Spring 2010

The Imagined Moment

Inderjeet Mani

The Imagined Moment

FRONTIERS OF NARRATIVE

SERIES EDITOR: David Herman

Ohio State University

The Imagined Moment

Time, Narrative, and Computation

INDERJEET MANI

University of Nebraska Press ¦ Lincoln and London

Library of Congress
Cataloging-in-Publication Data
Mani, Inderjeet.
The imagined moment: time,
narrative, and computation /
Inderjeet Mani.
p. cm.—(Frontiers of narrative)
Includes bibliographical references
and index.
ISBN 978-0-8032-2977-8
(cloth: alk. paper)
1. Time in literature. 2. Fiction—
History and criticism. I. Title.
PN3352.T5M35 2010
809.3'933—dc22
2009049418

Set in Minion Pro and Meta by
Bob Reitz.
Designed by A. Shahan.

CONTENTS

ILLUSTRATIONS

TABLES

ACKNOWLEDGMENTS

I am deeply grateful to the following scholars for their extensive comments on earlier drafts of the manuscript: David Herman, Ohio State University; Birte Lönneker-Rodman, University of Hamburg; Nick Montfort, MIT; and James Pustejovsky, Brandeis University. David Elson at Columbia University also provided valuable feedback. Thanks are also due to the helpful staff at the Lauinger Library at Harvard University and the Hayden Library at MIT. I would also like to thank my editor at the University of Nebraska Press, Kristen Elias Rowley, for many helpful suggestions. My copyeditor, Barb Wojhoski, also deserves thanks for her careful scrutiny of the manuscript. Any errors remaining are entirely my own doing.

The Imagined Moment

Introduction

Stories inform our lives. The prince who renounces his kingdom to seek the meaning of life; the warrior who returns home after many travails and is reunited with his faithful wife; the ingénue who learns to find her way in a corrupt world; and the valet who rescues his master from embarrassments and entanglements—these stories live on in our imagination, long after the author has disappeared; they make their presence felt in books, movies, speeches, excuses for war, campaigns for peace, and moments when we throw away the cares of life and seek humorous relief.

How is it that a few strokes of a pen or taps of a key can summon up entire worlds, immersing us in the triumphs and traumas of various characters, and even suggesting how exemplary lives should or should not be led? My book attempts to answer this question, focusing on one key slice of narrative: time.

I examine time in narrative primarily from the perspective of recent developments in linguistics and the computational sciences. I try to account for some of the delight we feel in reading, in particular explaining our identification with narrative characters and their journeys in time. I do this with the help of computer models.

Computers have had a long history of use in analysis of literary works, the best-known area being authorship identification. Recent successes have included the unmasking of the journalist Joe Klein as the anonymous author of the roman à clef *Primary Colors*. (What gave Klein away was his relatively frequent use of adjectives ending in "inous," like "vertiginous," and his hyperactive use of prefixes such as "hyper," "mega," and "quasi.") Almost all of this literary forensics work has been focused on vocabulary, spelling, and grammar, aided by statistical analyses.[1] These methods are very useful in revealing the style and personal stamp of the author. However, computer science today can dig much deeper. I hope to convince the reader that recent developments allow us to posit a richly textured, even scientific account of what it means to understand the flow of time in literary stories.

As everyone knows, new developments in computer technology have made it possible for anyone to publish content in a variety of new media and to search and communicate almost instantaneously. Children use these tools of our age to tell stories. All sorts of stories in a variety of languages can be found today on the Web, manifested in e-mails, travel blogs, newsgroups, text messages, medical forums, even the dossiers of human rights tribunals; the variety seems almost endless. Much attention has been paid in the news media to the ways in which one can get noticed online and to the many new communication styles and brogues that can be found in text messaging, chat, and blogs. However, there has been relatively little discussion of what the technologies for sifting through information in human language on the Web can tell us about the way stories are put together.

News about some of this sifting has surfaced in the media now and then, mostly under the rubric of "data mining" and "surveillance" technologies that can crunch through reams of data to keep tabs on people's behavior. The newsworthiness of these technologies has so far been mainly related to fears about the further rupturing of the few remaining shreds of privacy we still cling to in our hyper-networked and highly surveillance-oriented society. However, the underlying techniques for processing data, particularly data in human languages, are of considerable relevance for narrative understanding. The main goal of this book is to show how these techniques can be of value beyond their use by Big Brother and Big Business, in particular, as applied to understanding how stories make use of time.

Although these techniques were aimed originally at applications favored by the military-industrial complex, the goals of this book are entirely humanistic, in the broad sense of seeking to enrich our literary culture, which is concerned, ultimately, with human self-expression and freedom. Such a humanistic approach can hardly be avoided when the subject is time.

We think of time as an additional dimension offered by the brilliant conceptual lens through which we view life. This dimension speaks to us of incessant change, of epochs and eras, of trajectories marked by beginnings and endings, births and extinctions. In addition to this rather somber message, time is also a harbinger of freedom, allowing us to escape from the eternal present, to resurrect a past that includes victory and failure, pleasure and suffering, and also giving us the wherewithal to prepare for and envision a better future. These twin characteristics of finitude and escape can be viewed as being in opposition, the first representing the heavy tug of reality, the second representing a flight toward freedom. Thus, many of our narratives of life, the stories we tell one another, reflect the uniquely human qualities of this tension. The computational theory I propose illuminates these very human aspects of timelines, allowing us to track characters filled with lofty or insidious goals, experiencing events that unfold along a timeline, reacting with pleasure or sadness, reconstructing a past and gazing out into a future.

