

Do Story Agents Use Rocking Chairs?

The Theory and Implementation of One Model for Computational Narrative

Kevin M. Brooks

MIT Media Lab

20 Ames St., Rm. E15-430

Cambridge, MA 02139

(617) 253-9787

brooks@media.mit.edu

<http://www.media.mit.edu/people/~brooks>

ABSTRACT

Narrative structure models are useful tools for understanding how and why narratives of any medium affect an audience's level of participation in their role of story reconstruction and understanding. [14] With the advent of the computer comes the potential to negotiate through many such models and variables within models as a means of generating multiple narratives quickly and semi-autonomously.

The computational narrative model presented in this paper offers one approach to narrative generation, that of splitting the task into three functional areas: 1) defining an abstract narrative structure; 2) defining a collection and organization of story pieces with some representation of their meaning; and 3) a navigational strategy or reasoning through that collection of story pieces. Agent Stories is a story design and presentation environment for non-linear, multiple-point-of-view cinematic stories. It is designed to be placed in the hands of the non-linear story writer to use as a tool to promote structuring and re-writing of non-linear narratives before and as they are realized in audio and video.

KEYWORDS

computational narrative, interactive narrative, hypermedia, entertainment

INTRODUCTION

Writers of stories for both print and screen have a deeply ingrained tendency to construct stories for an audience to experience the finished work in a fixed linear fashion. Although there are starting to be some examples of fixed non-linear multimedia works, viewing a cinematic story must always be linear, as a linear sequence of pictures and sounds conveying some meaning. However, it should be

possible to structure a story non-sequentially for the purpose of providing many different sequential playouts. Computational processes can assist and affect both production and viewing. With this purpose in mind, this paper examines cinematic story construction through the use of computer based storytelling systems. Questions guiding this research are:

- 1) How can computational processes assist in the development and presentation of stories?
- 2) What computational processes can meaningfully affect different presentations of a story, and therefore, different experiences of it?
- 3) How can user input feed into these processes?

BACKGROUND

Traditional Context

The terms *narrative* and *story* have been interpreted differently by different people, depending on their application in various media. The word *story* is used often in common colloquial speech to stand for any explanation of events, people, and/or things. *Story* is also frequently used as a reference to some general or abstract description of a meaningful collection of events, people and/or things. Meaningful, in this case, denotes a causal or temporal linkage between the people, events and things. By stating that a certain event happened either before or after another event, or that a certain person caused an event to take place, a structure is described within which story elements are interconnected. *Story*, then, can be thought of as a system of associations between elements, composed of events, people, and things. The scholar, Seymour Chatman, describes *story* in a similar way:

Story, in my technical sense of the word, exists only at an abstract level; any manifestation already entails the selection and arrangement performed by the discourse as actualized by a given medium. There is no privileged manifestation. [5, p. 37]

Permission to make digital/hard copies of all or part of this material for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copyright is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires specific permission and/or fee.

ACM Multimedia 96, Boston MA USA

© 1996 ACM 0-89791-871-1/96/11 ..\$3.50

To “tell a story” means to choose a medium for the telling and to construct it by imparting some amount of editorial power over all the possible material which could be used. That choice involves what to tell and how to tell it. According to Edward Branigan, herein lie the definitions of the terms *narrative* and *narration*. Branigan characterizes the relationship between narrative and narration as that of the organization of information and the manner of expression for that organization. [2, p. 106] Or put another way, narrative represents the universe of story elements for a given story - the collection of possibility — while narration represents a specific navigation through that universe.

The computer is a tool well suited for decision making. With it comes the possibility of making complex editorial decisions quickly; however, this requires us to provide the computer with directions for doing so. If these directions could be provided, and if the computer could make these decisions quickly, then the computer could quickly construct an entire narrative (made up of many such decisions) for one person and remake those decisions in a different way for another person. A storytelling system is not a magic box which creatively makes up a story when asked (that would be called a parent), but a system of specially stored and organized narrative elements which the computer retrieves and assembles according to some expressed form of narration. The question is, is this a reasonable thing for a computer to do? At what level would the storytelling system play a valuable role? These are the questions which my research attempts to answer.

New Context

When we as members of an audience view a movie, we experience it in a linear form, from title and opening credits to ending credits. We experience the movie’s various characters as they are introduced one after another, and watch them as they develop through various scenes, situations and causally linked events. Our job as audience is one of reconstruction – to fit the movie’s events together into a structure that we think we understand, which will be partially based on events in our own lives. That is, we identify with the movie. The movie experience helps an audience make some sense of the world by expressing life, our own lives, as a simplified linearly structured narrative experience. [4, p. 46] [12]

Yet, when a screenwriter writes a movie script, the linear flow of the finished screenplay is finally achieved only through disciplined nonlinear work. Instead of creating the fully developed movie in all its detail from beginning to end, a screen writer often starts with a single simple abstract idea. Often progress through the story is made from that single central idea outward in a non-sequential fashion, toward both the beginning and the end. Work on a screenplay may proceed from the middle of the story backward to the beginning and then forward toward the end. The screenplay may

first take the form a fleshed-out outline. Subsequently, as the writer creates the script, she or he combs through the work many times, making many writing and editing passes through the entire work, spending focused efforts working on certain parts, and even writing many different versions of some scenes in order to pick just the right one.

While the task of screenwriting falls under the description of preproduction, the production and post-production processes (recording the screenplay to film or video and the editing) are also usually done in a non-sequential fashion. For efficiency of time and money, a film’s scenes are usually shot out of order, and sometimes shot with two film crews working at the same time in different locations. Therefore, while viewing a cinematic story may be linear, with no choices for redirection involved, the creation process for that same story is quite nonlinear and non-sequential.

Throughout the entire production process, decisions are made about the script and production elements of the finished film through many layers of *feedback*. Feedback is that form of communication when work done is reviewed and commented on by someone with a specific point-of-view, skill set, or special interest. The nature of specialty of the feedback source is important because as feedback is collected from many people, it is the progression through the various types of reviewers that affects the nature of the work.

