK')],
        appearance='a typically nondescript character',
        state={
            'GENDER': 'FEMALE'
        }
    ),

    Thing( 'CLOAK',
        called=('a', 'velvet', 'cloak'),
        parent=('ON', 'OPERAGOER'),
        adjs=['handsome', 'dark', 'black', 'velvet', 'satin'],

```

```

nouns=['cloak','cape'],
qualities=['CLOTHING'],
appearance=[
    'a handsome cloak, of velvet trimmed with satin, slightly spattered with
raindrops',
    'its blackness V_be_S so deep that it V_seem_S to suck light from the
room'
],
state={
    'LIGHT': -0.3
}
),

Room( 'FOYER',
called=('the', 'foyer of the opera house'),
adjs=['opera', 'splendidly', 'decorated', 'red', 'gold', 'spacious'],
nouns=['foyer', 'hall', 'house'],
appearance=[
    'S_? V_see O_? sanding in a spacious hall, splendidly decorated in red
and gold, with glittering chandeliers overhead',
    'the entrance from the street V_be_S to the north, and there V_be_P
doorways south and west'
],
exits={
    'SOUTH':'BAR',
    'WEST':'CLOAKROOM'
},
prohibited={
    'impel (operagoer operagoer north *)':
    'S_? V_have only just arrived, and besides, the weather outside V_seem_S
to be getting worse',
    'configure (operagoer cloak * * in foyer)':
    'D_THIS V_not_be_S the best place to leave a smart cloak lying around'
}
),

```

```

Thing( 'HOOK',
  called=('a', 'small', 'brass', 'small brass', 'hook'),
  parent=('PART','CLOAKROOM'),
  adjs=['small','brass'],
  nouns=['hook','peg'],
  allowed=[('ON','CLOAK')],
  qualities=['METAL'],
  appearance='a small brass hook, screwed to the wall'
),

Room( 'CLOAKROOM',
  called=('a', 'cloakroom'),
  adjs=['small','cloak'],
  nouns=['room', 'cloakroom'],
  appearance=[
    'S_? V_see that clearly, the walls of this small room were once lined
with hooks, though D_NOW only one V_remain_S',
    'the exit V_be_S a door to the east'
  ],
  exits={
    'EAST':'FOYER'
  }
),

Room( 'BAR',
  called=('a', 'foyer bar'),
  adjs=['foyer', 'rough', 'rougher', 'empty'],
  nouns=['bar'],
  appearance=[
    'the bar, much rougher than S_? would have guessed after the opulence of
the foyer to the north, and completely empty',
    'there V_seem_S to be some sort of message scrawled in the sawdust on the
floor'
  ],
  state={
    'LIGHT': 0.6
  },
  exits={
    'NORTH':'FOYER'
  }
),

```

```

    Scrawled_Message( 'MESSAGE',
        called=('a', 'scrawled message'),
        parent=('PART','BAR'),
        adjs=['scrawled','trampled'],
        nouns=['sawdust','message','floor','scrawl'],
        appearance='the message, neatly marked in the sawdust, which V_read_S ...
*** S_? V_win_PERF ***',
        state={
            'INTACT': True
        }
    )

]

pfn = {}
pfn['focalized'] = 'OPERAGOER'
pfn['commanded'] = 'OPERAGOER'
pfn['indicate_known_directions'] = True

#- OPERAGOER'S -#
#---- WORLD ----#
#-----#

exd = {}
for e in existents:
    exd[str(e)] = e

operagoer_existents = [
    deepcopy(exd['OPERAGOER']),
    deepcopy(exd['CLOAK']),
    deepcopy(exd['FOYER'])
]

operagoer_existents[2].exits['SOUTH'] = '?A'
operagoer_existents[2].exits['WEST'] = '?B'

#- FINAL WORLD SETUP -#
#-----#

focalizers = [
    ('OPERAGOER',operagoer_existents,{})
]

```

16 Appendix D

Example nn Code

Excerpted from *Adventure*

```
"""Adventure, by Will Crowther and Don Woods
```

```
This version of Adventure is unfaithful to the original in a few ways:  
The pirate stands around at End of Road and waves, the dwarf wanders  
above ground, and there is a new character, the Dungeon Master,  
provided so she can be focalized and commanded. The basics of the  
original have been implemented, though. Note also that the room  
descriptions/appearances are not done as, for instance, in Lost One.  
They are written so as to match the original descriptions closely.  
For this reason, recounting may not be as clear as it is in Lost One.
```

```
"""
```

```
__author__ = 'Nick Montfort <nickm@nickm.com>'
```

```
__version__ = '0.1'
```

```
from Models.discourse import Message, command_lists, compass, main_directions,  
directives, relation_names, text_to_relation, state_names  
from Models.world import *  
from Simulator.events import *  
import copy  
from random import choice
```

```
##- DISCOURSE -#
```

```
#-----#
```

```
frontmatter = (  
    'Adventure in Style',  
    'based on the classic by Will Crowther and Don Woods',  
    'Nick Montfort',  
    'Welcome to Adventure!!!'  
)
```

```
message = Message()

compass['ACROSS'] = ['across']
compass['BARREN'] = ['barren']
compass['BED'] = ['bed']
compass['BEDQUILT'] = ['bedquilt']
compass['BROKEN'] = ['broken']
compass['CANYON'] = ['canyon']
compass['CAVERN'] = ['cavern']
compass['CLIMB'] = ['climb']
compass['COBBLE'] = ['cobble']
compass['CRACK'] = ['crack']
compass['CRAWL'] = ['crawl']
compass['DARK'] = ['dark']
compass['DEBRIS'] = ['debris']
compass['DEPRESSION'] = ['depression']
compass['DOWNSTREAM'] = ['downstream']
compass['ENTER'] = ['enter']
compass['ENTRANCE'] = ['entrance']
compass['FLOOR'] = ['floor']
compass['FOREST'] = ['forest']
compass['FORK'] = ['fork']
compass['GIANT'] = ['giant']
compass['GULLY'] = ['gully']
compass['HALL'] = ['hall']
compass['HILL'] = ['hill']
compass['HOLE'] = ['hole']
compass['HOUSE'] = ['house', 'building']
compass['IN'] = ['in']
compass['JUMP'] = ['jump']
compass['LEAVE'] = ['leave']
compass['LEFT'] = ['left']
compass['LOW'] = ['low']
compass['NOWHERE'] = ['nowhere']
compass['ONWARD'] = ['onward']
compass['ORIENTAL'] = ['oriental']
compass['OUTDOORS'] = ['outdoors']
compass['OVER'] = ['over']
compass['PIT'] = ['pit']
compass['PLOVER'] = ['plover']
compass['PLUGH'] = ['plugh']
compass['RESERVOIR'] = ['reservoir']
compass['RIGHT'] = ['right']
```

```

compass['ROCK'] = ['rock']
compass['ROOM'] = ['room']
compass['SECRET'] = ['secret']
compass['SHELL'] = ['shell']
compass['SLAB'] = ['slab']
compass['SLIT'] = ['slit']
compass['STAIR'] = ['stair']
compass['STREAM'] = ['stream']
compass['SURFACE'] = ['surface']
compass['TUNNEL'] = ['tunnel']
compass['UPSTREAM'] = ['upstream']
compass['VALLEY'] = ['valley']
compass['VIEW'] = ['view']
compass['WALL'] = ['wall']
compass['XYZZY'] = ['xyzyz']
compass['Y2'] = ['y2']

```

```

#- WORLD -#
#-----#

```

```

class Lamp(Thing):
    def react(self, world, e):
        new_actions = []
        if e.get(0) == 'MODIFY' and e.get(2) == str(self) and \
            e.get(3) == 'LIT':
            if e.get(5):
                modify = Modify_Event(e.get(1), str(self), 'LIGHT', \
                    0.0, 0.6, e.timestamp, e.id)
                light_action = Action([modify], e.id)
            else:
                modify = Modify_Event(e.get(1), str(self), 'LIGHT', \
                    0.6, 0.0, e.timestamp, e.id)
                light_action = Action([modify], e.id)
            new_actions += [light_action]
        return new_actions

```

```

class Wanderer(Actor):
    def act(self, world):
        if choice([0,1,1,1,1]) == 1:
            return None
        direction = choice(world.ex[str(self)].r(world).exits.keys())

```

```

        to_tag = world.ex[str(self)].r(world).exit(direction)
        impel = Impel_Event(str(self), str(self), direction, 50, 'walk',
world.clock.ticks(), str(self) + '-ACTING', 5, 5)
        configure = Configure_Event(str(self), str(self), 'IN',
world.ex[str(self)].room(world), 'IN', to_tag, world.clock.ticks() + 10,
impel.id, 5, 5)
        return Action([impel, configure], str(self) + '-ACTING', ['GO',
str(self), direction])

```