I hope the book will not only alert lovers of literature to the power and promise of our tools today but also excite the interest of those who might want to encourage the further development of these models. To make the book interesting to the general reader, the ideas are presented here without any mathematical notation, using pictures to illustrate technical points.

THE THEORY IN A NUTSHELL

In order to understand how time is used in narrative, one needs a computational model of how time is used in language. My perspective on narrative is informed by research in computational linguistics. Developments in this field have resulted in systems that can answer questions as to when particular events in a text narrative are purported to happen. These systems create arbitrarily complex timelines that track mentions of events and times in any given document. Such systems, which can represent hypothetical past or future events, can anchor and order events in time, determining when an event happens, or whether one event occurs before another, irrespective of the order in which events are narrated.

Given this background, this book explores a computational theory of time in fictional narratives. I begin with a characterization of the timelines of literary texts. Here one must take into account not only the ordering and anchoring of events in time, but the embedding of stories within stories, and correspondences between entities of differing fictive status, such as events or times that are stated to occur or are merely remembered or imagined. These timelines can capture more than the succession of events: they can also represent their durations, and their distances from one another and from the time of narration, when such information is available. Given that such information is often not available, the model is still able to detect when the cloak of vagueness is present. All in all, this first layer is able to answer the question, given any text from Aesop's fables to the *nouveau roman*, of when something happened.

Narratives usually are more than an ordered list of events: they involve a story line that establishes a connection between the events. A reader can hazard a guess as to the motives behind the actions of particular characters and the outcomes and ramifications of those actions. However, there are many sequences of events in a person's life, such as my morning rituals of lifting the cap off the toothpaste tube and then brushing my teeth, that even if explained thoroughly in terms of when and why and whereof, are of no interest whatsoever (a point that some of today's bloggers seem not to have grasped). An event in a story is only significant if it means something, in a fundamental sense. And what it means, I believe, is tied to theories of what is good, desirable, avoidable, bad, and so forth in our thoughts and actions. A man may be a devoted father but a tyrannical ruler, and those behaviors could alter as the narrative evolves. A character might be uninteresting at the beginning but may grow in complexity. These are judgments that the reader makes, and these judgments can evolve over the course of a narrative. A layer is needed, therefore, to track, over the course of a narrative, changes in our emotional reactions to particular characters. Following Aristotle, I characterize the (idealized) reader's emotion only in terms of sympathy for or antipathy toward a character's situation. This layer, therefore, addresses the question of how a character evolves over time in the eyes of the actual reader. The actual reader is a member of what Peter Rabinowitz calls the "actual audience," namely, the "flesh-and-blood people who read the book" (1988, 20).[2] As Rabinowitz argues, each member of the actual audience can differ in the way she interprets a given work; clearly, interpretation can vary due to many factors, including the reader's class, ethnicity, historical period, familiarity with the work and its author, and so on. However, I argue that it is possible to arrive at generalizations about members of the actual audience based on comparing the reactions of one reader with another.

My theory thus characterizes the manner in which time is used in narrative to help characters construct and evolve identities that readers are happy to pass judgment on. The theory does not require that the interpretation of a text be completely determined, or even that there be a single interpretation. Of course, there is interpretation, and much can be learned that is of literary value from the study of the hermeneutic process of interpretation.

These two layers do not exhaustively cover the role that time plays in narrative. As the phenomenological philosopher Paul Ricoeur has argued, time and narrative can be viewed as intrinsically linked: "Time becomes human time to the extent that it is organized after the manner of a narrative; narrative, in turn, is meaningful to the extent that it portrays the features of temporal experience" (1984, 3). In Ricoeur's theory this linking of time and narrative occurs at three levels, tied to a theory of narrative as mimesis, the "imitation of action." Narrative comes (is "prefigured") with a variety of capabilities allowing for expression of aspects of action, including who did what to (or what happened to) whom; the agents, events, and entities involved; and the temporal order of actions. This level feeds into a second level of "imaginative configuration" of those action-related elements into a coherent structure by means of plot, providing an explanation for why things happen in the narrative, and allowing for considerable flexibility in the temporal locations, orders, and tempos in terms of which chronologies can be expressed. Finally, this level is mapped to a third level of "refiguration," which allows the reader to interpret the text in terms of his or her own experience—"the final indicator of the refiguring of the world of action under the sign of the plot" (77).

My account, while stopping short of any grand, all-encompassing theory, addresses aspects of these three levels within a computational framework. In terms of prefiguration, I delve into how computational tools can compute a basic level of linguistic analysis of the meaning of sentences in a narrative, one that covers the events, participants, their times, and places, and that provides a basic timeline for the text. In terms of configuration, my analysis includes the ability to detect various kinds of orders in which events are narrated, as well as their tempos.