Making a movie is an extremely collaborative process requiring interactive communication throughout. Once the movie is complete, the feedback mechanisms become somewhat less interactive and direct, if they exist at all. Some examples like additional test screenings, estimated box office receipts, actual box office receipts, and measurements like the Nielson ratings for television, offer some forms of feedback. But none of these forms puts the power of changing the narrative or production elements of the movie into the hands of the audience.

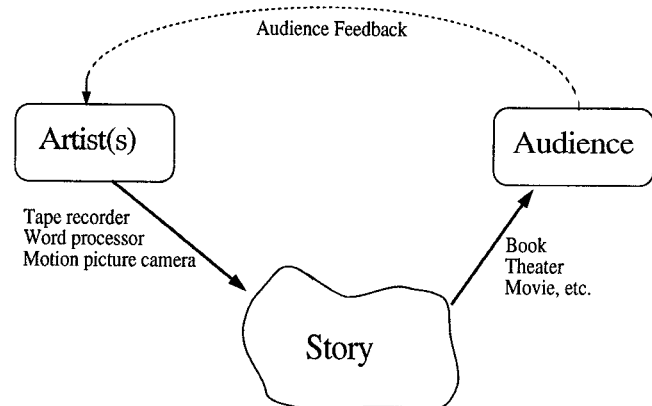


Figure. 1 Simple artist-story-audience structure, with feedback. [10]

Figure 1 shows a typical system of artist, story, and audience — which is not unlike any commercial system involving creator, product, and consumer. In this system, the artist creates the story, which is then produced into a specific form of media and presented to an audience. Once the audience experiences this single instance form of the story (book, theater, movie, etc.), there is sometimes some type of feedback from the audience to the artist. While a book is fixed, with audience feedback possible only after the fact (through the publisher), feedback in the theater is a little more fluid. Actors, for example, can often *read* the audience for needed adjustments during their performance, as well as interpret the level of applause they receive. However, as shown in the diagram, the audience's feedback is usually limited to only after they have experienced the story — especially true in the case of movies. But in general for all media, the audience can not give feedback during the story in a way which will change any of the narrative structure or its production elements.

It is interesting to note that while many companies like Loews, Sony, Time Warner and others are spending lots of money and effort trying to develop interactive story systems for theaters and personal computers, what has proved to be perhaps the most successful “interactive” movie in history allows the audience to effect neither the film’s production nor outcome. Audiences attending screenings of *A Rocky Horror Picture Show* have been packing movie houses for over 20 years to recite the dialogue along with the characters on the screen, throw story-relevant objects during the movie at specific times, and come to the theater dressed as their favorite characters. The audience’s active participation in effect alters their experience of the movie, and people largely come to the theater expecting the standard movie experience to be altered. This phenomenon is one piece of evidence which indicates that interaction in a movie experience does not necessarily require modifying the movie itself. It also suggests that to inspire human interaction with a narrative, while machine intelligence could be one important element, nothing can replace narrative creativity.

The feedback loop in figure 1 illustrates the fundamental property of story; that being the connections between story-teller, story, and story viewer.

Narrative structure is in any case not associated with the short term elementary experience and actions which have served us as examples, but pertains to longer-term or larger-scale sequences of actions, experiences, and human events. ... it can be argued that narrative involves more than just a certain temporal organization of events. To our concept of a narrative

belongs not only a progression of events but also a story-teller and an audience to whom the story is told.
[4, p. 46]

THE COMPUTER’S ROLE

Where does the computer fit into this system of creation, presentation and feedback? Computers have already been incorporated into the story writing and production process

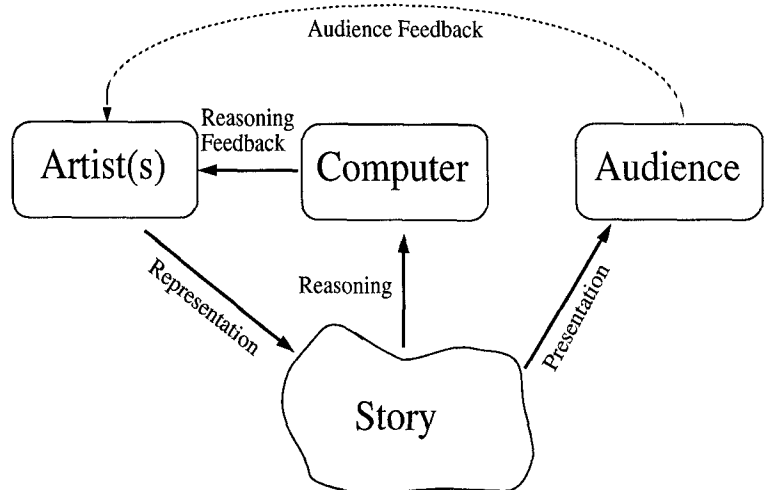


Figure 2 Computer assistance in the simple artist-story-audience structure.

— most notably with the use of word processing tools. Spell checkers and grammar checkers have a representation of words and sentence grammar which allow them to use the written text as input for comparison with their rules. The software offers the writer feedback regarding where its representation of correctly spelled words or proper grammar do not match up with the writer’s text. These actions compose a feedback loop between writer, the story/text, and the computer. However, these systems have little to do with the narration of the story and nothing to do with the audience. That is, the presentation to the audience and the feedback from the audience are not directly affected by the computer’s presence. Spell checkers and grammar checkers do not make the giant bigger, or greener, or Jack a faster runner, but just “clean up” what the writer has done. Figure 2 illustrates such a system.

Figure 2 differs from Figure 1 in that there are paths marked with the words *Representation*, *Presentation*, and *Reasoning*. *Representation* conveys the existence of story content as well as story description. The story is described to the computer in a way which allows the computer to understand it enough to facilitate simple manipulation. [1, p. xiii] There is more on representation later in this paper. *Presentation* conveys that the form of media for the story is not something which is predetermined. That is, the story exists first in a somewhat amorphous or unrealized state. It

is the presentation process which forces a choice of medium on the story. *Reasoning* conveys the existence of something which makes logical inferences about the story based on the description (representation) of the story. The reasoning *engine* reads the story description, makes its inferences, and feeds the results of those inferences back to the artist.