```

class Waver(Actor):
    def act(self, world):
        misc = Misc_Event(str(self), 'wave', '-', '', '-', None,
world.clock.ticks(), str(self) + '-ACTING', 5, 5)
        return Action([misc], str(self) + '-ACTING', ['WAVE', str(self)])

```

```

existents = [

```

```

    Actor( 'ADVENTURER',
        called=('the', 'adventurer'),
        parent=('IN', 'END-OF-ROAD-1'),
        adjs=['my'],
        nouns=['self', 'person', 'individual', 'human', 'adventurer'],
        allowed=[('OF', '*')],
        qualities=['PERSON', 'MAN'],
        appearance='a nondescript adventurer',
        state={
            'GENDER': 'MALE'
        }
    ),

```

```

    Waver( 'PIRATE',
        called=('a', 'pirate'),
        parent=('IN', 'END-OF-ROAD-1'),
        adjs=['very', 'piratical'],
        nouns=['pirate'],
        allowed=[('OF', '*')],
        appearance='a man, a very piratical individual',
        qualities=['PERSON', 'MAN'],
        state={
            'GENDER': 'MALE'
        }
    ),

```

```

Room( 'END-OF-ROAD-1',
      called=('the', 'end of the road'),
      adjs=['surrounding'],
      nouns=['area'],
      appearance=[
        'S_? 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'
      ],
      exits={'HILL':'HILL-2', 'WEST':'HILL-2', 'UP':'HILL-2', 'ENTER':'BUILDING-3',
            'HOUSE':'BUILDING-3', 'IN':'BUILDING-3', 'EAST':'BUILDING-3',
            'DOWNSTREAM':'VALLEY-4', 'GULLY':'VALLEY-4', 'STREAM':'VALLEY-4',
            'SOUTH':'VALLEY-4', 'DOWN':'VALLEY-4', 'FOREST':'FOREST-5', 'NORTH':'FOREST-5',
            'DEPRESSION':'OUTSIDE-GRATE-8'}
      ),

Wanderer( 'DWARF',
          called=('a', 'dwarf'),
          parent=('IN', 'HILL-2'),
          adjs=['dwarflike', 'stout'],
          nouns=['dwarf'],
          allowed=[('OF', '*')],
          appearance='a stout dwarf',
          qualities=['PERSON', 'MAN'],
          state={
            'GENDER': 'MALE',
            'NUMBER': 'SINGULAR'
          }
          ),

Room( 'HILL-2',
      called=('a', 'hill in road'),
      adjs=['surrounding'],
      nouns=['cave', 'room'],
      appearance=[
        'S_FC V_walk_PERF up a_hill, still in the_forest', 'the_road V_slope_S
back down the_other_side of the_hill', 'there V_be_S a_building in
the_distance'],
      exits={'HILL-2':'END-OF-ROAD-1', 'HOUSE':'END-OF-ROAD-1', 'ONWARD':'END-OF-ROAD-1',
            'EAST':'END-OF-ROAD-1', 'NORTH':'END-OF-ROAD-1', 'DOWN':'END-OF-ROAD-1',
            'FOREST':'FOREST-5', 'SOUTH':'FOREST-5'}
      ),

Room( 'BUILDING-3',
      called=('the', "building's interior"),

```

```

    adjs=['surrounding'],
    nouns=['area'],
    appearance=[
        'S_FC V_be inside a_building, a_well_house for a large spring'],
    prohibited={'GO DOWNSTREAM':'The stream V_flow_S out through a pair of 1
foot diameter sewer pipes, too small to enter.','GO STREAM':'The stream
V_flow_S out through a pair of 1 foot diameter sewer pipes, too small to
enter.','ENTER STREAM':'The stream V_flow_S out through a pair of 1 foot
diameter sewer pipes, too small to enter.'},
    exits={'ENTER':'END-OF-ROAD-1', 'LEAVE':'END-OF-ROAD-1', 'OUTDOORS':'END-
OF-ROAD-1', 'WEST':'END-OF-ROAD-1', 'XYZZY':'DEBRIS-ROOM-11', 'PLUGH':'Y2-33'}
),

```

```

Thing( 'KEYS',
    called=('some', 'metal', 'glinting', 'keys'),
    parent=('IN', 'BUILDING-3'),
    adjs=['key','of','ring','typical'],
    nouns=['keys','key','keyring','ring'],
    allowed=[],
    appearance='keys, on a ring',
    qualities=['DEVICE','METAL'],
    state={
        'NUMBER': 'PLURAL'
    }
),

```

```

Thing( 'FOOD',
    called=('', 'food'),
    parent=('IN', 'BUILDING-3'),
    adjs=['tasty'],
    nouns=['food'],
    allowed=[],
    appearance=[
        'D_THIS V_be_S just some food'],
    state={
        'CONSUMED': False
    }
),

```

```

Thing( 'BOTTLE',
    called=('a', 'glass', 'clear glass', 'bottle'),
    parent=('IN', 'BUILDING-3'),
    adjs=['clear','glass'],

```

```

nouns=['bottle'],
appearance='a clear glass bottle, currently A_BOTTLE_OPEN',
state={
  'OPEN': False,
}
),

Lamp( 'LAMP',
called=('a', 'brass', 'brass', 'brass carbide', 'shiny brass carbide', \
  'lamp'),
parent=('IN', 'BUILDING-3'),
adjs=['shiny', 'brass', 'carbide'],
nouns=['lamp', 'lantern', 'light'],
qualities=['DEVICE', 'METAL'],
appearance=[
  'a brass carbide lamp, the kind often used for illuminating caves',
  'S_LAMP V_be shiny and A_LAMP_LIT'
],
state={
  'LIT': False,
  'LIGHT': 0.0
}
),

Room( 'VALLEY-4',
called=('the', 'valley'),
adjs=['surrounding'],
nouns=['area'],
appearance=[
  'S_FC V_be in a_valley in the_forest beside a_stream tumbling along
a_rocky_bed'],
exits={'UPSTREAM':'END-OF-ROAD-1', 'HOUSE':'END-OF-ROAD-1', 'NORTH':'END-
OF-ROAD-1', 'FOREST':'FOREST-5', 'EAST':'FOREST-5', 'WEST':'FOREST-5',
'UP':'FOREST-5', 'DOWNSTREAM':'SLIT-7', 'SOUTH':'SLIT-7', 'DOWN':'SLIT-7',
'DEPRESSION':'OUTSIDE-GRATE-8'}  ),

Room( 'FOREST-5',
called=('the', 'forest'),
adjs=['surrounding'],
nouns=['area'],
appearance=[
  'S_FC V_be in open forest, with a deep valley to one side'],
exits={'VALLEY':'VALLEY-4', 'EAST':'VALLEY-4', 'DOWN':'VALLEY-4',

```

```

'FOREST':'FOREST-6', 'WEST':'FOREST-5', 'SOUTH':'FOREST-5'}    ),

Room( 'FOREST-6',
      called=('the', 'forest'),
      adjs=['surrounding'],
      nouns=['area'],
      appearance=[
        'S_FC V_be in open forest near both a valley and a road'],
      exits={'HILL-2':'END-OF-ROAD-1', 'NORTH':'END-OF-ROAD-1',
'VALLEY':'VALLEY-4', 'EAST':'VALLEY-4', 'WEST':'VALLEY-4', 'DOWN':'VALLEY-4',
'FOREST':'FOREST-5', 'SOUTH':'FOREST-5'}    ),

Room( 'SLIT-7',
      called=('a', 'slit in the streambed'),
      adjs=['surrounding'],
      nouns=['area'],
      appearance=[
        'at P_FC feet all the water of the stream V_splash_S into a 2-inch slit
in the rock','downstream the streambed V_be_S bare rock'],
      exits={'HOUSE':'END-OF-ROAD-1', 'UPSTREAM':'VALLEY-4', 'NORTH':'VALLEY-4',
'FOREST':'FOREST-5', 'EAST':'FOREST-5', 'WEST':'FOREST-5',
'DOWNSTREAM':'OUTSIDE-GRATE-8', 'ROCK':'OUTSIDE-GRATE-8', 'BED':'OUTSIDE-GRATE-
8', 'SOUTH':'OUTSIDE-GRATE-8'},
      prohibited={'GO SLIT':"You don't fit through a two-inch slit!","GO
STREAM':"You don't fit through a two-inch slit!","GO DOWN':"You don't fit
through a two-inch slit!"}    ),