Although I certainly believe that plot is crucial to any theory of narrative, I sidestep the modeling of plot, because more foundational research must be done in the field of artificial intelligence before one can effectively compute plots. Finally, regarding refiguration, I focus on a single aspect of the reader's interpretation of the narrative: her evaluation of the character's situation at particular points in the narrative.[3]

Leveraging the power of modern computational tools, my theory can allow a system to track when a character gazes back or ahead in time to reflect on experience, to contemplate roads not taken and roads that might be taken. My theory can also track the varying speeds at which time passes in a narrative, the modulation of which allows the reader to experience events at different levels of detail and to await long-expected outcomes with added suspense. It can also address the question of the span of time a story covers, given that novels said to be set on a single day often have references to times well before and after that. Pace and span together contribute to the reader's sense of how fast the narrative has flown.

Reading a story with pleasure involves a constructive activity by the reader; indeed, "expression creates being," as the French critic Gaston Bachelard (1964, xxiii) has observed. The experience of reading a narrative is similar, in some respects, to performing the music of a composer, except that reading involves an act of simulation. Time in narrative is especially interesting in this regard because the reader is carrying out her simulation in time, at her own pace, squeezing the reading into free moments. The reader can vicariously experience, as she travels about with other characters, entirely different experiences and thoughts in the past and the future, passing quickly over decades and even centuries; meanwhile, she also gratefully recognizes experiences and thoughts that echo her own. Most importantly, she can fast-forward or rewind—all of time is before her; she seems to be free at last of the weight of time. But at some point this freedom ends; the clock reminds one it's time to put away the book or the computer and to instead return to the day job, to run another errand or satisfy another master.

This contrast between real and imagined worlds becomes even more vivid when stories are digitized in various ways, from electronic texts to computer-mediated or computer-generated novels, allowing for nonlinear reading styles, and the co-opting of the reader in selecting or creating alternative futures. Overall, new modes of storytelling, along with the technologies of immersive game playing, have introduced rather novel ways of exploring time. These temporal aspects also have a variety of implications, some of them disturbing, for the way stories are and will be experienced. Chapter 9 ("Digital Storytelling") explores such issues.

PSYCHOLOGICAL CLAIMS

Having spent several pages outlining this theory, I would like to briefly indicate what this theory is not about. The computational theory is not intended to model how humans understand narrative. A prerequisite for a theory of how humans understand narrative, given that my account is focused on language, would be an adequate explanation, backed by empirical evidence, of how humans understand language. Despite some very interesting proposals, linguistics is not quite there yet. Further, the computer programs that implement at least part of my computational theory rely on both induction, that is, having the system learn generalizations based on statistics collected from large bodies of linguistic examples, as well as rules. How rules and statistics interact in communication and language acquisition is far from clear.[4]

The computational approach also includes various modules that carry out arithmetic calculations related to calendars and that make logical deductions. It is extremely unlikely that the particular calculations, statistics, and logics involved in a model are actually those implemented in the circuits of the brain. Some of the evidence for specifics of time processing in the brain supports the idea of a system of circuits that oscillate periodically. My computational model, which does not even adopt the interesting biologically oriented perspective of cognitive linguistics, is far removed from any grounding in such biological mechanisms. This lack of grounding, while hardly uncontroversial, is nevertheless quite common in functionalist approaches to the philosophy of mind, where the mind is viewed as an abstract symbol-processing device; such a functionalism has been a strong influence in research in computational linguistics. The theories of cognitive linguistics, in turn, do not see much of a role for logic in representing meaning and, further, are somewhat sketchy when it comes to analyzing the details of how to represent time.

However, a weaker psychological claim can be made. Aspects of this computational model underlie some of the representations humans use in narrative understanding, as uncovered by psychological experiments. For example, psychologists have shown that in reading, humans carry out a simulation of the events in a story, modeling the characters, their events, and their passage in time, while experiencing feelings of empathy, repulsion, and so on at various points. Neuroscientists have provided further evidence for such simulation in experiments that reveal that the same areas of the brain are activated when reading about an action as would be activated when carrying out that action. These simulations have many points of correspondence with the computational model I propose. Studies of how children generate narrative show a progression from jumbled thoughts to simple timelines and eventually to plots and character models. This ordering supports my claim of timelines being more basic than plot and character. Further, it is simulation that contributes, in my theory, to the perceived tempo of events in narrative. Finally, my theory makes several predictions about narrative that can be tested experimentally. Even though much can be learned from human models, I do not believe it is essential for an artificial model to resemble the human one.

DEFENDING THE LINGUISTIC PERSPECTIVE

The focus in this book is on language rather than on extralinguistic factors such as the context of creation of a work, the time in which it was written, its relationship to writer and reader, its relation to prior literature, references to salient historical events or cultural allusions, and so forth. These extralinguistic factors, beginning with the reader, are of the utmost importance. Theorists beginning with Louise Rosenblatt in the 1930s have drawn attention to the importance of the cultural presuppositions that the reader brings to a text, allowing for a variety of meanings among readers. Postmodern critics such as Roland Barthes (1977) have taken this view to an extreme, minimizing the importance of the author and his or her biography to the understanding of a literary work, emphasizing exclusively the work as a linguistic object and the role of the reader in interpretation. Barthes buttresses his argument with an account of how writers like Mallarmé, Valery, and Proust went to great pains to blur the role of the writer, so much so that the narrator in Proust's magnum opus is a young man who has not yet written the book that is being read. The linguistic perspective Barthes advocates is in fact consistent with the view adopted in this book.