A much more sophisticated example of this same structure would be the software program *Dramatica*, from Screenplay Systems [21]. *Dramatica* provides the writer with a gigantic sophisticated questionnaire. As the writer fills out the questionnaire, making high level narrative choices, *Dramatica* searches through its list of known narrative structures for matches. It's main goal is to force the writer into making the most detailed decisions about the story as possible, such that *Dramatica*'s matches for story structure and style come down to just one. When there is only one match, then the system can provide the user with whatever general information about their screenplay is still undecided, according to the description of the found match.

Dramatica seems well suited as a narrative feedback mechanism, as long as the user's goal is to create a type of story in line with *Dramatica*'s "expertise" — that being linear Hollywood movie narratives. *Dramatica*'s knowledge is stored as static rules about linear screenplay structure and character definition. In artificial intelligence (AI) terms, *Dramatica* takes a *knowledge-based* approach to the problem domain of screenplay structural analysis. It is well known in the AI community that one of the weaknesses of knowledge-based AI (KBAI) is that its structures become brittle when faced with a dynamic problem domain or any problem domain which it was not specifically designed to handle. [11] [15]. An alternative approach would be a behavior-based approach. This approach is described well by Pattie Maes in her paper entitled *Behavior-Based Artificial Intelligence*, where she compares and contrasts the two forms. [15] In it she lists characteristics which typify the knowledge-based and behavior-based approaches. What follows are those characteristics which I find most appropriate in the domain of narrative structure:

A Knowledge Based Approach

- Models isolated and advanced or specialized competences (i.e. medical diagnosis or chess playing). The knowledge-based approach would rather provide "depth" than "width" in its expertise.
- Solves one problem at a time, usually with no time constraints for solving that problem.¹ Also the problem domain is static, remaining unchanged.

¹ This is certainly true in *Dramatica*'s case. In the Hollywood film industry, the writing process is so far

- Usually does not concern the developmental aspect, or how the knowledge structures got there in the first place, or how they change over time. Therefore, the knowledge-based approach does not have to be adaptive.

Alternatively,

A Behavior-Based Approach

- Has multiple integrated competences, such as locomotion or navigation for robots. For stories, these competences may choose conflicts, decide on story resolutions, or, with regards to presentation, decide how to smooth out audio transitions, for example.
- Is a system "situated" in its environment. For robots, this means that they are directly connected to their problem domain through sensors and effectors. For computational narratives, it means that the system navigates an environment of story representation and is "open" in that it is always accepting of user feedback².
- Emphasizes the behavior exhibited by the system rather than the system's knowledge.
- Emphasizes the system's adaptive ability, which means that the system improves over time.

Behavior-Based AI (BBAI), as an alternative to knowledge-based AI, represents a fundamentally different way of thinking about a problem domain. Where the knowledge-based approach makes an a priori attempt to capture the rules for successfully solving or navigating a domain, the behavior-based approach instead relies on a set of lower level competences which are each "experts" at solving one small part of the larger problem domain. BBAI constitutes the theoretical groundwork for the notion of *Autonomous Agents* [16] Autonomous agents are software intelligences that embody the ideas of BBAI. Autonomous agents are typically designed to control some sort of mobile robot, animant, or more recently, maintain certain "personalities" as a member of a MUD (Multi-User Dungeon or Domain) [8] or as a set of behaviors in a software interface. [17] An agent must be

removed from the presentation process that as far as the *Dramatica* software is concerned, the writer has all the time in the world to create and analyze his or her screenplay before production begins. Therefore, *Dramatica*'s reasoning system really does not need to work any faster than what good human interface design dictates.

² That is, the system is open to user feedback whether it is responsive to that feedback or not.

able to maneuver around obstacles without getting stuck in an awkward space or lost in an endless loop, oscillating between multiple goals. They are designed to deal with domains that are not entirely known, where unexpected things can happen. To this end, care must be taken to ensure that there exists a set of competences within the agent which handle the low-level tasks necessary for the operating environment; i.e. stepping, talking, or communicating over a network.

Can BBAI and autonomous agents be applied to story construction? How would one map the elements and functionality of autonomous agents into the story realm? An agent which deals with story material³ would have to operate within an environment of story — that is, the representation of story — where the story’s structural elements make up the *story agent’s* terrain which it would navigate. Navigating the story material means choosing a sequence of story pieces based on the agents particular set of behaviors. Also, with a rich set of competences, an agent would be able to communicate in different ways to both the artist and the audience. Figure 3 shows the next diagrammatic step towards a more fully functioning system. The Computer has been replaced by the Agent, which has the task of reasoning about the story — that is, navigating its representation — and offering feedback to the artist. The nature of that reasoning and the feedback will be addressed later, but the important point here is that the artist and agent are members of a feedback loop where the story can benefit from multiple iterations of change. With each iteration, the agent will

bring to bear its reasoning about the story representation and present its findings back to the artist.

For an intelligent system to improve and evolve over time as Maes suggests, the computer would have to receive feedback itself from somewhere. Figure 3 shows the Agent in the position of presenting the story to the audience as well as receiving feedback from the audience. What this means is that the final narrative is computed or reasoned out by the agent, shown to the audience, and the audience then has the opportunity to offer feedback, not to the artists or producers, but to the agent. The agent is in a position to respond to such feedback by either directly modifying the presentation or passing the feedback to the artist. The agent acts as an *agent of information* between the artist, story, and audience. Its pivotal role touches each part of the narrative-generation process. Herein lies the definition of a computational narrative: a narrative whose story representation, structure, and presentation are so intertwined with the functioning of a computational tool that the nature of the narrative reflects the nature of the tool.