Room( 'OUTSIDE-GRATE-8',
      called=('the', 'area outside the grate'),
      adjs=['surrounding'],
      nouns=['area'],
      appearance=[
        'S_FC V_be in a 20-foot depression floored with bare dirt','set into the
dirt V_be_S O_GRATE mounted in concrete','a dry streambed V_lead_S into the
depression'],
      exits={'FOREST':'FOREST-5', 'EAST':'FOREST-5', 'WEST':'FOREST-5',
'SOUTH':'FOREST-5', 'HOUSE':'END-OF-ROAD-1', 'UPSTREAM':'SLIT-7',
'GULLY':'SLIT-7', 'NORTH':'SLIT-7', 'ENTER':'GRATE', 'DOWN':'GRATE'}    ),

Thing( 'GRATE',
      called=('a', 'strong steel grate'),
      parent=('PART', 'OUTSIDE-GRATE-8'),
      adjs=['strong','steel'],
      nouns=['grate','grating','grill','grille','barrier'],

```

```

    qualities=['DOORWAY','METAL'],
    allowed=[('THROUGH','*')],
    appearance='a grate, placed to restrict entry to the cave, which is
currently A_GRATE_OPEN',
    state={
        'OPEN': False,
        'LOCKED': True,
        'TO': 'BELOW-GRATE-9',
        'KEY': 'KEYS'
    }
),

Room( 'BELOW-GRATE-9',
    called=('the', 'area below the grate'),
    adjs=['surrounding'],
    nouns=['cave','room'],
    appearance=[
        'S_FC V_be in a small chamber beneath a 3x3 steel grate to the
surface','a low crawl over cobbles V_lead_S inward to the west'],
    state={
        'LIGHT': 0.7
    },
    exits={'LEAVE':'OUTSIDE-GRATE-8', 'EXIT':'OUTSIDE-GRATE-8', 'UP':'OUTSIDE-
GRATE-8', 'CRAWL':'COBBLE-CRAWL-10', 'COBBLE':'COBBLE-CRAWL-10', 'IN':'COBBLE-
CRAWL-10', 'WEST':'COBBLE-CRAWL-10', 'PIT':'SMALL-PIT-14', 'DEBRIS':'DEBRIS-
ROOM-11'}    ),

Room( 'COBBLE-CRAWL-10',
    called=('the', 'cobble crawl'),
    adjs=['surrounding'],
    nouns=['cave','room'],
    appearance=[
        'S_FC V_crawl_PROG over cobbles in a low passage','there V_be_S a dim
light at the east end of the passage'],
    state={
        'LIGHT': 0.5
    },
    exits={'LEAVE':'BELOW-GRATE-9', 'SURFACE':'BELOW-GRATE-9',
'NOWHERE':'BELOW-GRATE-9', 'EAST':'BELOW-GRATE-9', 'IN':'DEBRIS-ROOM-11',
'DARK':'DEBRIS-ROOM-11', 'WEST':'DEBRIS-ROOM-11', 'DEBRIS':'DEBRIS-ROOM-11',
'PIT':'SMALL-PIT-14'}    ),

Thing( 'CAGE',
    called=('a', 'wicker cage'),

```

```

parent=('IN', 'COBBLE-CRAWL-10'),
adjs=['wicker'],
nouns=['cage'],
appearance='a wicker cage, about the size of a breadbasket, currently
A_CAGE_OPEN',
state={
  'OPEN': True
}
),

Room( 'DEBRIS-ROOM-11',
called=('a', 'debris room'),
adjs=['surrounding'],
nouns=['cave', 'room'],
appearance=[
  'S_FC V_be in a room filled with debris washed in from the surface','a
low wide passage with cobbles V_become_S plugged with mud and debris here, but
an awkward canyon V_lead_S upward and west','a note on the wall V_say_S "MAGIC
WORD XYZZY"'],
state={
  'LIGHT': 0.0
},
exits={'ENTRANCE':'BELOW-GRATE-9', 'CRAWL':'COBBLE-CRAWL-10',
'COBBLE':'COBBLE-CRAWL-10', 'TUNNEL':'COBBLE-CRAWL-10', 'LOW':'COBBLE-CRAWL-
10', 'EAST':'COBBLE-CRAWL-10', 'CANYON':'AWKWARD-CANYON-12', 'IN':'AWKWARD-
CANYON-12', 'UP':'AWKWARD-CANYON-12', 'WEST':'AWKWARD-CANYON-12',
'XYZZY':'BUILDING-3', 'PIT':'SMALL-PIT-14'}
),

Thing( 'ROD',
called=('a', 'black rod'),
parent=('IN', 'DEBRIS-ROOM-11'),
adjs=['black', 'iron', 'rusty', 'sinister'],
nouns=['rod'],
appearance=[
  'an ordinary sinister black rod, one that is a bit rusty'],
state={
  }
),

Room( 'AWKWARD-CANYON-12',
called=('an', 'awkward canyon'),
adjs=['surrounding'],
nouns=['cave', 'room'],
appearance=[

```

```

    'S_FC V_be in an awkward sloping east/west canyon'],
state={
    'LIGHT': 0.0
},
exits={'ENTRANCE':'BELOW-GRATE-9', 'DOWN':'DEBRIS-ROOM-11',
'EAST':'DEBRIS-ROOM-11', 'DEBRIS':'DEBRIS-ROOM-11', 'IN':'BIRD-CHAMBER-13',
'UP':'BIRD-CHAMBER-13', 'WEST':'BIRD-CHAMBER-13', 'PIT':'SMALL-PIT-14'}    ),

Actor( 'DM',
called=('a', 'dungeon master'),
parent=('IN', 'AWKWARD-CANYON-12'),
adjs=['dungeon'],
nouns=['master','dm','chick'],
allowed=[('OF','*')],
qualities=['PERSON','WOMAN'],
appearance='a totally elite dungeon master woman, radiating light',
state={
    'LIGHT': 60,
    'GENDER': 'FEMALE'
}
),

Room( 'BIRD-CHAMBER-13',
called=('the', 'bird chamber'),
adjs=['surrounding'],
nouns=['cave','room'],
appearance=[
    'S_FC V_be in a splendid chamber thirty feet high','the walls V_be_P
frozen rivers of orange stone','an awkward canyon and a good passage V_exit_P
from east and west sides of the chamber'],
state={
    'LIGHT': 0.0
},
exits={'ENTRANCE':'BELOW-GRATE-9', 'DEBRIS':'DEBRIS-ROOM-11',
'CANYON':'AWKWARD-CANYON-12', 'EAST':'AWKWARD-CANYON-12', 'TUNNEL':'SMALL-PIT-
14', 'PIT':'SMALL-PIT-14', 'WEST':'SMALL-PIT-14'}    ),

Thing( 'BIRD',
called=('a', 'little bird'),
parent=('IN', 'BIRD-CHAMBER-13'),
adjs=['little','cheerful'],
nouns=['bird'],
appearance=[

```

```

    'Just a bird'],
state={
    }
),

Room( 'SMALL-PIT-14',
    called=('the', 'top of the small pit'),
    adjs=['surrounding'],
    nouns=['cave', 'room'],
    appearance=[
        'at P_FC feet V_be_S a small pit breathing traces of white mist', 'an east
passage V_end_S here except for a small crack leading on'],
    state={
        'LIGHT': 0.0
    },
    prohibited={'GO WEST': 'The crack is far too small to follow.', 'GO
CRACK': 'The crack is far too small to follow.'},
    exits={'ENTRANCE': 'BELOW-GRATE-9', 'DEBRIS': 'DEBRIS-ROOM-11',
'TUNNEL': 'BIRD-CHAMBER-13', 'EAST': 'BIRD-CHAMBER-13', 'DOWN': 'HALL-OF-MISTS-
15'}    ),

Room( 'HALL-OF-MISTS-15',
    called=('the', 'hall of mists'),
    adjs=['surrounding'],
    nouns=['cave', 'room'],
    appearance=[
        'S_FC V_be at one end of a vast hall stretching forward out of sight to
the west',
        'there V_be_P openings to either side',
        'nearby, a wide stone staircase V_lead_S downward',
        'the hall V_be_S filled with wisps of white mist swaying to and fro
almost as if alive',
        'a cold wind V_blow_S up the staircase',
        'there V_be_S a passage at the top of a dome behind O_FC'
    ],
    state={
        'LIGHT': 0.0
    },
    exits={'LEFT': 'NUGGET-ROOM-18', 'SOUTH': 'NUGGET-ROOM-18',
'ONWARD': 'FISSURE-EAST-17', 'HALL': 'FISSURE-EAST-17', 'WEST': 'FISSURE-EAST-17',
'STAIR': 'HALL-OF-MOUNTAIN-KING-19', 'DOWN': 'HALL-OF-MOUNTAIN-KING-19',
'NORTH': 'HALL-OF-MOUNTAIN-KING-19', 'UP': 'SMALL-PIT-14', 'Y2': 'JUMBLE-34'}    ),

