

Narrative Meaning Creation in Interactive Storytelling

Joshua Tanenbaum*, Angela Tomizu

School of Interactive Arts and Technology, Simon Fraser University, Surrey, BC, Canada
{joshuat, atomizu}@sfu.ca

Abstract. Many of the difficulties in authoring interactive stories lie in explicitly determining how and where narrative meaning is created. In many ways, these difficulties parallel challenges in intelligent system design, especially in systems that must infer meaning from user interactions. We describe Scarlet Skellern and the Absent Urchins(SSAU); a prototype interactive storybook that utilizes a user-modeling approach which we have implemented experimentally in Flash. SSAU explores issues of narrative meaning creation by mapping user interactions to three distinctive narrative moods, and changing dramatic presentational elements to shift story mood, while leaving the pre-authored elements of the narrative unmodified.

Keywords: interactive storytelling, narrative theory, meaning creation, emergent narrative, user modeling, common-sense heuristics.

1 Introduction

In this article we consider the role of *meaning* in mediated experiences in general, and in Interactive Storytelling in specific. Narrative studies have long been concerned with understanding the process by which stories emerge from a “sequence of events in cause-effect relationship occurring in time and space”. [5] Each sophisticated narrative medium has its own set of conventions and tropes; techniques which have developed specific cultural meanings. Skillful storytellers, be they filmmakers, novelists, playwrights, or game designers employ the conventions specific to their

* Corresponding Author. Email: joshuat@sfu.ca.

medium in order to increase the likelihood that a reader of their work will have the narrative experience they had in mind.

In Intelligent Systems, one of the open questions in the field is how to infer meaning from user interactions with a system. [15] For instance, when a shopper looks at a book on Amazon.com, does that particular interaction represent an interest in the book's author? Is the user interested in books about that particular subject, or in that particular genre? Are they shopping for themselves, or for a friend or family member? In order for the system to make an intelligent recommendation to the user based on this interaction, it must first make some assumptions about the meaning of that interaction. Much like the conventions used to communicate meaning in a narrative medium, intelligent systems require interaction conventions that can reliably increase the probability of *inferring* meaning correctly from interactions.

Interactive Storytelling requires a command of both of the above domains in order to create a satisfying experience. One of the challenges facing designers of intelligent storytelling systems is how to balance what has been described as a "tension between user and author control". [17] [26] Many of the current approaches to interactive storytelling operate on an event or goal based architecture designed to either dynamically assemble a story [21] or adapt an existing series of plot events to a user's interactions. [25] [26] These *plot-centric* approaches put the author and the reader in direct opposition by making the unfolding of the story the subject of a computer mediated conflict between their individual agendas. Solutions to this conflict often come at the expense of the author, who is relegated to a role of storyworld designer. [23] [20] By contrast, *character-centric* systems, such as FearNot! [1] and Friends [6] eliminate the author almost entirely by treating narrative meaning as an *exclusively* emergent phenomenon that arises out of a reader's interpretation of the behaviors of virtual actors. We propose a compromise that allows for the pleasures of an authored narrative while still permitting a satisfying level of responsiveness to user interactions.

In this article we describe *Scarlet Skellern and the Absent Urchins* (SSAU), a prototype interactive story which we have created to explore these questions. Unlike the previous approaches we have described, which are primarily concerned with computationally modifying plot events and character behaviors, our system responds to user interaction by changing the presentational elements of the story. In SSAU, plot and character events are fixed, and instead it is the story *mood* which is variable, using visual and auditory conventions from film, theater, and visual design. *Scarlet Skellern* is implemented in Flash as an interactive comic-book story. Our system employs a simple user model with several experimental "common-sense" heuristics for inferring an *affective preference* from an user interaction, and some simple mechanisms for representing the overall story mood back to the reader.

2 Narrative Meaning Creation and Representation

In order to discuss how meaning emerges in *Scarlet Skellern and the Absent Urchins* we must first

discuss the process by which meaning is constituted in a non-interactive narrative experience. At first glance it is tempting to simply say that an author encodes meaning in a text, which is then decoded—unmodified—by the reader. However, literary theory from the last 100 years has a more interactive approach to this process. Instead of a static phenomenon, narrative meaning can be considered an emergent phenomenon that arises from a reader's interpretation of, and interaction with, a text.