Barthes' argument is in keeping with the postmodern agenda of celebrating the subjectivity of interpretation. In principle two works created in different periods with exactly the same "text" might be viewed as quite different, a point made with effervescent wit by Borges in his story of Pierre Menard, the author who re-creates Don Quixote word for word: "He dedicated his conscience and nightly studies to the repetition of a pre-existing book in a foreign tongue. The number of rough drafts kept on increasing; he tenaciously made corrections and tore up thousands of manuscript pages" (1963, 54).

However, as mentioned earlier, I do require that interpretation be not entirely subjective to the point where anything goes; language does substantially constrain interpretation, in specific ways that are addressed in this book.

No matter what Barthes and other postmodernists might claim, situating a work in terms of the background of the writer is sometimes crucial. Reading *The Sound and the Fury* after perusing Faulkner's letters makes all the difference to its appreciation. Knowing the author's background is essential for understanding stories from cultures that the reader is utterly unfamiliar with. Of course, reading a book with hindsight makes a tremendous difference; imagine not knowing about the events of the 1930s and the Second World War when reading the repulsive *Mein Kampf* (first published in 1925).

Computer modeling capabilities today cannot satisfactorily address many of the extralinguistic contextual factors. However, the techniques of network analysis, in particular the computational analysis of various kinds of links (social or other) between individuals, can help build an emergent model of some of these factors. Such an analysis has been put to use in the tracking of academic citations (the counts of which are all too influential in measuring scholarly achievement and granting academic tenure); it is also used in the Page Rank algorithm of search engines like Google. With the growth of social networking sites on the Web, this area has also drawn the interests of governments. Given the history of such investments, and given the various U.S. and European government funding initiatives, new technologies relating to the contexts in which narratives (e-mails, calls to arms, and the like) are interpreted are likely to follow.

Turning away from the extralinguistic to the linguistic: narrative does essentially and inescapably involve language, and the ability to model textually provided linguistic aspects is clearly a necessary step. The modeling of narrative structure must address the popular languages of today as well as those endangered languages in which many of the stories that form our collective cultural patrimony reside. One disadvantage of focusing on language, however, is that linguistic constructions found in one particular version of a work may not be found in other versions, leading to inferences that may be invalid across versions.

Folktales are notorious in this respect. We find the animal stories from the Sanskrit *Panchatantra*, for example, traveling around the world, being altered as they go along; often there are multiple versions of the same tale with each rendering having a somewhat different treatment. The Greek myths, of course, have many different forms; the story of Demeter and Persephone, for example, has dozens of variants. *The Conference of Birds*, that great twelfth-century Persian poem by Farīd od-Dīn Attār, owes its framing narrative to the *Panchatantra* as well as *The Recital of the Bird*, by the tenth-century Persian physician and philosopher Avicenna (or Ibn Sīnā); a similar story is also seen in Chaucer's *Canterbury Tales*. The Indian epic the *Ramayana* has many different versions, in Epic Sanskrit (the original version), as well as other Indian languages such as Hindi, Tamil, Telegu, and Bengali, and, as it traveled east, in live performances in Thai, Malay, Javanese, Balinese, Burmese, Lao, Khmer, and so on.

In the world of printed books, writers often go on editing and rewriting their work, while printers and publishers try to keep pace, meanwhile hoping for a windfall from a yet another "new" or "critical" edition. Which version of *Hamlet* should we draw our inferences from—the second quarto, which Shakespeare may have seen, or the posthumous but official first folio? Does one prefer the first (1913) edition of Proust's *Du côté de chez Swann* or the more authoritative 1987 Pleiades edition? The problem of "critical" editions becomes much more an issue when one goes back to ancient texts, where manuscripts are often lost, damaged, or copied and reedited by unknown hands.

A similar problem is posed by oral narratives. The classic work of Milman Parry and Albert Lord, described in Lord's *Singer of Tales* (1960), illustrates vividly how certain varieties of epic poetry (from the songs of Serbo-Croatian bards to Homer) are transmitted through a process of creative improvisation by "illiterate" singers, so that stories change in each telling. Later research (most notably by Ruth Finnegan [1992]) has explored the varying degrees of improvisation in oral storytelling across cultures.

Discussions of world literature involve analyses of translations. This can weaken any inferences drawn from the language of a particular translation, since the devices found in a translation are just as likely to have been injected by the translator rather than by the author (as Borges quipped, "The original is unfaithful to the translation"). One way to address this latter problem is to consider multiple translations.

Fortunately, all these problems are faced by the reader as well, and many a reader experiences a much-loved work only through the prism of a particular edition and translation. Readers of English rather than Russian or Spanish should think of (and thank) Constance Garnett when reading *Crime and Punishment* (or Margarshack, or Pevear and Volokhonsky), and they should be eternally grateful to Gregory Rabassa for *One Hundred Years of Solitude*.