Room( 'FISSURE-EAST-17',

```

```

called=('the', 'east bank of the fissure'),
adjs=['surrounding'],
nouns=['cave','room'],
appearance=[
    'S_FC V_be on the east bank of a fissure slicing clear across the hall',
    'the mist V_be_S quite thick here, and the fissure V_be_S too wide to
jump'
],
state={
    'LIGHT': 0.0
},
exits={'HALL':'HALL-OF-MISTS-15', 'EAST':'HALL-OF-MISTS-15', 'OVER':'WEST-
SIDE-OF-FISSURE-27'}    ),

Room( 'NUGGET-ROOM-18',
called=('the', 'nugget of gold room'),
adjs=['surrounding'],
nouns=['cave','room'],
appearance=[
    'this V_be_S a low room with a crude note on the wall','the note V_say_S
, "You won\'t get it up the steps"'],
state={
    'LIGHT': 0.0
},
exits={'HALL':'HALL-OF-MISTS-15', 'LEAVE':'HALL-OF-MISTS-15',
'NORTH':'HALL-OF-MISTS-15'}    ),

Thing( 'NUGGET',
called=('a', 'nugget of gold'),
parent=('IN', 'NUGGET-ROOM-18'),
adjs=['large','sparkling','of'],
nouns=['nugget','gold'],
appearance='a large gold nugget'
),

Room( 'HALL-OF-MOUNTAIN-KING-19',
called=('the', 'Hall of the Mountain King'),
adjs=['surrounding'],
nouns=['cave','room'],
appearance=[
    'S_FC V_be in the Hall of the Mountain King, with passages off in all
directions'],
state={

```

```

        'LIGHT': 0.0
    },
    exits={'STAIR':'HALL-OF-MISTS-15', 'UP':'HALL-OF-MISTS-15', 'EAST':'HALL-
OF-MISTS-15', 'NORTH':'ROOM-32', 'SECRET':'SECRET-EAST-WEST-CANYON-74'}    ),

. . . . .

pfn['focalized'] = 'ADVENTURER'
pfn['commanded'] = 'ADVENTURER'

#- ADVENTURER'S -#
#---- WORLD ----#
#-----#

adventurer_existents = [
copy.deepcopy(existents[0]),
copy.deepcopy(existents[1]),
copy.deepcopy(existents[2])
]

for dir in adventurer_existents[2].exits.keys():
    adventurer_existents[2].exits[dir] = '-'

#- DUNGEON-MASTER'S -#
#----- WORLD -----#
#-----#

dm_existents = copy.deepcopy(existents)

#----- PIRATE'S -----#
#----- WORLD -----#
#-----#

pirate_existents = copy.deepcopy(existents)

#----- DWARF'S -----#
#----- WORLD -----#
#-----#

dwarf_existents = copy.deepcopy(existents)

```

```
#- FINAL WORLD SETUP -#  
#-----#  
  
focalizers = [  
  ('ADVENTURER',adventurer_exists,{}),  
  ('DM',dm_exists,{}),  
  ('PIRATE',pirate_exists,{}),  
  ('DWARF',dwarf_exists,{})  
]
```

17 Appendix E

Example Plans for Narrating

Default

```
#- PLAN FOR --#
#- NARRATING -#
#-----#

pfn = {
  'focalized' : 'FC',
  'commanded' : 'CC',
  'order' : 'CHRONICLE',
  'speed' : .75,
  'frequency' : [('DEFAULT', 'SINGULATIVE')],
  'time' : 'SIMULTANEOUS',
  'explicit_I' : 0,
  'implicit_I' : 0,
  'explicit_you' : 0,
  'implicit_you' : 0,
  'narratee' : 'FC',
  'narrator' : None,
  'window' : 'CURRENT',
  'progressive' : False,
  'perfect' : False,
  'time-words' : True
}

output_filter = None
```

Hesitant

```
#- PLAN FOR --#
#- NARRATING -#
#-----#

from random import randint, choice
import re

class uh_filter:
    def token_filter(self, ts):
        if randint(1,6) == 1:
            prefix = choice(['', uh, ',', uh, ',', uh, ',', um, ',', um, ',',
er, '])
            ts = prefix + ts
        return ts
    def sentence_filter(self, ss):
        if ss[0:2] == ', ':
            ss = ss[2:]
            ss = re.sub('\, *\,','',ss)
            ss = re.sub('\, *\,','',ss)
            ss = re.sub('\, *\,','',ss)
            ss = re.sub('\, *\,','',ss)
        return ss
    def paragraph_filter(self, ps):
        return ps

output_filter = uh_filter()

pfn = {
    'focalized' : 'FC',
    'commanded' : 'CC',
    'order' : 'CHRONICLE',
    'speed' : .75,
    'frequency' : [('DEFAULT','SINGULATIVE')],
    'time' : 'SIMULTANEOUS',
    'explicit_I' : 0,
    'implicit_I' : 0,
    'explicit_you' : 0,
    'implicit_you' : 0,
    'narrator' : None,
    'narratee' : 'FC',
```

```
'window' : 'CURRENT',
'progressive' : False,
'perfect' : False,
'time-words' : True
}
```

Personal

```
#- PLAN FOR --#
#- NARRATING -#
#-----#
```

```
pfn = {
    'focalized' : 'FC',
    'commanded' : 'CC',
    'order' : 'CHRONICLE',
    'speed' : .75,
    'frequency' : [('DEFAULT', 'SINGULATIVE')],
    'time' : 'SIMULTANEOUS',
    'explicit_I' : 0,
    'implicit_I' : 0,
    'explicit_you' : 0,
    'implicit_you' : 0,
    'narratee' : None,
    'narrator' : 'FC',
    'window' : 'CURRENT',
    'progressive' : False,
    'perfect' : False,
    'time-words' : True
}
```

```
output_filter = None
```

Previous

```
#- PLAN FOR --#  
#- NARRATING -#  
#-----#
```

```
pfn = {  
    'focalized' : 'FC',  
    'commanded' : 'CC',  
    'order' : 'CHRONICLE',  
    'speed' : .75,  
    'frequency' : [('DEFAULT', 'SINGULATIVE')],  
    'time' : 'SUBSEQUENT',  
    'explicit_I' : 0,  
    'implicit_I' : 0,  
    'explicit_you' : 0,  
    'implicit_you' : 0,  
    'narratee' : None,  
    'narrator' : None,  
    'window' : 'CURRENT',  
    'progressive' : False,  
    'perfect' : False,  
    'time-words' : True  
}
```

```
output_filter = None
```

Prophetic

```
#- PLAN FOR --#  
#- NARRATING -#  
#-----#
```

```
pfn = {  
    'focalized' : 'FC',  
    'commanded' : 'CC',  
    'order' : 'CHRONICLE',  
    'speed' : .75,  
    'frequency' : [('DEFAULT', 'SINGULATIVE')],  
    'time' : 'PREVIOUS',  
    'explicit_I' : 0,  
    'implicit_I' : 0,  
    'explicit_you' : 0,  
    'implicit_you' : 0,  
    'narratee' : None,  
    'narrator' : 'FC',  
    'window' : 'CURRENT',  
    'progressive' : False,  
    'perfect' : False,  
    'time-words' : True  
}
```

```
output_filter = None
```

Retrograde

```
#- PLAN FOR --#  
#- NARRATING -#  
#-----#
```

```
pfn = {  
    'focalized' : 'FC',  
    'commanded' : 'CC',  
    'order' : 'RETROGRADE',  
    'speed' : .75,  
    'frequency' : [('DEFAULT', 'SINGULATIVE')],  
    'time' : 'SIMULTANEOUS',  
    'explicit_I' : 0,  
    'implicit_I' : 0,  
    'explicit_you' : 0,  
    'implicit_you' : 0,  
    'narratee' : None,  
    'narrator' : 'FC',  
    'window' : 3,  
    'progressive' : False,  
    'perfect' : False,  
    'time-words' : True  
}  
  
output_filter = None
```

Surprise

```

#- PLAN FOR --#
#- NARRATING -#
#-----#

from random import randint, choice
import re

class surprise_filter:
    def token_filter(self, ts):
        return ts
    def sentence_filter(self, ss):
        chosen = randint(1,8)
        if chosen == 1:
            ss = choice(['whoa, ', 'dude, ',]) + ss
        elif chosen == 2:
            ss = ss + choice([' man', ' dude',])
        ss = ss + '!'
        ss = re.sub(' !', '!!', ss)
        ss = re.sub(' *\\', ' ', ss)
        return ss
    def paragraph_filter(self, ps):
        chosen = randint(1,3)
        if chosen == 1:
            ps = ps + choice(['Amazing!', 'Wow!', 'Awesome!'])
        return ps

output_filter = surprise_filter()

pfn = {
    'focalized' : 'FC',
    'commanded' : 'CC',
    'order' : 'CHRONICLE',
    'speed' : .75,
    'frequency' : [('DEFAULT', 'SINGULATIVE')],
    'time' : 'PREVIOUS',
    'explicit_I' : 0,
    'implicit_I' : 0,
    'explicit_you' : 0,
    'implicit_you' : 0,
    'narratee' : None,
}
```