In 1935 Mikhail Bakhtin explored this idea with his essay *Discourse in the Novel*. By framing the experience of reading as dialectic, he set the stage for the study of how a reader's context can affect narrative meaning. Bakhtin introduced the concept of *heteroglossia* for evaluating the different perspectives and conventions of language use that are evident in the experience of reading a novel, both from the author's and the reader's perspectives. [2]

Almost 30 years later, Umberto Eco discussed this phenomenon more explicitly in his 1962 book *The Open Work*. Inspired by the recent rise in experimental and indeterminate musical compositions, Eco went on to describe all artistic experiences as “open”.

“[a reader's] comprehension of the original artifact is always modified by his particular and individual perspective. In fact, the form of the work of art gains its aesthetic validity precisely in proportion to the number of different perspectives from which it can be viewed and understood....A work of art, therefore, is a complete and closed form in its uniqueness as a balanced organic whole, while at the same time constituting an open product on account of its susceptibility to countless different interpretations which do not impinge on its unadulterable specificity.” [12]

Eco's framing of this argument is of particular use to us in that it takes into account the responsibility of the author, who defines the field of possible interpretations for the reader, while not minimizing the role of the reader as a unique interpreter and co-creator of meaning.

By contrast, Roland Barthes makes a significantly stronger claim about narrative meaning. His 1967 *Death of the Author* regards textual meaning as elusive and unassignable. The Author, in Barthes' opinion, is merely a “scriptor”, arranging and re-arranging text from a limitless “ready-formed dictionary” making it impossible to author meaning with language. The sole responsibility for meaning creation lies in the hands of the reader.

“Thus is revealed the total existence of writing: a text is made of multiple writings, drawn from many cultures and entering into mutual relations of dialogue, parody, contestation, but there is one place where this multiplicity is focused and that place is the reader, not, as was hitherto said, the author. The reader is the space on which all the quotations that make up a writing are inscribed without any of them being lost; a text's unity lies not in its origin but in its destination.” [4]

Barthes' argument represents the approach taken by proponents of emergent narrative and *character-centric* approaches to interactive storytelling. What it denies, however, is the possibility of a quality authored narrative experience, or at least the possibility of judging any one narrative artifact in comparison to another. While this view offers the reader greater freedom and responsibility for the narrative experience, it does not account for the difficulties inherent in any creative

endeavor, and it does not provide the basis for any aesthetic or design criteria that might allow designers of interactive narratives insight into how to improve their craft.

More recently, Janet Murray discusses this emergent phenomenon in terms of *belief creation*. In her 1995 book *Hamlet on the Holodeck* Murray identifies the phenomenon of *immersion* as one of the core aesthetic principles of digital media. Immersion has since become one of the most difficult aesthetics to quantify in a mediated experience, and one of the most highly sought after. [4]

“Since the 1970s literary critics have been increasingly aware of the role of the reader in actively shaping the experience of a novel...Some critics focus on the reader’s perception of the formulas of the narrative itself...Others focus on the emotional projections of the reader...Still others emphasize the cognitive activity of the reader in fitting the work of art into existing schemata, or frames of reference.” [20]

Murray’s immersion is rooted in what Coleridge characterized as the “willing suspension of disbelief”, however she augments it, arguing that immersion is not simply a suspension of one faculty, but the exercising of the active faculty of belief creation. She connects immersion directly with the process of emergent meaning creation described above.

The four theorists above describe what amounts to a sliding scale of emergent meaning. At its most extreme, in Barthes’ writing, meaning creation is the sole responsibility of the reader. In the work of the other three, meaning creation is shared between reader and author to varying degrees.

3 Visual Storytelling, Narrative Conventions, and Interaction

3.1 McCloud and Comics

Comic books are one of the storytelling mediums where this interpretative process is most explicitly responsible for creating narrative meaning. Comics creator and semiotician Scott McCloud describes a phenomenon called *closure*, which he defines as “the phenomenon of observing the parts but perceiving the whole.” [19] It is this process that allows us to experience motion and continuity across comic panels, rather than perceiving them as discrete entities. Visual storytelling in comics relies on the reader’s mind to fill in the spaces between panels, and has developed its own symbolic language for specifying how the reader will interpret these spaces which are empty of any tangible narrative content. [19]

Other narrative mediums also rely on closure to suture the transition between isolated narrative events into a whole story within the reader’s mind. Comics, however, rely on it, even in “moment to moment” transitions. [19] Because meaning creation in comics relies on this process of closure more explicitly than film, theatre, or the written word, we have taken techniques from comics as the basis for our own design choices. It is important to note the similarities between the process of

closure and the process of emergent meaning creation described in the previous section: both deal with the cognitive process by which a reader creates narrative meaning out of a narrative artifact.