AGAINST THE INDETERMINACY OF INTERPRETATION

This emphasis on analyzing a particular version of a work raises the broader question of how different versions of a story relate to one another. Barbara Herrnstein-Smith (1980) has pointed out that the idea of a single version of a story, such as "Cinderella," is problematic when one tries to pin the story down in terms of a unique summary. She cites a folkloric compendium of 345 documented variants of the story of Cinderella! Although I certainly agree that there are likely to be many versions, I find this particular number suspicious. Folkloric classifications—such as the Aarne-Thompson classification—emphasize commonalities across collections of folktales at the expense of delineating differences; they are thus likely to be overinclusive. Also, the problem of a nonunique summary is entirely familiar to me from my work in automatic summarization, where getting humans to agree on a summary is hard. But agreement on summaries of, say, news stories is not random and is usually better than agreement across unrelated material.

The existence of multiple accounts or versions of a narrative (whether hundreds or thousands) does not imply a complete indeterminacy of interpretation, as Herrnstein-Smith appears to have mistakenly inferred.

Although there may not be necessary and sufficient conditions that define what it means for a story to be that of Cinderella, some versions of a story could still be more Cinderellaish than others. Some versions of Cinderella could (and in fact are likely to) agree on common properties much more than those versions compared, say, to versions of "Little Red Riding Hood." For example, if one clusters a set of stories by similarity, using features drawn from an enumeration of all properties contained in any of them, there are likely to be groupings that amount to the de facto Cinderella stories. Such clusters may have a centroid (a statistical average) that corresponds to a "typical" Cinderella story, or there may be a particular exemplar (or exemplars) chosen as a Cinderella "prototype" (prototypes).

Of course, version comparisons presuppose some way of comparing and enumerating properties (concepts). Once one settles on an account of what sorts of concepts one is considering, one can empirically determine the extent of human agreement across versions that will indicate what it means for a pair of stories to have similar sets of concepts. In other words the notion of whether two stories are versions of Cinderella is dependent on one's theory and classification of concepts (of plot, story, etc.) but is hardly indeterminate.

NARRATIVE SCIENCE

My perspective on computational modeling has been influenced by several decades of research in computational linguistics, a field that lies at the intersection of computer science, linguistics, and artificial intelligence (AI), a field richly influenced by developments in cognitive science and philosophy.[5] Another influence on my work comes from structuralist theories of narrative, developed mainly in France in the 1950s and 1960s and subsequently refined. The structuralist accounts (which in turn developed as an extension of earlier work by the Russian formalists in the 1920s) were descriptive, and like traditional models of say, the sound system in a language, provided elaborate classifications based on contrasting features of various kinds. In phonetics we have the *p* pronounced in "Paul" as a voiceless bilabial plosive, whereas the *b* pronounced in "ball" is a voiced bilabial plosive. Similarly, in structuralist narratology not only do we have a first-person versus a third-person narration but also mimesis (representation, or showing) versus diegesis (narration, or telling) — a distinction that originates from Plato — and a homodiegetic narrator (a narrator who is part of the story world, i.e., a narrator as a character) versus a heterodiegetic narrator (where the narrator isn't part of the story, roughly, a disembodied narrator). In some stories the narrator switches from one to the other, as in Nabokov's *Pnin*, where the heterodiegetic narrator suddenly becomes homodiegetic in the final chapter, along with a change from third person to first. Although these classifications were insightful, given the atmosphere in which they were conceived, they did not result in precise models with which one could compute. In this book we not only identify patterns in narrative but also describe programs that can identify those patterns.

Research in computational linguistics began with a Chomskyan perspective, based on rules for grammar that analyzed not only the forms of language, as structuralist linguistics had, but also tried to account for underlying similarities between different forms. Chomsky's earliest grammar formalism was developed in the mid-1950s, at a time when cybernetics and computing were revolutionizing the way scientists thought about systems and symbol processing systems like language. His approach was based on introspection about which linguistic forms were natural (grammatical) versus unnatural (ungrammatical). In parallel the research of the logician Richard Montague at UCLA developed a new view of language based on the idea of systematic correspondences between form and meaning, addressing many classic problems of reference that were a central concern in the philosophy of language and in the development of logic.

The theoretical paradigms developed by Chomsky and Montague were popular in computational linguistics since they were inherently computational and were stated algorithmically in precise and mathematically rigorous ways. They also were very expressive and allowed for the generation of an infinite variety of linguistic forms. Narratology, rooted in structuralist models, has not really made such a transition toward more expressive foundations.

Subsequently, in the early 1990s computational linguistics took a seriously empirical turn, moving away from introspection as a means of generating data to the use of large samples of real language, called corpora. In doing so researchers were taking advantage of the availability of inexpensive high-capacity computer storage and increasing access to vast universes of information. (As we shall see, corpora can be used to automatically induce linguistic structures by statistical methods.) The emphasis here was on forms that actually occurred in real data, as opposed to introspectively defined criteria for grammaticality, or to devoting attention mainly to rare phenomena that happened to illustrate interesting properties of language. The theories themselves were formally restricted in various ways, to allow for more efficient algorithms. Today both form and meaning are analyzed in terms of structure, but statistics as well as logic can be used in the analysis.