```
'narrator' : 'FC',  
'window' : 'CURRENT',  
'progressive' : False,  
'perfect' : False,  
'time-words' : True  
}
```

18 Appendix F

Evaluation Texts and Instructions

The 14 texts included after these instructions are typeset as they were when presented to annotators.

Instructions (1)

In this first pass, you are being asked to rate how natural a text appears on a scale of 1 (least natural) to 10 (most natural), using your judgment as a reader of English. Read the text word for word at a normal pace and try to understand what it communicates. You do not need to review or closely study it. Please indicate your rating on the sheet provided.

For the purposes of this evaluation, naturalness means being written in ordinary, fluent English. The more awkward and disfluent the text is, the less natural it is.

Each text is an excerpt that appears on its own page. It is the naturalness of the entire excerpt, not individual phrases or sentences within them, that you will be evaluating. Even if all the sentences seem natural when considered individually, you may still find that the overall text, considered as a whole, is not very natural.

Some of the excerpts are from dialogues in which a command such as ">go south" is followed by a reply. You should read these commands and take into account that they are part of the dialogue, but you should not consider whether these commands that occur after the ">" are written in natural English. Base your evaluation only on the quality of the replies.

Instructions (2)

In this second pass, you are being asked to mark all the events that are narrated in these texts. Please underline the main verb associated with each event. For instance, in the sentence "The train raced past the barn," *raced* should be underlined.

Only verbs associated with events that are happening in the narrative should be underlined; those associated with description, states, or propositions should not be. For instance, nothing should be underlined in "The books are on the shelf," "The street was quiet," or "Two plus two is four."

If you believe an event has been narrated and there is no verb to underline, underline whatever word is most appropriate.

Instructions (3)

In this final pass, you are being asked to indicate the order in which the narrated events occurred. In the left margin next to each verb you have underlined, please write what number that event is in chronological sequence.

If there is more than one verb underlined on a single line, write the numbers of those events next to each other in the margin.

If you know that two events occurred simultaneously, write the same number for both of them.

If you cannot tell where an event occurred in chronological sequence, write “?” for that event. You do not need to indicate any other information, even if you can tell something about the chronology — that this event must have happened after this other event, for instance. As long as you cannot find a definite point where an event occurs relative to other events whose chronology is known, write “?”.

1.

>get book

You sit down at the desk.

Taken.

Come to think of it, as with a hangover, dehydration is probably the problem. Your mouth is dry wool.

>read book

The Schmendrick's Guide to Desert Camping -- you've been studying it for weeks now. Vast alien reams of information: perspiration rates, sunburn factors, tent guying, and how to make a condensation still out of a trashbag and a tin can. You could probably write the book out longhand from memory; and none of it means a damn thing to you.

(That's why you're going on this trip, after all. It is, of course, the Death Valley Om -- half arts festival, half cult, a week in the deep desert where people show off, have sex, take drugs, and maintain a twenty-four-hour constant OM. Sand, heat, and thousands of throats. A space outside the world; with no dues except that you join in the chant, as much as is in you to chant.)

>get glass

You stand up, and step into the kitchen nook.

Taken.

2.

>unlock south door with key
Unlocked.

>go south
(Opening the south door)
Alexis' Room

This appears to be a actual "room," that is, it feels like a real person lived here. The nondescript furniture--a metal desk, a bookcase, a dresser and a short bed--through the most subtle alterations, give the room a warm feeling. Everything from the fading smell of lilacs to the small, glowing bud vase on the table, speak of an occupant determined to make a home in this sterile building. The bathroom door is to the west and a second door to the north leads out.

>touch vase

As your fingers touch the crystal, you hear the voices of David and Alexis behind you, but the sound fades in and out. You turn and see them lounging on the bed, the air of tension and discipline you've sensed in the other scenes completely absent here.

"You know, I think Jonas would kill us if he found out I was here," says David, grinning slowly.

"Look at us," laughs Alexis, "Acting like two frightened teenagers."

3.

>get letter

(the hastily written letter)

You pick up the hastily written letter. It consists of several pages torn from a yellow legal pad and stapled in the corner; each page is covered on both sides, top to bottom, with erratic, frightened handwriting. It is dated two days ago -- the day you arrived in Anchorhead.

>open drawer

You open the drawer, revealing a hastily written letter and a tarnished bronze key.

Your score has just gone up by two points.

>unlock drawer

You go through all the keys on your keyring, trying each one in turn, and after several false attempts you discover that the small steel key fits the lock.

You unlock the drawer.

>open drawer

It seems to be locked.

>take bronze key.

You pick up the tarnished bronze key. It is heavy and apparently very old, crusted over in places with a thin rime of verdigris.

4.

You looked at the newspaper sheet.

You saw the summary texts LEADER WORKING THROUGH NIGHT FOR COUNTRY, MONUMENT NEARS COMPLETION, and PURITY ACCOMPLISHED.

Before that, the punk kicked the tree.

Then, the flaneur wandered around briefly.

Before that, you looked at the northern area.

You saw the space north of the plaza's center, which was particularly barren of vegetation and ornament.

You saw the statue, the flaneur, and the rock.

Before that, the trash collector took something.

Then, the punk kicked the tree.

5.

>take rod

You take the black rod.

>go west

You go to the awkward canyon.

Then, you look at the awkward canyon.

- Awkward canyon -

You are in an awkward sloping east/ west canyon.

West is the bird chamber. Down is the debris room.

>go down

You walk to the debris room.

Then, you look at the debris room.

- Debris room -

You are in a room filled with debris washed in from the surface. A low wide passage with cobbles becomes plugged with mud and debris here, but an awkward canyon leads upward and west. A note on the wall says "MAGIC WORD XYZZY."

West is the awkward canyon. East is the cobble crawl.

6.

I brandished the black rod at the dungeon master.

I picked up the wicker cage.

I looked at the area below the grate.

I was in a small chamber beneath a 3x3 steel grate to the surface. A low crawl over cobbles led inward to the west.

I looked at the awkward canyon.

I was in an awkward sloping east/west canyon.

I saw the dungeon master.

I headed over to the area below the grate.

I looked at the area outside the grate.

I was in a 20-foot depression floored with bare dirt. Set into the dirt was the strong steel grate mounted in concrete. A dry streambed led into the depression.

I saw the strong steel grate.

I turned on the brass carbide lamp.

I picked up the black rod.

7.

A few stay to clean up the lobby, smiling as they fold the table and walk off, some taking the table and chairs, others carrying their rifles on their shoulders. With the event over, the whole corridor is left empty.

Someone fires a starting gun. The students reach for the parts and begin – as Sarah knows they will - to assemble guns, some sort of modern carbines. After a flurry of activity, one of them finishes by clipping in an empty magazine and slams the gun down on the table to the cheers of the crowd. The others finish putting together their weapons, then remove their blindfolds.