One of our primary goals with *Scarlett Skellern* was to find an alternative to plot events and character behaviors which would give us some control over narrative meaning, but which could be adapted to respond to reader interactions. We found, within McCloud's most recent book, *Making Comics*, a discussion that led us to explore different presentational techniques. McCloud posits a continuum between two opposing goals in visual storytelling: clarity and intensity. *Clarity* is described as the set of techniques which make reader comprehension their ultimate goal, while *intensity* is the set of visual techniques which add contrast, dynamism, graphic excitement, and a sense of urgency to a panel. [18] He goes on to provide a framework of *design choices* of which a visual storyteller should be aware, in order to better communicate in either of these modes.

McCloud identifies five design choices that influence the amount of *clarity* in visual communication: choice of moment, choice of frame, choice of image, choice of word, and choice of flow. He identifies seven elements for increasing *intensity* in visual communication: depth contrast, graphic contrast, diagonals, extreme poses, 4th wall breaks, frame variation, and surface appeal. [18]

In considering these twelve visual communication techniques, we contend that only four of them—choice of moment, choice of image, and choice of word, and extreme poses—deal with *which events should be portrayed*. The remaining eight are concerned with *how to present* the details of the narrative. For example, a scene with two characters conversing in a restaurant could be made to communicate suspense by employing unsteady angles, and stark contrasting shadows. The same scene could communicate romantic interest by softening the lighting and shadows, and perhaps lingering on pleasing aspects of each character's physiology. Within McCloud's toolkit, narrative meaning is as much a function of presentation as it is of content.

3.2 Acting and Film Theory

Comics are not the only medium that has conventions for visual storytelling which lend themselves to interactive story design. Film and theater also have a rich palette of visual techniques, with their own "grammar". Cinematographer Roy Thompson describes this phenomenon in *The Grammar of the Shot*:

"People 'read' pictures as they do words, and pictures, like sentences, have their own grammar. Once that grammar is changed or omitted, either by design or ignorance, then many elements in the complex process of perception, of reading and understanding, are themselves changed." [28]

SSAU is not the first interactive storytelling system to explore how these conventions can impact story meaning. Magy Seif El-Nasr takes a presentational approach to interactive storytelling in her project *Mirage*. *Mirage* retells the Greek tragedy *Electra* as a 3D interactive drama, using techniques from theater and film to inform the design of its adaptive lighting and camera system.

El-Nasr proceeds from the idea that by modifying aspects of the visual design, it is possible to influence the viewer's psychological and aesthetic perception of a scene, while also steering visual focus to narratively salient features. [14]

In SSAU we draw on this research in addition to the design tools discussed by Scott McCloud to form the basis for our own design choices. *Scarlet Skellern*'s initial implementation employs changes in color and brightness to simulate changes in lighting and time of day, as well as using ambient audio and music to influence the creation of mood. We will discuss these specific design choices in detail in the second half of this article.

3.3 Salen & Zimmerman's "Four Modes" of Interaction

In their 2004 book *Rules of Play*, Salen and Zimmerman describe a "Multivalent Model of Interactivity" with four *modes of interaction*, or levels of engagement, that a person might have with an interactive system. [30] We briefly paraphrase them here.

The first mode—*cognitive interactivity*—occurs primarily in the mind of the user. This is the dialectical interaction between a person and a system. The second mode—*functional interactivity*—occurs at the mechanical, or utilitarian level. This mode includes such literal interactions as page turning, and button clicking. The third mode—*explicit interactivity*—is described as "participation with designed choices or procedures" and is differentiated from *functional interactivity* by its connection to long term effects in the system over time. The final mode—"*beyond-the-object*" *interactivity*—is a cultural form of interaction. This is most commonly seen in fan culture through such venues as fan fiction, and "cosplay". [30]

In video games, designers have traditionally put their efforts into providing users with what Salen and Zimmerman term as *meaningful choices*. [30] This is done through an emphasis of the third mode—*explicit interactivity*. The other three modes do not lend themselves as readily to the creation of meaningful choices in that the first and fourth modes fall outside the control of the designer, while the second mode is purely utilitarian, and is assumed as a given in order to design explicit interactions.