This striving for expressiveness as well as empirical validation, which will be explored in this book, is highly applicable to linguistic aspects of narrative. Just as we have come to rely on the ability to construct large-scale scientific models of all kinds of data—biological, meteorological, astronomical, and economic—we now have the computational tools to analyze vast universes of linguistic information, including the variety found on the Web.

The multiple disciplines that contribute to my account have convinced me that there is enough common ground to allow for a firm foundation for a scientific theory of narrative, at least as it impinges on temporal aspects, the focus of this book. Such a theory must develop precise, formalizable concepts (formality being merely a way of stating ideas precisely) and also be able to adequately test hypotheses and predictions against collections of example narratives in a replicable manner, rather than relying on informal, anecdotal observations (which nevertheless may provide a rich motivational background for a theory). Clearly, such a methodological requirement may be viewed variously as reductionist, foolhardy, and so forth, but it makes eminent sense from a computational standpoint.

To rebut one such accusation right away: such a scientific theory is not in any sense a reductionist one. There are many ineffable aspects of literary appreciation that will, by their very nature, defy any attempt at formal theorizing. Ultimately, the qualitative states that correspond to our feelings about a literary work could never be reproduced by a machine, for the simple reason that they are inherently subjective. In fact, it is hard to see how the quality of any mental state (as opposed to its occurrence, correlation, or cause) can be characterized by some particular physical state of a computer. As the philosopher Thomas Nagel pointed out more than thirty years ago in his famous essay "What Is It Like to Be a Bat?":

> We must consider whether any method will permit us to extrapolate to the inner life of the bat from our own case, and if not, what alternative methods there may be for understanding the notion. . . . It will not help to try to imagine that one has webbing on one's arms, which enables one to fly around at dusk and dawn catching insects in one's mouth; that one has very poor vision, and perceives the surrounding world by a system of reflected high-frequency sound signals; and that one spends the day hanging upside down by one's feet in an attic. In so far as I can imagine this (which is not very far), it tells me only what it would be like for me to behave as a bat behaves. . . . Yet if I try to imagine this, I am restricted to the resources of my own mind, and those resources are inadequate to the task. (1974, 439)

The qualities of our mental states do seem sacrosanct, privileged in some respect; our minds have enshrined them with ineffability. It is doubtful if any of the programs described here or anywhere else, today or in the future, can savor the peculiar and particular thrills we experience in reading.

This does not mean, however, that we cannot model key aspects of our emotions and analyze features of a narrative that give rise to sadness rather than joy. It is possible to construct a biologically relevant computational theory of emotion that is entirely symbolic, one that makes specific predictions (in fact, AI researchers have done precisely that), without the computer program having a body and being able to experience any of the emotions we have of fear, anger, happiness, sadness, or disgust.

AI AND NARRATIVE

Since this book is all about artificial models of time in narrative, it is worthwhile reviewing the lessons learned from related AI work, and how my particular approach differs from it.

AI has had a long history of building narrative systems, that is, programs to understand and generate stories. Some of the inspiration for these approaches can be traced back to the 1930s, when the Cambridge psychologist Frederic Bartlett conducted experiments in which he asked subjects to repeat a given story over time. Bartlett found that the story was reconstructed in memory according to what he called “an active organization of past reactions, or of past experiences” (1932, 201), or a schema. The schema imposes a level of abstraction on the information in a story, favoring gists over details and biasing the assimilation of new information toward conformity with the schema, resulting in details irrelevant to the schema being left out.

The AI researcher Roger Schank and his colleagues gave these schemas a rich computational embodiment in terms of scripts, which are structures that represent a sequence of stereotypical activities, such as eating at a restaurant. As described in Schank and Abelson (1977), a restaurant script will involve characters typically found in a restaurant, such as waiters and patrons, and events such as arriving, perusing the menu (and perhaps other patrons), ordering, eating, paying a bill, and so forth, that may be arranged in a sequence. Characters have wants (such as food, i.e., they are hungry), and they know that they can carry out actions (such as paying for the food). These are some of the preconditions that a script requires in order to be applicable. When a script is “executed,”[6] postconditions also result, such as the person’s hunger being satisfied. Every script will have to specify the characters, event sequence, preconditions, and postconditions.

The postulation of scripts as a means of organizing the experience of stories in memory is intuitively appealing.[7] However, for a computer program to use scripts, the programmer must be precise as to the set of scripts and the contents of each one. One of the basic challenges with scripts is determining which script applies in a given situation.[8] Even when a script can be identified for a given story, figuring out what the script should contain isn’t easy. This is because of the large number of commonsense inferences that can be involved. If someone enters a restaurant, it sets up expectations that the person is going to eat there (unless the person is a waiter or an assassin). If the person pays up, one can infer that she has already eaten (unless she is in a place where one pays first). A person may eat with the intent of not paying (this results in an extremely uncomfortable outcome, as anyone who has been on the receiving end of this situation can attest). Of course, the presence of a waiter is only a default: fast-food places usually don't have waiters.