NEW SOLUTION

The goal of my current research project, called *Agent Stories*, is to provide a story design and presentation environment for nonlinear, multiple point-of-view stories. The approach taken with *Agent Stories* is to assemble narratives in either textual or QuickTime movie form by making use of the three components of computational storytelling mentioned earlier:

- 1) The structure of the narrative;
- 2) The collection and organization of story pieces with some representation of their meaning;
- 3) A navigational strategy through that collection of story pieces, with *style* and *purpose* - that is, the narrative construction is a product of deliberate decisions and not random choices.

The hope is that by designing a tool that knows something about the writing process and about what has been written, a symbiotic relationship can develop between writer and writing tool which would foster the process of nonlinear writing.

While this paper focuses on the theoretical foundations of *Agent Stories*, efforts have been made toward implementing such a system. A prototype of Agent Stories has been written using HyperCard for the structural and presentational environments described below. For the representational environment, a tool for describing content bases called *ConArtist* is used, written by Media Lab research assistant

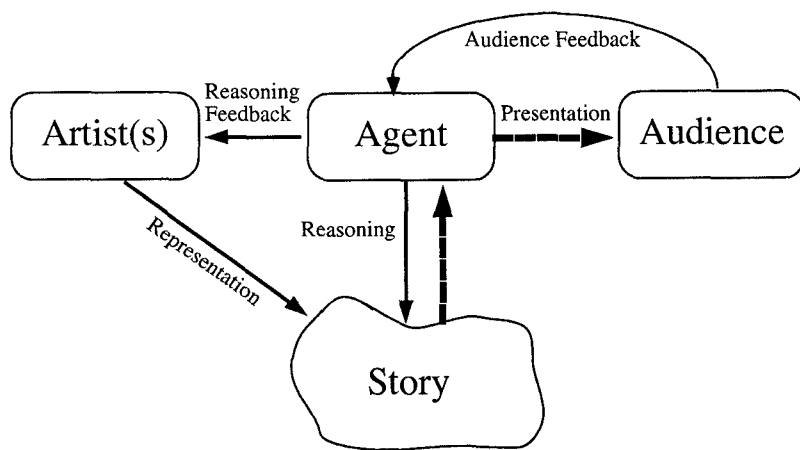


Figure 3 A agent’s role in the creative process. An interim step in the progression toward a fully functioning system diagram.

³ By “story material” I refer first to narrative descriptions of characters and events, as these are the fruits of the writing process. But I also refer to cinematically realized material such as motion picture shots, sounds, etc.

Michael Murtaugh. For content, a collection of stories written by this author entitled *Crossing the Street* is used as the navigated story material. Work is now underway on creating a more functional version of Agent Stories using a more powerful software development platform on the Macintosh and with new or augmented story material.

The Three Environments

Agent Stories tackles the task of narrative construction, or more accurately - narrative sequencing and orchestration - by separating the job into three parts, each with its own construction environment. Each environment maps to one of the three components of computational storytelling mentioned earlier. These environments are called:

- 1) The Structural Environment, in which the structure of the narrative is described in simple abstract terms;
- 2) The Representational Environment, in which knowledge of the various story elements is captured in the form of relationships between story events;
- 3) The Presentational Environment, in which software agents work as text/video editors, intelligently sequencing and orchestrating the different story elements according to an agent's individual stylistic preferences.

Agent Stories allows a story designer to create a simple structure or *framework* for a story and then use that framework to create multiple narratives from the same collection of story elements. The multiple narratives are created when different software agents, each with unique editing/sequencing styles, make clip sequencing decisions in accordance with the story framework, the viewer's preferences, and the existing story material.

In theory, the agents learn to tell better stories in as much as they know what "better" is. There is no way to qualitatively describe to a piece of computer software what a *good* or *bad* story is. In a computer, every element, every decision, is represented by a quantifiable value. How would one even begin to quantify the nature of good and bad? However, by providing a method for agents to quantitatively judge how good a job they have done at fulfilling the requirements of the story structure, as specified by the story designer, then, given the agent's stylistic goals, it should be possible for the system as a whole to move toward narratives which are both coherent and in line with the story designer's vision. It is ultimately left to the human story designer and the audience to make the judgment of good or bad.

The Structural Environment

The Structural Environment introduces the notion of Story Framework. A story framework is a construction of abstract

story element descriptions, referred to as narrative primitives. Such constructions are nothing new. Edward Branigan uses one for film narratives in his writings [2, p. 14]; while Joseph Campbell offers another in his interpretation for hero myths [22]. It is because of Branigan's extensive work with modern cinematic narratives that his narrative scheme is especially relevant here:

- 1 introduction of setting and characters;
- 2 explanation of a state of affairs;
- 3 initiating event;
- 4 emotional response or statement of a goal by the protagonist;
- 5 complicating emotions;
- 6 outcome;
- 7 reactions to outcome

The order of these elements is important, as they progress from the beginning of the narrative through to the end. Elements one and two introduce the narrative, bringing us quickly up to speed with the rules, physical attributes, and even the physics of the environment, as well as the state of the story world and its important characters. Element three, the initiating event, is the spark which sets the affairs of the story world even more off balance than they may have already been. Branigan's element four represents a direct or nearly direct statement by a main character, which focuses the entire narrative around the stated goal of this main character.⁴ Elements five and six are part of a causal relationship stemming from the initiating event, in that the initiating event happened and caused certain emotions and outcomes. Element seven is then part of a causal relationship with element six. Recognizing such causal relationships, or in Branigan's terms, focused causal chains, are important for helping to give the audience a handle on understanding and identifying.

Focused causal chains are not just sequences of paired story events in time and space, but embody a desire for pairing events and the power to make pairs. Narrative causes are thus principles of explanation, or criteria for grouping elements, which are derived from cultural knowledge as well as from physical laws: the human

⁴ No other of Branigan's narrative elements illustrates more clearly that this narrative scheme is specific to stories/movies of the American culture. Many European movies, for example, do not offer the audience a direct statement of the protagonist's goals. Instead, the narrative progression of many non-American movies is based much more on the strength of the characters and character interplay alone. It is as yet unclear if Agent Stories is best suited for American narratives or not.

plans, goals, desires, and routines — realized in action sequences — which are encouraged, tolerated, or proscribed by a community.
[2, p. 116]

In the early part of this century, Russian formalist Vladimir Propp proposed a set of 31 narrative schemes in order to provide a method for understanding and cataloguing Russian fairy tales. [18] Each of his narrative schemes work like functions or classifiable actions taken by characters in the fairy tale. Much of the power behind Propp's work is that it offers detailed patterns of narrative events with a mathematics-like symbol system of representation. However, it is difficult to accurately apply Propp's work to modern narratives, let alone computational narratives, because its form of sequencing is quite rigid.