Sarah pauses. Here in the lobby area, a contest is about to start. Students sit blindfolded at a long table, grinning. Before each lies a pile of metal parts. A small crowd is standing about, chatting and looking on.

Sarah is in an interminable corridor that cuts all the way through the building. The floor is highly polished but old and uneven in places. Pipes, painted gray, run exposed, high along the walls. She is near the middle of this long hallway, where it widens a bit into a sort of lobby. The exit to the center of campus is northeast.

8.

You relinquished the cloak.

Before that, you looked at yourself.

You saw a typically nondescript character.

You saw the velvet cloak.

Before that, you looked at the cloakroom.

You saw that clearly, the walls of this small room were once lined with hooks, though then only one remained. The exit was a door to the east.

You saw the small brass hook.

Before that, you went to the cloakroom.

Before that, you looked at the foyer of the opera house.

You saw yourself standing in a spacious hall, splendidly decorated in red and gold, with glittering chandeliers overhead. The entrance from the street was to the north, and there were doorways south and west.

9.

You have so far completed 7 of the 8 tasks assigned you.

You have accomplished these important job-related tasks:

- Rebooting Nora.
- Rebooting Ester.
- Rebooting Sophie.
- Rebooting Wendy.
- Rebooting Cindy.
- Compiling and installing the crucial update.
- Offlining the rooted server.

10.

The adventurer looked at the valley.

The adventurer was in the valley in the forest beside the stream tumbling along the rocky bed.

Before that, the adventurer went to the valley.

Before that, the adventurer looked at the forest.

The adventurer was in open forest, with a deep valley to one side.

Before that, the adventurer headed over to the forest.

Before that, the adventurer looked at the hill in road.

The adventurer had walked up the hill, still in the forest. The road sloped back down the other side of the hill. There was the building in the distance.

Before that, the adventurer walked to the hill in road.

Meanwhile, the dwarf headed over to the end of the road.

11.

>look

You look at the center of the plaza.

Your senses are humming as you view the broad, circular, encircling Plaza of the Americas. The morning has concluded. It is midday now.

From here, you see a statue and a flaneur to the north, a fountain to the east, a trash collector to the southwest, a ball and a boy to the northeast, an obelisk to the south, and some punk and a tree to the west.

Then, the punk kicks the tree.

Meanwhile, the flaneur conveys himself to the northern area.

Then, the boy throws the ball.

Then, the trash collector picks up something.

Then, the ball falls to the ground.

>look at flaneur

You look at the flaneur.

You see a man who seems dedicated to strolling about.

Then, the punk kicks the tree.

Then, the boy takes the ball.

Then, the trash collector goes to the western area.

>go east

You walk to the eastern area.

Then, the punk kicks the tree.

Then, you look at the eastern area.

You see the space east of the plaza's center.

You see a shredded shirt, a newspaper sheet, and the fountain.

From here, you see the statue and the flaneur to the northwest, the boy to the north, the obelisk to the southwest, and the tree across the plaza.

Then, the trash collector picks up something.

12.

>go north

Gus interposes himself between you and the exit. Then he raises his arms above his head, flails them about, and says “BEEP BEEP, ACCESS DENIED!” while assuming the facial expression of the president of the United States.

>use the computer

Deftly you request a photo of Gus for purposes of positive ID, reset his account, and store in your short-term memory his cryptic-yet-memorable temporary password. Now you just need to share this important bit of information with Gus, and you’ll have executed your responsibilities most gloriously.

And people wonder why you fantasize about being a primatologist.

>tell gus his password

You tell Gus his cryptic-yet-memorable temporary password.

His face lights up with some combination of relief and prophetic rapture. He pulls a miniscule PDA from one of the pockets on his photographer’s vest and taps happily, presumably gaining access.

You get that satisfied feeling that you usually get after accomplishing a work-related task.

Some glitch, neurological or digital, causes Gus to scream “ODYSSEUS” and run toward the exit. This involves his running right over you, knocking you down. Your head strikes the floor in a painless and consciousness-eliminating way.

13.

You have so far scored 81 out of a possible 125, in 768 turns, giving you the rank of full-fledged housebreaker.

The score is made up as follows:

- 2 forging your first link
- 1 making the dancers dance
- 5 eradicating a rat
- 2 redecorating the foyer
- 3 having an interesting dream
- 5 forging your first reverse-link
- 4 turning on the universe
- 4 success in a culinary venture
- 3 finally getting over your hunger
- 5 floating the tea recipe
- 3 getting rid of a tangle of roses
- 2 snooping in Marie's papers
- 2 snooping in the Count's papers
- 20 collecting items of some value
- 20 visiting various places

14.

I conveyed myself to the end of the road.

I looked at the end of the road.

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

I walked to the valley.

I looked at the valley.

I was in the valley in the forest beside the stream tumbling along the rocky bed.

I walked to the slit in the streambed.

I looked at the slit in the streambed.

At my feet all the water of the stream splashed into a 2-inch slit in the rock.

Downstream the streambed was bare rock.

I conveyed myself to the area outside the grate.

I looked at the area outside the grate.

I was in a 20-foot depression floored with bare dirt. Set into the dirt was the strong steel grate mounted in concrete. A dry streambed led into the depression.

I saw the strong steel grate.

I unlocked the strong steel grate.

I opened the strong steel grate.

19 Bibliography

19.1 Works of Interactive Fiction

With the exception of works marked with (†), all of the following pieces of interactive fiction are available for free at the Interactive Fiction Archive, <<http://ifarchive.org>>. They can be located and downloaded most easily using *baf's Guide to the IF Archive*, <<http://www.wurb.com/if/>>.

- Adams, Douglas and Steve Meretzky (1984) *The Hitchhiker's Guide to the Galaxy*. Infocom.
Developed in ZIL. †
- Adams, Scott (1978) *Adventureland*. Adventure International. Developed in BASIC and the Adventure editor.
- Adams, Scott and Alexis Adams (1980) *Pirate's Adventure*. Adventure International. Developed in BASIC and the Adventure editor.
- Anderson, Dave and Liz Cyr-Jones (1986) *Hollywood Hijinks*. Infocom. Developed in ZIL. †
- Anderson, Timothy, Mark Blank, Bruce Daniels, and Dave Lebling (1977-78) *Zork*. Developed in MDL.
- Austin, Pete (1987) *Knight Orc*. Programmed by Mike Austin. Developed in A-code. Level 9/Rainbird. †
- Barlow, Sam (1999) *Aisle*. Developed in Inform 6.
- Barrett, Sean (2001) *Heroes*. Developed in Inform 6.
- Berlyn, Michael and Patricia Fogelman (1983) *Infidel*. Infocom. Developed in ZIL. †
- Berlyn, Michael (1983) *Suspended*. Infocom. Developed in ZIL. †
- Blank, Mark (1982) *Deadline*. Infocom. Developed in ZIL. †
- Blank, Mark and Dave Lebling (1980) *Zork I*. Infocom. Developed in ZIL. †
- (1983) *Enchanter*. Infocom. Developed in ZIL. †
- Cadre, Adam (1999) *Varicella*. Developed in Inform 6.
- (2000) *Shrapnel*. Developed in Inform 6.
- Crowther, Will and Don Woods (1977) *Adventure*. Developed originally in PDP-10 Fortran, ported

by Linards Ticmanis in 1998 to Fortran 77.

Devlin, Jason (2005) *Vespers*. Developed in Inform 6.

Disch, Thomas (1986) *Amnesia*. Programmed by Kevin Bently. Cognetics Corporation/Electronic Arts. Developed in King Edward. †

Finley, Ian (1997) *Babel*. Developed in TADS 2.

——— (1999) *Exhibition*. Developed in TADS 2.