However, returning to our discussion of meaning creation from earlier, we see a parallel between the emergent constitution of meaning described by Bakhtin, Eco, and Barthes, and the mode one *cognitive interactivity* described by Salen and Zimmerman. Both of these approaches place the reader in the position of *interpreter* and treat the act of reading as an interaction which yields a personal and unique narrative experience. However, unlike the texts described by the three earlier theorists, interactive stories are capable of a new dimension of dialectic interaction: they can respond directly to the interactions of the reader. In SSAU we attempt to do exactly that, by engaging the user in what we describe as an "affective dialogue" with the story. Unlike a static novel or comic story, our system attempts to infer the reader's preferences from her interactions with the story, and reciprocate by slowly changing to reflect them. This serves to extend the process of *cognitive interaction* into the realm previously reserved for *functional* and *explicit interactions*.

All of the above discussion deals primarily with how media studies, literary theory, and game studies treat meaning in a narrative experience. It has been primarily concerned with how a reader can co-create meaning through the act of interacting with a narrative artifact. In the following section we look at the difficulties surrounding the process of computationally inferring meaning from a reader's interactions with a system.

4 User Modeling

One of the areas where intelligent systems are most concerned with the task of inferring meaning from user interactions is in the growing field of user modeling. User modeling is a young discipline with research that bridges Human-Computer Interaction and Artificial Intelligence. Gerhard Fischer provides a general overview of the development of user modeling for HCI. He defines user models as “the models that systems have of users that reside inside a computational environment”. [15] In order to make a case for the value of user modeling techniques, Fischer describes what he calls an *implicit communication channel* “which supports communication processes that require the computer be provided with a considerable body of knowledge [about the domain in which the interaction occurs]”. [15] HCI struggles with how to achieve this type of communication in a way which will provide the system with meaningful data about the user's interactions. Building a computational model of the user is one proposed solution to this problem, however determining how and what to model is no small problem. Fischer describes several techniques for finding out what the “user really knows and does” including soliciting direct feedback from the user (in the form of questionnaires, preference settings or configuration modifications), inferring meaning from user actions, and providing the system with additional information about the external context in which it is embedded. [15]

4.1 Meaning Inference

Of these three techniques, the second is both the most problematic and the most promising for interactive narrative applications. An interactive storytelling engine that is able to infer from a user's interaction what their knowledge of the subject is, what they like and dislike, and what mood they are in, has the potential to deliver a story that either compliments or challenges their experience of the media in a profound way. An interactive story that is able to successfully infer from a reader's interactions some understanding of how they are constructing narrative meaning—how they are interpreting the story—offers new possibilities for personalizing and customizing the experience, beyond the simple decision tree models found in many branching interactive stories.

In the final chapter of *Human-Machine Reconfigurations*, Suchman discusses issues of human-machine communication and provides insight into how to better understand this process of interpretation. She frames human to human communication as a highly contingent process rooted in an interactive exchange of action and interpretation. The significance of any action lies only

interactive exchange of action and interpretation. The significance of any action lies only partly in the action itself. “Every action assumes not only the intent of the actor, but also the interpretive work of the other in determining its significance.” [27] The process of interpreting actions is often invisible in human interaction, however, in designing interactions for humans and computers this process cannot be taken for granted. Suchman argues that for a machine system to make meaningful interpretations of human action, it must have an understanding of the context in which the action is taken. She presents two alternative perspectives on face-to-face interaction. In the first, the success of the interaction relies on each participant being able to anticipate the actions of the other based on a preconceived model of their possible responses. In the second, interactional success is rooted in being able to appropriately respond to the *unanticipated* actions of the other interactor. These perspectives can be applied to human-machine interaction. The first requires designers to anticipate possible user action at design time, while the second requires systems to adapt to user behavior at run time. [27]