Similar to Branigan, in Agent Stories the story framework of the structural environment is expressed as set of narrative elements or primitives:

- 1 Speaker Introduction
- 2 Character Introduction
- 3 Conflict
- 4 Resolution
- 5 Diversion
- 6 Ending

While Branigan's schema elements serve as a model for the Agent Stories' structural primitives, the primitives differ from Branigan's elements in order to facilitate the notion of recombination. Branigan's schema is specifically designed to describe linear narratives constructed of specially designed and fine tuned story pieces and not narratives constructed "on the fly" from story pieces designed to serve more than one purpose, that is, more than one narrative.

Agent Stories primitives one and two are similar in scope to Branigan's elements one and two, respectively. Primitives three and four together are generalized versions of Branigan's initiating event. Most initiating events are or include conflicts of some sort. They may be a single conflict, a series of conflicts or conflict/resolution pairs. Primitive four, Resolution, is something like Branigan's elements four and six. A resolution is probably closer to an *outcome* than *complicating emotions*. In the Agent Stories scheme, resolutions are directly tied to conflicts — conflicts can have multiple resolutions, and resolutions, multiple conflicts. Primitive five, diversion, is a story element which deviates or digresses from the plot, which is often driven by conflict/resolution pairs. Diversions act as periods of information transference and tension relief, but do not directly drive the plot. Comic relief is one example of a diversive period in a narrative. The ending is an overall resolution to the narrative.⁵

Each of these narrative primitives describe sections of the writer's intended story. Together, they offer the writer familiar elements for making the narrative "flow" from beginning to end. In the Agent Stories software, the writer builds the framework using simple colored blocks on the screen, as shown in figure 4. These primitive blocks act as class prototypes for the six narrative primitives. When the user clicks and drags a class prototype primitive into the Story Framework, a new instance of that block type is created, numbered, and is then able to be spatially ordered along with other narrative element instances. It is the order




Figure 4 A simple structural framework using Agent Stories.

of these elements which determine much of the flow or narration of a narrative.

The order of these elements also suggest certain narrative genres. For instance, it would make sense to start a narrative with a speaker introduction, so that the audience would immediately have a sense of who is telling the story, followed by a character introduction, during which the characters and setting are introduced, followed by the narrative's

⁵ When showing this project to a person with experience working as both an engineer and a family therapist, his suggestion was to include a seventh primitive: *Negotiation*. While this research is certainly open to discovering new primitives for describing robust narratives, this person's suggestion also brings up two important points: 1) That conflict resolution is often a negotiated hard earned resolution and not one that simply falls from the sky; and 2) That there is a universal quality to narrative structure which allows people to comfortably apply it to non-narrative real life situations.

first conflict. However, if the order of just these three simple narrative elements were rearranged to be: character introduction, conflict, and then speaker introduction, the resulting structure would resemble that of the beginning of a typical murder mystery; where first one sees the characters and setting, then the murder as the first conflict, followed by the introduction of the detective, around whose point of view the story usually revolves. This is a commonly used structure for film and television mystery narratives like: Agatha Christie, Sherlock Holmes, or Colombo, and is recognizable by most western audiences.

Another important attribute of the story's structure is the linkage between conflicts and resolutions. The structural environment provides a way for the writer to specify whether a conflict should be resolved at the next available resolution in the framework or perhaps strung out until a later resolution. By making such adjustments to the narrative structure, it is possible to affect the rhythm of the narrative by either repeatedly introducing and resolving a number of conflicts or introducing many conflicts one after the other so that narrative tension is built to a higher level before any resolutions.

The Representational Environment

How is representation, to which I have alluded, useful in the context of narrative? Representation is here taken from the term *Knowledge Representation* (KR), an area of research within artificial intelligence. In the very beginning of their article entitled "What Is Knowledge Representation," Randall Davis and his colleagues give five roles which KR plays. These roles serve as functional definitions of KR. The three most relevant roles KR plays for narrative are:

- 1) *As a surrogate or substitute for the thing itself; that is, a tool to help one to reason about the world rather than having to take action directly on it.*
- 2) *As a set of ontological commitments, that is, an answer to the question: In what terms should I think about the world.*
- 3) *As a medium of human expression, that is, a language in which we say things about the world.*
[7, p. 25]

As a surrogate or a substitute for the thing itself, a representation of story is something other than the text of the story. For representation, as it is relevant to this project, the text of a story can be part of the representation for that story; that is, the story text can be woven into the system of reasoning describing the story. The important part of the representation is not the story text, but the structures which store infor-

mation about the story. These structures provide the means for deciding what should happen next in the narrative. The better the representation, in that the more authored choices which can be made among well defined elements, the better the resulting narrative. If these choices are made well, then the audience will not just passively sit through the narrative, but actively want to know what happens next. This, according to E. M. Forester, is a story's chief merit — that it makes the audience want to know what happens next. [9, p. 18]

As a set of ontological commitments, a representation of story offers a way for the computer to "think about" or reason about the story elements. For a computer program to be able to read and in any way "understand" story text, it would have to know a broad range of things having to do with people and the world — a difficult if not impossible task.⁶ Instead, by asking the question, *In what terms should I think about the world?*, one is asking for a set of nouns and verbs, or things and functions of things, which can be used to describe the story world.