Firth, Roger (1999) *Cloak of Darkness*. Developed originally in Inform 6. Released with source code and specification to facilitate re-implementation.

Gentry, Michael (1998) *Little Blue Men*. Developed in Inform 6.

——— (1998) *Anchorhead*. Developed in Inform 6.

Ingold, Jon (2001) *All Roads*. Developed in Inform 6.

Mason, Joe (1996) *In the End*. Developed in Inform 6.

Mateas, Michael (1997) *Fastfood World*. An Oz Text World. †

Mateas, Michael and Andrew Stern (2005) *Façade*. For Windows, later ported to Mac. †

Meretzky, Steve (1983) *Planetfall*. Infocom. Developed in ZIL. †

——— (1985) *A Mind Forever Voyaging*. Infocom. Developed in ZIL. †

Montfort, Nick (1999) *Winchester's Nightmare: A Novel Machine*. Developed in Inform 6.

——— (2000) *Ad Verbum*. Developed in Inform 6.

Moriarty, Brian (1985) *Wishbringer*. Infocom. Developed in ZIL. †

Nelson, Graham (1993) *Curses*. Developed in Inform 1-6.

O'Neill, Jeff (1987) *Nord and Bert Couldn't Make Heads or Tails of It*. Infocom. Developed in ZIL. †

Plotkin, Andrew (1996) *A Change in the Weather*. Developed in Inform 6.

——— (2000) *Shade*. Developed in Inform 6.

Pontius, Andrew (2000) *Rematch*. Developed in TADS 2.

Pinsky, Robert (1984) *Mindwheel*. Programmed by Steve Hales and William Mataga. Synapse/Brøderbund. Developed in BTZ. †

Rees, Gareth (1995) *Christminster*. Developed in Inform 6.

Schmidt, Dan (1999) *For a Change*. Developed in Inform 6.

Shiovitz, Dan (1999) *Bad Machine*. Developed in TADS 2

Short, Emily (2000) *Metamorphoses*. Developed in Inform 6.

- (2000) *Galatea*. Developed in Inform 6.
- (2005) *Bronze*. Developed in Inform 7.
- Starkey, Scott (2001) *The Beetmonger's Journal*. Developed in TADS 2.
- Wilson, James (2001) *Schroedinger's Cat*. Developed in Inform 6.

19.2 Other References

- Aarseth, Espen (1997) *Cybertext: Perspectives on Ergodic Literature*. Baltimore, Maryland: Johns Hopkins University Press.
- Aarseth, Espen (2001) Comparative Media Studies seminar, MIT, Cambridge, February 8.
- Adams, Scott (1980) "Pirate's Adventure." [Article with source code.] *BYTE* 5 12:192-212, December.
- Adams, Jon-K (1999) "Order and Narrative." In *Recent Trends in Narratological Research* ed. John Pier, 111-127. Tours: GRAAT.
- Androutsopoulos, Ion, Spiliotopoulos Dimitris, Stamatakis Konstantinos, Dimitromanolaki Aggeliki, Karkaletsis Vangelis and Spyropoulos Constantine (2002) "Symbolic Authoring for Multilingual Natural Language Generation." *Methods and Applications of Artificial Intelligence, Proceedings of the 2nd Hellenic Conference on Artificial Intelligence (SETN 2002)*, 131-142. Thessaloniki, Greece.
- Bal, Mieke (1997) *Narratology: Introduction to the Theory of Narrative*. 2nd ed. Toronto: University of Toronto Press.
- Barrett, Sean (2002) Conversation on ifMUD with Nick Montfort *et al.*, January 7.
- Barthes, Roland (1975). "An Introduction to the Structural Analysis of Narrative." *New Literary History* 6, 2:237-272.
- Blank, Mark and Stu Galley (1980) "How to Fit a Large Program Into a Small Machine." *Creative Computing* 5, 7:80-87, July.
- Bringsjord, Selmer and David A. Ferrucci (2000) *Artificial Intelligence and Literary Creativity: Inside the Mind of BRUTUS, a Storytelling Machine*. Hillsdale, NJ: Lawrence Erlbaum.
- Brooks, Cleanth Jr. and Robert Penn Warren (1943) *Understanding Fiction*. New York: Appleton-Century-Crofts.
- Buckles, Mary Ann (1985) "Interactive Fiction: The Computer Storygame 'Adventure.'" Ph.D.

- Thesis, University of California San Diego.
- Cadre, Adam (2002) "Re: Default parser responses: how do they affect the gaming / authorship experience?" <news://rec.arts.int-fiction> Message-ID <a11jtp\$8kk\$1@drizzle.com>. January 3.
- Callaway, Charles B. and James C. Lester. (2002) "Narrative Prose Generation." In *Artificial Intelligence* 139, 2:213-252, August.
- Carbol, Roger (2001) "Locational Puzzle Theory." <news://rec.arts.int-fiction> Message-ID <82675075.0110020840.12d1c469@posting.google.com>. October 2.
- Chatman, Seymour (1975) "Towards a Theory of Narrative." *New Literary History* 6, 2:295-318. Winter.
- (1978) *Story and Discourse: Narrative Structure in Fiction and Film*. Ithaca, NY: Cornell.
- Correia, Alfred (1980) "Computing Story Trees." *American Journal of Computational Linguistics* 6:135-149.
- Coyne, Bob and Richard Sproat (2001) "WordsEye: An Automatic Text-to-Scene Conversion System." *Proceedings of ACM SIGGRAPH 2001*, 487-496.
- Cox, Gregory (1999) "Re: [Theory] When is a puzzle a puzzle?" <news://rec.arts.int-fiction> Message-ID <7i3ogc\$kc3\$1@news4.svr.pol.co.uk>. May 21.
- Eskelinen, Markku (1998) "Omission Impossible — The Ergodics of Time." *Digital Arts and Culture* 1998, Bergen, Norway, <http://cmc.uib.no/dac98/papers/eskelinen.html>.
- Genette, Gerard (1980) *Narrative Discourse: An Essay in Method*. Trans. J. E. Lewin. Ithaca, NY: Cornell University Press.
- (1988) *Narrative Discourse Revisited*. Trans. J. E. Lewin. Ithaca, NY: Cornell University Press.
- Gervás, Pablo, Birte Lönneker-Rodman, Jan Christoph Meister, and Federico Peinado (2006) "Narrative Models: Narratology Meets Artificial Intelligence." In *Proceedings of Satellite Workshop: Toward Computational Models of Literary Analysis, 5th International Conference on Language Resources and Evaluation*, eds. Roberto Basili and Alessandro Lenci, 44-51. 22nd May, Genoa, Italy.
- Giner-Sorolla, Roger (1996) "Puzzles, problem-solving, and IF." <news://rec.arts.int-fiction> Message-ID <Pine.SUN.3.91.960405121744.28920B-100000@xp.psych.nyu.edu>. April 5.
- Herman, David (2002) *Story Logic: Problems and Possibilities of Narrative*. Lincoln: University of

- Nebraska Press.
- Hovy, Eduard H. (1988) *Generating Natural Language under Pragmatic Constraints*. Lawrence Erlbaum, Hillsdale, NJ.
- Hoy, Matt and Dennis G. Jerz (2001) "Storytelling in Computer Games Past, Present and Future (UWEC Panel featuring Scott Adams)." September 2.
<<http://jerz.setonhill.edu/if/adams/>>.
- James, Henry (1972) *Theory of Fiction: Henry James*. Edited by James E. Miller, Jr. Lincoln: University of Nebraska Press.
- Jerz, Dennis G. (2001) "An Annotated Bibliography of Interactive Fiction Scholarship." August 25,
<<http://jerz.setonhill.edu/if/bibliography/>>.
- Juul, Jesper (2001) "Games Telling Stories?" *Game Studies* 1:1.
<<http://www.gamestudies.org/0101/juul-gts/>>.
- Kantrowitz, Mark (1990) "GLINDA: Natural Language Text Generation in the Oz Interactive Fiction Project." Technical Report CMU-CS-90-158, School of Computer Science, Carnegie Mellon University.
- Klein, Sheldon, J.F. Aeschlimann, D. F. Balsiger, S. L. Converse, C. Court, M. Foster, R. Lao, J. D. Oakely and J. D. Smith (1973) "Automatic Novel Writing: A Status Report," UWCS Technical Report No. 186.
- Koller, Alexander, Ralf Dubusmann, Malte Gabsdil, and Kristina Striegnitz (2004) "Put my galakmid coin into the dispenser and kick it: Computational Linguistic and Theory Proving in a Computer Game." *Journal of Logic, Language and Information* 13, 2:187-206.
- Lebling, P. David, Marc S. Blank, and Timothy A. Anderson (1979) "Zork: A Computerized Fantasy Simulation Game." *IEEE Computer* 12, 4:51-59. April.
- Lebling, P. David (1980) "Zork and the Future of Computerized Fantasy Simulations." *BYTE* 5, 12:172-182. December.
- Lebowitz, Michael (1984) "Creating Characters in a Story-Telling Universe." *Poetics* 13, 171-194.
- Liu, Hugo and Henry Lieberman (2005) "Metafor: Visualizing Stories as Code." *Proceedings of the ACM International Conference on Intelligent User Interfaces, IUI 2005*, January 9-12, 2005, San Diego, CA, USA.
- Lönneker, Birte. 2005. "Narratological Knowledge for Natural Language Generation." *Proceedings*

- of the 10th European Workshop on Natural Language Generation (ENLG 2005), 91-100.
Aberdeen, Scotland, August 8-10.
- McKeown, Kathleen R. (1985) *Text Generation*. Cambridge University Press.
- Mateas, Michael (1997) "Computational Subjectivity in Virtual World Avatars." *Working notes of the Socially Intelligent Agents Symposium, AAAI Fall Symposium Series*. Menlo Park, Calif.: AAAI Press.
- (2002) "Interactive Drama, Art, and Artificial Intelligence." Ph.D. Thesis. Technical Report CMU-CS-02-206, School of Computer Science, Carnegie Mellon University.