Both of these approaches correspond to different types of user models. Suchman’s first approach describes a *stereotyping user model*, which is a model comprised of several templates for possible user types which are used to pigeonhole the interactor. The second perspective is similar to an *adaptive user model*, which builds a unique model for each user, by recording and storing their pattern of responses. This model type does not require any “typing” of the user, instead relying on the interpretive strength of whatever algorithm is in place to decode the model to the rest of the system. Each approach has its own strengths and weaknesses. *Stereotyping models* are capable of more carefully crafted responses, because the potential responses are limited to the number of user types recognized by the system, however they do not perform well when confronted with a user who does not conform to a known type. *Adaptive models* are capable of greater personalization and flexibility; however they run the risk of incoherence and confusion, especially if the user does not display a clear interaction preference.

4.2 User Models in SSAU

The user model in *Scarlet Skellern* has elements of both types of the above described approaches, but it is primarily an *adaptive model*. There is a set of three storage variables, corresponding to the three different extremes of mood that the system can measure and represent. There is a second set of variables for tracking raw interaction counts, and there are three exploratory common-sense heuristics, which are used to determine how much a given interaction weighs the model. There is also an algorithm for translating the model values into parameters which control three different modalities of feedback on the “stage”.

The three different heuristics currently implemented in SSAU are a pre-weighted heuristic, an inertial heuristic, and a choice order heuristic. These are represented to the reader in the form of two different types of utilitarian interaction which we term as Major choices, and Minor choices. Major choices are mandatory, non-repeatable interactions which are required of the reader in order

to move forward with the story. Minor choices are repeatable, optional interactions which are “hidden” in the environment of the story. Major choices are pre-weighted at design time, whereas Minor choices use the inertial and choice order heuristics to determine their values. [Figure 1]

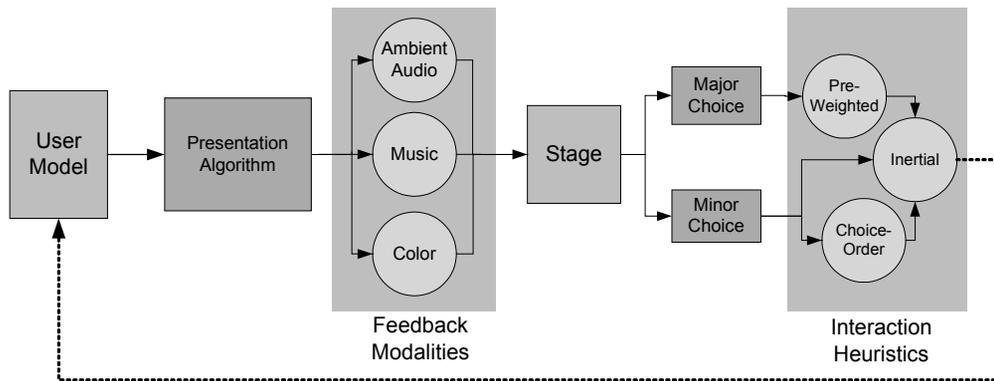


Fig. 1. SSAU system architecture

4.2.1 Major Choices

Major choices present themselves as events within the story, and require reader action in order to move the story along. They are always presented in sets of three, with each one reflecting one of the three possible story moods. For example, at the beginning of the story, the reader is asked “What music am I in the mood for?” She is presented with three possible soundtracks, each corresponding with one of the three possible story moods. Upon selecting which song to play, the story continues, and the soundtrack begins in the background. There are five major choices over the course of *SSAU*. These include a choice of which music to listen to and which food to eat in the first scene, where to search for evidence and who to ask for assistance in the second scene, and a choice of dialogue possibilities in the third scene.

4.2.2 Minor Choices

Minor choices, on the other hand, are embedded within the environment of the story, and are optional. The user is invited to explore the page for these minor choices. When a user encounters a minor choice, they are rewarded with a small bit of meta-information about the setting of the story, in the form of a small animation, a bit of text, or a sound effect. There are currently twenty possible minor choices within the first two scenes, which prevents us from listing them all here. To give an example of one of them, in Scene Two much of the action takes place within Scarlet’s home, which is decorated with a number of objects. On the windowsill are three flowers: a tulip, a wilted dead flower, and a venus flytrap. When the reader clicks on the wilted dead flower, they are rewarded with the text: “Sometimes my happy flowers don’t survive the hostile conditions brought

about by my experiments. I really should throw this one away.” Clicking on the other flowers will also reveal information about the character and the world.