As a medium of human expression, a story representation offers a second way of understanding a set of story elements. When a writer extends the way he or she understands the fictional story world, they are also extending the expressive potential that they bring to bear on the creation process. This, in turn, adds to the possibilities of expression for the final form of the narrative — a notion which Glorianna Davenport calls *fluidity*:

Perhaps the most important aesthetic of a computational content environment is fluidity – fluidity of expression and fluidity of manipulation devices.
[6, p. 31]

Indeed, even the form of chosen media for a story exerts a form of representation on that story, and thereby, another way of thinking about it. For this form of representation, the system which does the interpreting and reasoning lies in the minds of the audience, as Edward Branigan describes with regard to film:

As a medium, film is a distinctive collection of techniques for representing time, space, and causality on the screen. Normally these techniques

⁶ This approach to computer knowledge, that is, the system storing a lot of connected details about the problem domain, is the basis of an area of artificial intelligence called Expert Systems. Edward Feigenbaum, one of expert system's most articulate proponents, believes that such a system is feasible and could be applied to a wide range of problems. The "only" thing holding such systems back from wide development and acceptance is the vast amounts of knowledge required. [13]

should be understood not as conveying a “meaning” in themselves, but as “instructions” relating to procedures and rules used by a spectator in constructing a set of interrelationships. Such procedures are neither true nor false, but are measured only by their success or failure with respect to some goal.
[2, p. 117]

The goal of the representational environment of Agent Stories is to express a useful and efficient way of intelligently reasoning about the elements in a story domain. In the representational environment, a *clip* is defined as a story element with its information conveyed from a single point of view (POV) and with a single or limited number of narrative meanings. For instance, the meaning (and title) of one clip might be: **Anne decides Michael is a klutz**. But this specific meaning is not literally represented by the system. For the purposes of sequencing this clip in a way which makes sense, it is not important to try to have the computer understand what a *klutz* is or what *decides* means. Instead what is represented is each clip’s relationship to other clips. By defining different types of relationships or links between clips, the interconnected clips become members of a web of story elements, which all relate to one another. In other words, clips are defined in terms of themselves. Examples of links which connect clips together are:

- 1) follows
- 2) precedes
- 3) must include
- 4) supports
- 5) opposes
- 6) conflict<->resolution

Each of these links describe a type of relationship between two clips, and each clip can have many such links. The *follows* and *precedes* links are sequence specifying links meant to identify pairs of clips, where information contained in one absolutely needs to be seen before the other. However, these links do not specify that one clip must immediately be followed or preceded by the other or even that the second clip must be included in the story, but simply that if both clips are chosen, then there is an order in which they must be viewed. The *must include* link specifies that if one clip is chosen, then the other must also be chosen, with no specified order to the clips. The *conflict<->resolution* link specifies that a conflict clip is resolved by a specific resolution clip or clips. Conflicts can have multiple resolutions and resolutions multiple conflicts. The *supports* and *opposes* links offer the system a way of “understanding” to some extent the relationship between the story’s characters by specifying that the meaning or message offered in one character’s clip is in opposition to another character’s clip, or that two conflict clips from different characters are support-

ive of each other. Through this collection of clips and links, a web of story is defined which can be navigated by traveling its links as narrative paths. The navigation of these paths happens in the presentational environment.

The Presentational Environment

The Presentational Environment represents the reasoning portion of *Story Agents*. As the *Structural Environment* offers a framework or guide for building a narrative, and the *Representational Environment* offers a system of linked story elements - the *stuff* of a narrative, it is in the *Presentational Environment* where sequencing choices are made according to the framework. The presentational environment performs the function of Branigan’s narration, in that it chooses and sequences the story elements, presenting them with a sense of *style* through the use of software agents. Maes defines agents as “... computer programs that employ Artificial Intelligence techniques in order to provide assistance to a user dealing with a particular computer application.”⁷ [17] *Story Agents* use this notion of agent because constructing a narrative is more than just “choosing one from column A and two from column B.” There must be some representation of the particulars of narrative construction while at the same time offer the viewer some flexibility, diversity, and ease of use. *Story agents* are the embodiment of the reasoning necessary to construct narrative in this computational model. They perform the sequencing work by making logical choices from among the collection of story pieces. Because each story piece or clip is linked in some way to one or more other clips, there are many different ways in which a story agent can navigate the story web. This is where the agent’s style comes in.

Style is generally defined here as a distinctive and purposeful aesthetic expression in some medium or activity. For narrative construction, having a purpose implies that a certain sequence of narrative elements is assembled because in order to satisfy the over all goal of the one doing the constructing, no other sequence is possible.

There are currently five story agents who do this narrative construction, named: *Bob, Carol, Ted, Alice, and Isadora*. They each have different styles of narrative construction, based on their distinct collection of behaviors. A behavior is a set of rules which describe how the agent should perform in certain situations. The rules match situation or context with some action and alternative action. For example, when given the task of constructing a narrative, each agent first chooses a main point-of-view (POV) character among the various characters represented in the representational environment. The POV character forms the basis of other choices the agent must make later in the narrative construction process. This choice of the POV character makes up

⁷ In the paper cited, Maes and Kozierok apply this definition mainly to Interface Agents, but the point of the definition is still valid.

the first situation for a behavior. Does the agent choose a character with a lot of oppositional links among its story elements? Or instead does the agent choose one with a lot of supportive links? (Note - these are not mutually exclusive.) Does the agent choose a character with the greatest number of over all links to other story events? Or perhaps the agent could choose a character with the greatest number of *precedes* and *follows* links — that is, characters which could be thought of as having clearly labeled sequential paths through their story events. Choosing a POV is just one situation handled by an agent behavior. Other situations would be when the agent must choose a story element according to any of the six narrative primitives of the structural environment. Story agents can have very different styles for constructing a story.

One example of a story agent is Ted, the “point-counterpoint” agent, which attempts to provide an equal amount of conflict and resolution story elements, but from opposing POVs. Ted is a good agent for constructing an entire narrative in the form of a debate. Another example is Carol, the “one-sided story” agent. Carol chooses a main POV character in the beginning of the narrative, and during sections of conflict, will first show the conflict from character POVs opposing the main POV, then show a conflict from the main POV, and finally during sections of resolution, will show only the resolution from the main POV. The narrative generated resembles that of a political commercial, in which POVs opposing the main POV are discredited because conflict elements from opposing POVs are never resolved.