Research in the 1970s and 1980s tried to address the problems with scripts by reorganizing and simplifying them and designing them on an as-needed basis, driven by the particular domain of interest.[9] And scripts alone were soon recognized as insufficient for narrative understanding in AI. Researchers came to the conclusion that characters in a story needed to be modeled along with their goals and plans, schemes that are realized or foiled, with outcomes that matter to them.[10] The TALE-SPIN system of James R. Meehan (1977) instantiated this integrated vision, generating stories in the domain of Aesop's fables. The TALE-SPIN programmer provides the system with a cast of characters and a geographical setting and then gives each character a goal to solve. The goal spawns further "subgoals" involving specific actions, including cooperative or adversarial activities between characters; as the actions are carried out, the state of the world is updated. Despite these capabilities, the stories produced by TALE-SPIN are limited in scope: like bland running commentaries or documentaries, they are mainly a trace of the thoughts and actions of characters.[11] What is notably missing from this character-driven approach to narrative is any overall, global structure introduced by the author.

Later work on story generation has emphasized plot. The UNIVERSE system of M. Lebowitz (1985) is similar to TALE-SPIN but uses authorial goals rather than character-driven goals to structure narratives. UNIVERSE can be initialized with an authorial goal that calls for a couple to divorce, and then, like a meddling god, UNIVERSE will insert relationship obstacles into the narrative. In doing so it may use an existing plot fragment to achieve this goal. UNIVERSE is, however, overly author-driven; actions cannot occur unless they satisfy an authorial goal. Further research on story generation, including story "grammars" and temporal narrative, is discussed in chapter 9 ("Digital Storytelling").

Unfortunately, despite techniques for teaching systems to acquire new knowledge, AI narrative systems have relied on knowledge that is custom built (inserted by the programmer) for each story or type of story. The result is that story understanding (and generation) systems are brittle, unsuccessful in understanding (generating) stories on topics for which they weren't prepared. Also, given the large number of possibilities that can be inferred, the reasoning processes required to carry out commonsense inferences tend to face an explosion of possibilities that result in inefficient computation. Finally, these AI systems have been notorious for a complete lack of attention to issues of evaluation.

Some of the major challenges in story understanding (not time related, I believe) will continue to haunt the field, but I want to make it clear that this book isn't addressing the larger story-understanding problem. Instead, we are squarely focused on time, which can be addressed computationally (using methods from AI and computational linguistics) in a tractable manner.

THE STATE OF IMPLEMENTATION

This book develops a computational theory of time in narrative. The question that arises is how much of the theory has actually been implemented in a computer program.

Given the corpus-based approach to computational linguistics, the best practices followed today, based on numerous successes—discussed later in chapter 7 ("Tracking Narrative Progression")—involve collecting corpora and marking them up with the kinds of information and structures that the computer will be required to produce from that data. For example, a corpus of personal biographies might record the person, the author, the date of the narrative, as well as the chronology of events and the names of people mentioned in each biography. This sort of markup will be carried out in context, so that the computer can "observe" the surrounding text next to each mention of a person or an event. Once an annotation scheme for such markup has been developed, it needs to be tested so that humans can agree on how to mark up a given example. In other words the data on which the computer program will be trained must be of high quality. Usually an annotation editing tool is used to insert tags into each document, in a way that does not damage or alter the original document (otherwise, a sudden fit of sneezing might change "Arms and the man I sing" to "Arms around the man I fling"). The first phase, therefore, is one of data preparation.

Once a sufficiently large corpus has been marked up (a large quantity of data is crucial, especially for statistical approaches), then a variety of more or less off-the-shelf components are trained statistically to produce that markup for any text, along with any custom components that need to be built. The systems then have to be evaluated for accuracy.

The annotated data functions as a source of training material as well as a "gold-standard" data set for evaluation (the corpus is usually split up so that training and testing data are distinct). The problem of evaluating something as fine-grained and detailed as time in narrative becomes considerably simplified when a computer program is used to compare system-generated markup against the gold standard, allowing the time-lining program to be automatically scored for accuracy.

For many problems the challenge is more in the generation of sufficient quantities of high-quality annotated data (assuming that a gold standard can be defined) and less in the discovery of new algorithms.

A variety of annotated corpora and automatic tools now exist for generating timelines. However, given that fiction mining is not a high priority for funding agencies, timelining research worldwide (a highly active field) has focused on "real stuff" rather than corpora of fiction. Thus, we have sources such as newspaper articles, crime and accident reports, medical case histories, meeting discussions, e-mail, summaries of movies, and so forth, in a variety of languages, that have all been annotated with timelines, providing "gold-standard" data sets.

As a result, while there are lots of fiction sources online (Project Gutenberg being one major resource for such collections), no large-scale fiction corpora have been annotated with timelines. Therefore, the best one can do at the moment is to take timelining tools trained on nonfictional data sets and apply them to the rather different material found in fiction. This situation is far from ideal, since statistical systems tend not to thrive when tested on material that is as vastly different, in linguistic terms, as fiction is from news (even if there are storytelling elements common to both). There is, as a result, also no gold standard for fiction evaluation. This book suggests specific steps to address this situation, which of course will include efforts to annotate fiction corpora.