4.2.3 Pre-Weighted Heuristic

The simplest heuristic for evaluating user-intent in SSAU is the pre-weighted model, which is implemented in the five major choices. Pre-weighted choices increment the user model by an amount determined at design-time, while also being counted on the set of variables that tracks the raw interaction counts. Because the major choices are mandatory and appear only in triads, we know that over the course of the reading, the reader will interact with each of them, and we know the context of that interaction (location along the narrative arc, opacity of associated emotion, etc.) This allows us to use the pre-weighted choices as a broad measure of control over the overall state of the model. For instance, if we want the model to very quickly reflect user interactions, we can heavily weigh the early choices. If we want the model to develop more gradually we can weigh them lower. Pre-weighted choices are determined at design-time, and have no mechanism for building confidence, instead relying on the strength of our *a priori* judgments as designers and on our confidence in the clarity with which we represent each choice to the reader. As such they have the greatest potential to misbalance the model, or misrepresent the user’s intent.

4.2.4 Inertial Heuristic

The inertial model, by comparison, is designed to be self correcting over time. This heuristic assumes that the reader will demonstrate interest in a particular mood via repeated patterned interactions with choices related to that mood. If a reader finds and interacts with every minor choice available in the story, the overall relative state of the model remains unchanged. Each choice, in this case, is valued equally. However, if a reader returns to some choices again and again, or if she consistently elects to interact only with choices assigned to a specific mood parameter, the accumulated value of those minor choices will increment that mood in the user model to a greater degree. The more one particular mood is weighted, in this model, the less the other two become worth. This gives the model a certain degree of “stickiness” or “inertia” and prevents random or greedy interactions from being treated as meaningful. The effect of this technique is to serve as noise filter for user interactions. Should the reader choose to regularly interact with happy objects, each minor interaction with a happy object will increase in overall value. If the reader instead interacts in a less directed manner, the value of each minor interaction remains relatively the same. In this way, patterns of interaction distinguish themselves from random or meandering explorations. Minor choices accumulate inertia the more times they are made, and each Major choice also contributes to model inertia.

4.2.5 Choice Order Heuristic

The third technique we have implemented for inferring reader preference is a choice order heuristic. This model works on the assumption that a reader will prioritize interactions in order of most interest to least interest. Unlike the inertial model, which operates on a principle of *accumulation*, this approach operates on a principle of *consumption*. When the reader enters a scene containing new interaction options, the system allocates a pool of “points” to the choice-order model. As the reader makes choices in the scene, the system deducts those points from the pool, and allocates them to the user model. This scales down over time, so that the first interaction in a scene might be worth 10 points, the second worth 8, the third worth 6, and so on until there are no more points to distribute. In our informal testing of the system this technique has shown the most promise of the three; however it currently can become messy when the inertial model is simultaneously active, requiring further study to determine which circumstances are best for each of them.

These three heuristics are built on three different assumptions about how to best infer meaning from user interactions. While we do not yet have conclusive results indicating which technique or combination of techniques is the most effective, preliminary explorations suggest that each one has some potential. In the next section we discuss our approach to affect, emotion and mood in greater detail.

5 Affect, Emotion, and Mood

Our treatment of the terms emotion, affect, and mood is drawn from Rosalind Picard’s 1997 book *Affective Computing*. Picard defines this emerging field as “computing that relates to, arises from, or deliberately influences emotions.” [24] When dealing with emotions, it is important to be clear about the meaning of terms. Picard, for instance, distinguishes between emotion and mood using temporal characteristics. Emotions are considered to be fleeting, lasting only a few minutes at most. Moods, on the other hand, are defined as sustained states of an emotion brought about by repeated instances of that emotion over time. [24] In *SSAU* any given interaction represents an instance of an emotion, whereas the aggregate of these interactions over time in the user model is treated as a mood. We also differentiate between single emotions, and “emotional states”, which we use interchangeably with “mood” with no significant difference intended.