It is intended that eventually the story agents will not just build a narrative in time, that is, not just sequence story events, but also build a narrative in space. The computer screen can be treated as a multi-dimensional stage on which main characters live and struggle through their narrative events. There is no reason to fill the screen with a single stream of visuals if there is no technological or aesthetic reason to be so restrictive. Just as in a movie theater where sounds may come and go, overlap each other, and “appear” in a particular spatial location in the presentation field, there is no reason not to do the same with visuals. Therefore, as part of their set of behaviors, story agents will have additional behaviors which describe where on the screen visual elements appear, how long they will persist, and how many can be on the screen at one time. The expected result could be thought of as a story mosaic or dynamic narrative collage, whose nature or style is dependent on the agent managing the process.

Integration of the Three Environments

The three environments, Structural, Representational, and Presentational, together form a system for organizing and orchestrating the presentation of story elements. The three work together and work off each other to perform the task of

narrative construction. The design of each of them places design requirements on the others.

Representation and reasoning are inextricably and usefully intertwined: A knowledge representation is a theory of intelligent reasoning.
[7, p. 25]

And, the design of the representation, as well as the collective design of the entire system, specifies the ways in which to think about or reason about everything the system represents.

We observed first that every representation embeds at its core a conception of what constitutes intelligent reasoning. Hence, any discussion of representation unavoidably carries along with it a (perhaps implicit) view of intelligent reasoning. [7, p. 29]

The KR of the representational environment offers us a web of represented story elements, all of which are connected in some way to other story elements. The web of connected story elements represents a way of thinking or reasoning about the story domain. The framework of the structural environment offers us a simple structure or guide for constructing a narrative. Together, these two environments go a long way toward defining a system of narrative. But the task of actually building a narrative involves filling out the framework’s requirements by searching through the web of story representation until there are no unknowns or “holes” in the framework.

Figure 5 shows the full system, with the agent occupying the central position. The added story framework is shown contributing to the agent’s input. There is also an additional arrow labeled *Representation* leading from the artist to the Story Framework. It is there to express that the story framework is created by the artist, as is the story. In a sense, the artist feeds the story agent with a framework and a story representation. The agent responds with feedback as to how it would blend both these elements. The reasoning feedback is essentially a description of a narrative for the artist — a constructed story along with a description of the employed reasoning. The artist then may decide which to make changes to: the story framework, the story, or both.

Once a framework and story web have been constructed,⁸ Agent Stories allows an audience user to sit down in front of

⁸ Actually, there is no reason for the story representation to ever be considered fully constructed. The collection of connected story elements can continue to grow and evolve, resulting in an ever changing narrative potential.

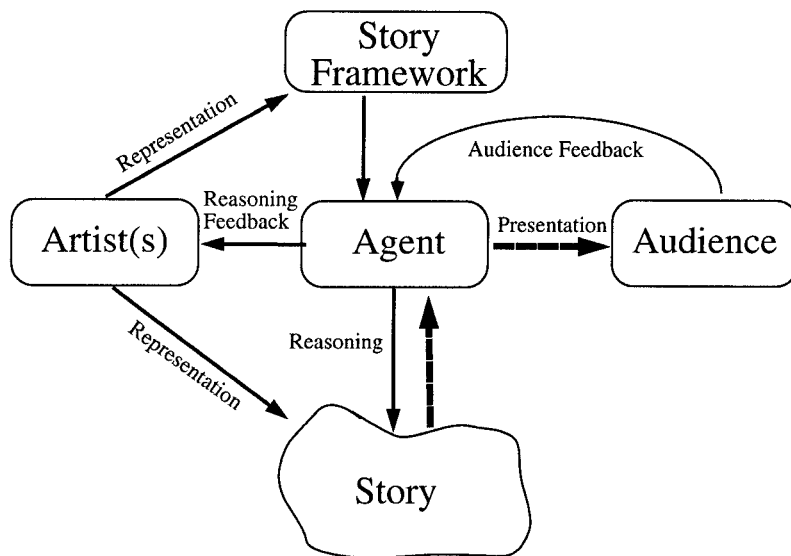


Fig. 5 The story framework, the story representation, and the story agent working together.

a monitor and choose a story agent by name and have that agent create a story play list. To do this, the chosen story agent looks at the characters in the story domain, chooses one as a main point of view character, then weaves a narrative in the method described earlier. Once all the clips have been chosen, the system plays them in sequence.

NEW CHALLENGES

What restrictions or challenges are presented when “conventional” narrative rules are used to construct such stories?

Granularity is here defined as the chosen unit size for building story. With it also comes the balance between power and efficiency: by using smaller story granules, there are more ways in which they can fit together, but more work is required for describing how all these pieces can fit together. The representation and reasoning tasks could easily be prohibitively complex. Conversely, the larger the story granules, the fewer number of ways they can fit together, but the easier it is to put them together. The bigger the story granules, the less reasoning required. A balance or compromise must be struck, keeping in mind the complexity required by the goal. In other words, to build a storytelling system, one needs to ask the question, *How complex does the system have to be in order to tell a good story?* Some systems are designed around large chunks of story and sometimes even use full stories as granules. [19] Agent Stories uses considerably smaller chunks, but there is no right answer.

Another issue of concern in computational narrative is that of the known or unknown nature of the story domain. Branch structures have been a popular means of organizing material for interactive stories. [3] And although they provide a stable organization of story pieces, the problems with branch structures are that they are extremely rigid and the

combinatorics of such structures explode when attempting to provide a robust story experience. The rigidity comes from having to “pre-think” all possibilities of story paths and outcomes with little possibility of adding new additional material at a later time. Agent Stories offers a different approach, that of storing the story elements in a web or semantic network. As mentioned above, the web structure is less rigid because it is much easier to add material to such a structure. The story domain may evolve. The story agents’ job then is to deal with the ever changing complexity of the story domain. However, it is still the job of the writer (or writers) to create each story element and place them in the story representation. This brings up the question, will the web structure of the story domain ever become so complex that it will

be difficult for the writer to comprehend the complete domain? If the writer is unable to hold the entire story domain in his or her head, then it is likely that he or she could be surprised by some of the feedback offered by the story agent. Combinations could be made, old and new material could be combined in ways that the writer may not have thought of. But it seems that as humans, we keep a lot of detail, and some would argue “story detail,” in our minds in normal everyday life. [20] If it is possible for the writer to always retain the story domain as he or she creates new elements, then the story agents have less usefulness. It would seem that only testing this with large collections of story elements would answer this question.