One reason we are not as far along in terms of processing fiction as we should be is, as mentioned earlier, the lack of a strong case being made to funding agencies. The situation was somewhat different in the 1970s, when story understanding and generation was a focus of more funded research. Such research, some of which was discussed earlier in this chapter, was mostly aimed at science for science's sake. Today, however, with the tremendous practical opportunities afforded by text mining, and the flow of dollars, euros, and yuans toward it, the idea of devoting time and energy to processing fiction seems almost quaint. Nevertheless, as we shall see, there is a somewhat livelier interest in the generation of interactive, multimedia stories, driven to some extent by the commercial and military interest in computer games.

THE FUTURE OF NARRATIVE COMPUTING

The fact that a computer could adequately model time in fiction shouldn't be taken to mean that computers might have substantially understood natural language. The focus here is on understanding time in narrative. That understanding is tested by asking the computer to answer questions about when and why events happen and how characters evolve. However, as any teacher knows, answering questions is not a sufficient condition for revealing understanding. The lively body of philosophical debate on the Turing Test, focused on whether a computer's answers to questions can be distinguished from a human's, leads to a similar conclusion.

One way of displaying understanding is, of course, to write an essay, to make a logical (or even, if one is inclined, an utterly Dada or antilogical) argument to tell a story. AI has had a distinguished history of story generation, with computers carrying out tasks such as the production of animal stories, fairy tales, murder mysteries, and Arthurian legends. By confining their output modalities just to text, these generation systems have not availed of the variety of multimedia content that can be synthesized with today's technology. They also do not involve the user.

A less ambitious role for AI in fiction is in machine-assisted authoring, where the system does not attempt to supplant the author's imagination. Already writers who aren't programmers are easily able to create Web content with a few clicks. Makers of home movies are today using digital video-editing tools to create professional-looking content. Just as people today make extensive use of such tools, it is only a matter of time before the average person will have AI-based story-authoring tools available. These authoring tools, which are already being created in laboratories, will allow authors to prepare conventional fiction as well as movies and multimedia stories, at first choosing from a palette of characters from preexisting video clips, and endowing them with sound bites, and eventually drawing from libraries of story lines and preprogrammed behavior elements that can be assembled together. Need a crooked cop? Or just some gumshoe-speak? How about an elegant and decadent aesthete? Such characters abound in fiction, are easily personified, and they all have only a certain number of ways of speaking.

Another important technological trend is the empowerment of the reader. In recent years systems for interactive fiction and interactive drama have started to gain prominence. As the technology of single-player and multi-player computer games starts to infiltrate fiction, authorial omniscience will give way to a view of fiction that is much more a collaboration between writer and reader. Today's interactive fiction and drama allow users not just to pause, rewind, or explore choice points, but also to completely change the story line and the perception of the characters. Clearly, there has to be some way of preserving authorial intent, while allowing for a high degree of interactivity and "player" autonomy.

These new trends face many challenges in terms of preserving the unity and authenticity of literary voice and vision. One way they can contribute to the creation of superior works of fiction is for topnotch "time management" tools to be used. Chapter 9 ("Digital Storytelling") discusses this issue further.

There are many aspects of narrative that we are still a long way from successfully modeling, such as the use of nonliteral language, which is a hallmark of narrative. Metaphor, humor, and irony have eluded computational solutions, even though research since the 1990s has explored small steps in this direction. The models of emotion that have been explored in computational work are still rather primitive, and the particular one I have used is extremely crude. Other researchers have explored richer artificial modes of emoting. Nevertheless, I see a time when computers might aim for higher aesthetic capabilities, such as analyzing the atmosphere of a scene, which is closely related to what some AI researchers call "situation awareness." Perhaps someday systems will even be able to appreciate the notions of "lightness" and "exactitude" that Italo Calvino (1988) identifies as among the enduring literary values.

This book sketches a platform for future research on time in narrative from both a theoretical and a practical perspective. For the interested scholar or budding scientist, I offer a list of topics that I believe are likely to yield considerable fruit in terms of near-term scientific investigation. For the reader who simply wants to enjoy fiction using computer aids, I recommend some tools that might be of use to her for semiautomated analyses of literary texts. All these resources are also available, along with live software demos, on the book's companion Web site. Unlike purely literary theories of narrative, this one comes with real embodiments.

TERMINOLOGY

A word on the often confusing topic of terminology. Since this book aims to be nontechnical, I have tried to keep definitions as simple as possible.

I consider a narrative to be any work of fiction or nonfiction in any medium that purports to tell a story, in the informal sense in which we use the term "story." For purposes of analysis, I focus on written texts, and I use the term "text" to stand for particular instances of these. I emphasize fiction, particularly novels and short stories, though other genres such as dramatic works and poetry, as well as cinema, are included on occasion. Plot is used in the Aristotelian sense, to express causal relations among events engaged in by characters with particular goals. A narrative, as embodied in a given written fictional text such as a particular edition of a novel, has a plot derived from that particular instance. Characters are particular agents in a narrative, whose behaviors can be tracked. Character, in the abstract sense, is the set of personal characteristics, virtues and vices, and dispositions that characters have; it can be something static, assumed, constructed initially, or developed through the narrative. These notions are further elaborated in the text; I commit to specific definitions, relating them to the narratological literature, in the glossary at the end of the book.