- Mathews, Harry and Alastair Brotchie, eds. (1998) *Oulipo Compendium*. London: Atlas Press.
- Matthiessen, Christian M. I. M. and John A. Bateman. (1991) *Text Generation and Systemic-Functional Linguistics: Experiences from English and Japanese*. London and New York: Frances Pinter Publishers and St. Martin's Press.
- Meehan, James (1976) "The Metanovel: Writing Stories by Computer." Ph.D. Dissertation. Yale University.
- Meister, Jan Christoph (2003) *Computing Action: A Narratological Approach*. Berlin: Walter de Gruyter.
- Montfort, Nick (2000-2001) "Cybertext Killed the Hypertext Star." *ebr (Electronic Book Review) 11*, <<http://www.electronicbookreview.com/ebr11/11mon/index.html>>. Winter.
- (2003a) *Twisty Little Passages: An Approach to Interactive Fiction*. Cambridge: The MIT Press.
- (2003b) "Toward a Theory of Interactive Fiction." Version 3.5, December 19. Version 1.0 published 8 January 2002. <<http://nickm.com/if/toward.html>>. Forthcoming in *IF Theory*, ed. Emily Short.
- (2004) "Interactive Fiction as 'Story,' 'Game,' 'Storygame,' 'Novel,' 'World,' 'Literature,' 'Puzzle,' 'Problem,' 'Riddle,' and 'Machine.'" In *First Person: New Media as Story, Game, and Performance*. Cambridge: The MIT Press.
- (2005) "Artificial Intelligence and Narrative." *The Routledge Encyclopedia of Narrative Theory*. Edited by David Herman, Manfred Jahn, and Marie-Laure Ryan. London: Routledge.
- (2006) "Natural Language Generation and Narrative Variation in Interactive Fiction." *Computational Aesthetics: Artificial Intelligence Approaches to Beauty and Happiness, Papers from*

- the AAAI Workshop*. Technical Report WS-06-04, 45-52.
- Montfort, Nick and Stuart Moulthrop (2003) "Face It, Tiger, You Just Hit the Jackpot: Reading and Playing Cadre's *Varicella*." *Fineart Forum* 17, 8. <http://www.msstate.edu/Fineart_Online/Backissues/Vol_17/faf_v17_n08/reviews/montfort.html>.
- Moore, Robert C. (1980) "Reasoning About Knowledge and Action." SRI Technical Note 191, October.
- Morningstar, Chip and F. Randall Farmer (2003) "The Lessons of Lucasfilm's Habitat." *The New Media Reader*, eds. Noah Wardrip-Fruin and Nick Montfort, Cambridge: MIT Press, 664-677.
- Motte, Warren F., Jr., ed. and trans. (1986) *Oulipo: A Primer of Potential Literature*. Lincoln: University of Nebraska Press.
- Murray, Janet (1995) "The Pedagogy of Cyberfiction: Teaching a Course on Reading and Writing Interactive Narrative." *Contextual Media: Multimedia and Interpretation*, ed. Edward Barrett and Marie Redmond, 129-162. Cambridge: The MIT Press, 1995.
- (1997) *Hamlet on the Holodeck: The Future of Narrative in Cyberspace*. New York: Free Press.
- Myers, David Samuel (2001) "PrologueComp." <<http://www.sinc.sunysb.edu/Stu/dmyers/PrologueComp.html>>. June.
- Nelson, Graham (1995) "The Craft of Adventure: Five Articles on the Design of Adventure Games." <<http://www.ifarchive.org/if-archive/programming/general-discussion/Craft.Of.Adventure.txt>> January 19.
- (2001) *The Inform Designer's Manual*, 4th ed., 1st print ed. <<http://www.gnelson.demon.co.uk/inform/DM4.pdf>>. St. Charles, Illinois: The Interactive Fiction Library.
- Olsson, Magnus (1997) "What is a meta-command?" <[news://rec.arts.int-fiction](mailto:rec.arts.int-fiction)> Message-ID <5g6fqg\$5t5@bartlet.df.lth.se>. Mar 12.
- Pagis, Dan (1996) "Toward a Theory of the Literary Riddle" in *Untying the Knot: On Riddles and Other Enigmatic Modes*, ed. Galit Hasan-Rokem and David Shulman, 81-108. New York: University of Oxford Press.
- Parker, Marnie (2000) "An Iffy Theory." Verson 2.25. <<http://members.aol.com/doepage/theory.html>> February 19.
- Pérez y Pérez, Rafael and Sharples, Mike (2001) "MEXICA: A Computer Model of a Cognitive

- Account of Creative Writing." *Journal of Experimental & Theoretical Artificial Intelligence* 13, 119-139.
- Peterson, Dale (1983) *Genesis II: Creation and Recreation with Computers*. Reston, Va: Reston Publishing Company.
- Pereira, Franciso C., Raquel Hervás, Pablo Gervás, and Amilcar Cardoso (2006) "A Multiagent Text Generator with Simple Rhetorical Habilities." *Computational Aesthetics: Artificial Intelligence Approaches to Beauty and Happiness, Papers from the AAAI Workshop*. Technical Report WS-06-04, 37-44.
- Pinsky, Robert. (1995) "The Poetics of Zork." *The New York Times Book Review*, March 19:3+.
- Plotkin, Andrew (2001) "Interactive Fiction." <<http://www.eblong.com/zarf/if.html>> May 25.
- Polyani, Livia (1985) *Telling the American Story: A Structural and Cultural Analysis of Conversational Storytelling*. Norwood, NJ: Albex.
- Prince, Gerald (1973) *A Grammar of Stories*. The Hague: Mouton.
- (1980) "Aspects of a Grammar of Narrative." *Poetics Today* 1, 3:49-63.
- (1982) *Narratology: The Form and Functioning of Narrative*. Berlin: Mouton.
- (1987) *A Dictionary of Narratology*. Lincoln: University of Nebraska Press.
- (1988) "The Disnarrated." *Style* 22, 1-8.
- (2003) *A Dictionary of Narratology*. Revised Edition. Lincoln: University of Nebraska Press.
- Rees, Gareth. "Puzzle classification." <<news://rec.arts.int-fiction>> Message-ID 1993May7.172931.23328@infodev.cam.ac.uk 7 May 1993.
- Reichenbach, Hans (1947) *Elements of Symbolic Logic*. New York: Macmillan.
- Reiter, Ehud and Robert Dale (2000) *Building Natural Language Generation Systems*. Cambridge University Press.
- Rimmon-Kenan, Shlomith (2002) *Narrative Fiction: Contemporary Poetics*. 2nd Edition. London: Routledge.
- Ryan, Marie-Laure (1991) *Possible Worlds, Artificial Intelligence, and Narrative Theory*. Bloomington: Indiana University Press.
- (2001) *Narrative as Virtual Reality*. Baltimore: Johns Hopkins University Press.
- (2003) Foreword. *Computing Action: A Narratological Approach*. By Jan Christoph Meister. Trans. Alastair Matthews. Berlin: Walter de Gruyter.

- (2005) "Possible-Worlds Theory." *The Routledge Encyclopedia of Narrative Theory*. Edited by David Herman, Manfred Jahn, and Marie-Laure Ryan. London: Routledge.
- (2006) *Avatars of Story*. Minneapolis: University of Minnesota Press.
- Schank, Roger (1975) "The Primitive ACTs of Conceptual Dependency." *Proceedings of the 1975 Workshop on Theoretical Issues in Natural Language Processing (TINLAP 1975)*, 34-37. Cambridge, Mass.
- Schmidt, Dan (1999) "For A Change beta diary." <news://rec.arts.int-fiction> Message-ID <wkwvrak0ba.fsf@thecia.net>. November 22.
- Short, Emily (2001) "What's IF?" <http://emshort.home.mindspring.com/whatsif.html> November 29.
- SmartMonsters, Inc. (2006) "Welcome to TriadCity!" <http://www.smartmonsters.com/TriadCity/index.jsp>
- Smith, Sean and Joseph Bates (1989) "Towards a Theory of Narrative for Interactive Fiction." Carnegie Mellon University Technical Report CMU-CS-89-121. February 20.
- Solomon, Eric (1984) *Games Programming*. Cambridge University Press.
- Sternberg, Meir (1990) "Telling in Time (I): Chronology and Narrative Theory." *Poetics Today* 11:901-948.
- Szilas, Nicolas (2001) "A new approach to interactive drama: From intelligent characters to an intelligent virtual narrator." *Proceedings of the AAAI Spring Symposium on AI and Interactive Entertainment*, 72-76.
- Turner, Scott R. (1994) *The Creative Process: A Computer Model of Storytelling and Creativity*. Hillsdale, NJ: Lawrence Erlbaum.
- Weizenbaum, Joseph (1966) "Eliza — a computer program for the study of natural language communication between man and machine." *Communications of the ACM* 9, 1:36-45. January.
- Winograd, Terry (1972) *Understanding Natural Language*. New York: Academic Press.
- Wittig, Rob for IN.S.OMNIA (1994) *Invisible Rendezvous: Connection and Collaboration in the New Landscape of Electronic Writing*. Hanover, N.H.: Wesleyan University Press.
- Young, R. Michael (2007) "Story and discourse: A bipartite model of narrative generation in virtual worlds." Forthcoming in *Interaction Studies* 8, 2:177–208.

Zimmerman, Eric (2000) "Aesthetics of Game Design." Panel at Computers and Video Games
Come of Age. MIT, Cambridge. February 11.

Zimmerman, Eric (2001) "Game-Stories: Simulation, Narrative, Addiction." Panel at SIGGRAPH
2001, Los Angeles. August 17.