CONCLUSION

In this paper I have outlined a path from more traditional narrative to one model of computational narrative. Computational narrative offers rich potentials for both the writer/creator and the audience. Done well, the processes of creation and viewing/participating could function like experiencing kinetic sculpture — designed from the beginning to be flexible and dynamic, but brought to life only by external interactive forces.

Story agents are not like human storytellers, in the sense that they do not have the same power of control over the narrative they relate. Human storytellers have the capacity to connect as they tell. Part of the storyteller’s special magic is that they make a connection with the audience. For instance in oral storytelling, that connection is made through the teller’s words and the rhythms of voice and body. The storyteller maintains those connections throughout the story, modulating them according to the teller’s sense of the audience’s responses.

So when we listen to storytellers, we are not just listening to the words, but also experiencing that connection between

teller and audience. This is a special kind of connection which I do not think a machine will ever be able to fully emulate. However, as needed research continues, especially in the areas of unobtrusive or naturalistic human-computer interaction methods and narrative database design, I believe that it will be possible for some natural and meaningful connection to be made between a computational storytelling tool and an audience.

Therefore, in answer to the question in the title, *Do story agents use rocking chairs?*, perhaps the best answer is: not yet.

ACKNOWLEDGMENTS

I'd like to thank my advisor Glorianna Davenport for her enthusiastic support of my research and my sometimes unorthodox research style, the other members of my general exam committee: Dr. Janet H. Murray, Ron MacNeil, and Prof. Ken Haas, who have pushed this research and my thinking in directions they needed to be pushed, the students and staff of the Interactive Cinema Group of the MIT Media Lab, Dr. Carol Strohecker for constructively challenging my thinking and the expression of my thinking, and Prof. Ricki Goldman-Segall whose paper listed in the bibliography inspired my thinking about graphical representations of this research.

BIBLIOGRAPHY

1. Brachman, R.J. and Levesque, H.J., ed. *Readings in Knowledge Representation*. 1985, M. Kaufmann Publishers: Los Altos, California.
2. Branigan, E. *Narrative Comprehension and Film*. Sightlines, ed. E. Buscombe. Vol. 1. Routledge, New York, New York, 1992.
3. Bruckman, A. *The Combinatorics of Storytelling: Mystery Train Interactive*. Interactive Cinema Group internal paper, MIT Media Lab (April, 1990). (Available by contacting the Interactive Cinema Group of the MIT Media Lab at the author's address.)
4. Carr, D. *Time, Narrative, and History*. 1 ed. Studies in Phenomenology and Existential Philosophy, ed. J.M. Edie. Vol. 1. Indiana University Press, Bloomington, 1986.
5. Chatman, S. *Story and Discourse - Narrative Structure in Fiction and Film*. 1 ed. Cornell University Press, Ithaca and London, 1978.
6. Davenport, G. Bridging Across Content and Tools. *Computer Graphics*. 28, 1 (1994), 31-32.
7. Davis, R., Shrobe, H., and Szolovits, P., *What is Knowledge Representation?*, in *AI Magazine*. 1993, p. 17-33.
8. Foner, L. *What's An Agent, Anyway?* Agents Memo 93-01 (May, 1993).
9. Forster, E.M. *Aspects of the Novel*. University of Cambridge. Trinity College Clark lectures, 1927, Harcourt, Brace & World, New York, 1954.
10. Goldman-Segall, R. Configurational Validity: A Proposal for Analyzing Ethnographic Multimedia Narratives. *Journal of Educational Multimedia and Hypermedia*. 4, 2 (1995), 163-183.
11. Kolodner, J.L. *Case-based Reasoning*. Morgan Kaufmann, San Mateo, 1993.
12. Kroeber, K. *Retelling/Rereading - The Fate of Storytelling in Modern Times*. Rutgers University Press, New Brunswick, New Jersey, 1992.
13. Lenat, D.B. and Feigenbaum, E.A. On the Thresholds of Knowledge. *Artificial Intelligence*. 47, (1991), 185-250.
14. Livo, N.J. and Rietz, S.A. *Storytelling - Process & Practice*. Libraries Unlimited, Inc., Littleton, Colorado, 1986.
15. Maes, P. Behavior-Based Artificial Intelligence. In *Proceedings of Second Animat Conference on Adaptive Behavior* (Hawaii), 1992, pp. 15.
16. Maes, P., *Situated Agents Can Have Goals*. 1990, Vrije Universiteit Brussel and Massachusetts Institute of Technology: Brussels, Belgium and Cambridge, MA. p. 43.
17. Maes, P. and Kozierok, R. Learning Interface Agents. In *Proceedings of AAAI '93* (Washington, D.C.), 1993, pp. 459-465.
18. Propp, V. *Morphology of the Folktale*. 2 ed. ed. L.A. Wagner. University of Texas Press, Austin, Texas, 1968.
19. Schank, R., et al. Agents in the Story Archive. Technical Report #27, Northwestern University (May 1992).
20. Schank, R.C. and Abelson, R.P. Knowledge and Memory: The Real Story. In *Knowledge and Memory: The Real Story - Advances in Social Cognition*. R.S. Wyer Jr., Editor. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1995.
21. Screenplay-Systems, *Dramatica*. 1994, Screenplay Systems: Burbank, California.
22. Vogler, C. *The Writer's Journey - Mythic Structure for Storytellers & Screenwriters*. Michael Wiese Productions, Studio City, CA, 1992.