While Picard’s work is primarily concerned with systems that are able to “read” human emotions through facial expressions and biofeedback techniques, we have found plenty of value for our own work in the examples that she gives, and in her discussion of the distinction between emotion, mood, and affect. Picard devotes an entire chapter to a discussion of imagined potential applications for affective computers, many of which parallel what we have attempted to do with *Scarlet Skellern*. Among these are the concept of a computer system as an “affective mirror”, of computers that map emotion to musical preference, and of agents that learn user preferences over time and build user models to better serve their needs. She also discusses systems which might

evaluate a given scene in a film for its affective content, based on framing and lighting conventions. [24]

In order to accomplish any of these goals, it is first necessary to have a system in place that can reliably determine the emotional state of the user. This is a computationally difficult problem that is made no easier by the highly individual and personal nature of human emotion. Because SSAU is intended for online viewing, we have not attempted any biofeedback or facial recognition approaches to this problem, instead relying on user interaction as the only communication channel between the reader and the system. This limitation, combined with the subjectivity of emotional experience, means that when we discuss modeling user affect and representing mood in the system, these ideas are shorthand for a more elusive phenomenon. It might be more accurate to say that our system attempts to gauge an aesthetic preference, and then responds by transforming the *genre* of the narrative artifact to match it. These aesthetic mappings roughly correspond to three moods: a melancholy, a happy, and a scary mood. We make no claims as to quantitative accuracy, or the universality of the mappings we have devised. We only claim to have attempted to implement the three different moods in a uniform and aesthetically consistent way throughout the system.

6 Presentation and Feedback

6.1 SSAU Project Description

We have implemented *SSAU* in Adobe Flash CS3. It is viewable via any standard web browser. Reading the story involves simply navigating forward from panel to panel. Occasionally, the reader will arrive at a panel where she is required to make a choice before the story will move forward. These are the major choices. By investigating any given panel, readers have the possibility of uncovering the additional minor choices embedded in the environment.

The default mood of *SSAU* is one of slight melancholy. Characters are line art paper cutouts, arranged on collaged paper backgrounds. The aesthetic draws heavily on Victorian illustration and children's literature, but seen through a macabre lens. Most of the characters are deformed or broken in some way. Scarlet Skellern herself is a sardonic skeletal girl of indeterminate age, and her sidekick and best friend—Errol—has no arms and towers on two stilt-like striped legs. They are accompanied by Petri, a guinea pig that is kept muzzled and straight jacketed for the protection of his companions and their belongings. Should the reader express no clear affective preference during their reading, the story retains this state of quirky dilapidation. However, if the reader is affectively engaged with the piece, then the tone changes dramatically, altering how these characters and settings might be interpreted by the reader in the process.

To either side of the default state we locate an emotional extreme: on one side we have the darker, scarier story, and on the other we have the lighter, happier story. For simplicity's sake we will refer to these emotional states as "happy", "scary", and "melancholy" over the course of this

paper, with the understanding that these are shorthand for more nuanced moods. Comic book and film convention provide us with a wide palette of techniques for conveying these moods. For our initial prototype we focused on two channels of feedback. The first is audio feedback in the form of soundtrack selection, instrumentation, and ambient sound. The second is visual feedback in the form of color and brightness.

6.2 Audio Feedback

“The most frequent example of mood inducement, however, is choosing forms of entertainment, especially music. We enjoy selecting a recording that affects our mood in a particular way—a piece to lift one’s mood or to console in a state of grief.” [24]

The role of music in creating and sustaining mood is widely acknowledged. Livingston and Brown write: “The ability of a film score to bring about a change in the user’s emotional state is profound.” [16] The functional role of music in multimedia is the subject of much ongoing research. [11] [13] [16] While surveying this is outside the scope of this paper, DiPaola and Arya have coined a particularly useful term for understanding how music and art impact emotions. This term is Affective Communication, which they describe as the aspect of a piece of art or music which causes or expresses different emotional states. [11] We regard the auditory aspects of the project not simply as aesthetic window dressing, but instead as a means by which we express various emotions to the reader, via the mechanism of Affective Communication.

Audio feedback in *SSAU* occurs at two levels: the soundtrack and the ambient environmental noise. The soundtrack elements can further be separated into two primary types of audio feedback, as described by Livingston and Brown. These are event-based audio and dynamic (or adaptive) audio. [16] In the case of *SSAU*, event based musical changes happen at the level of song selection. When the reader reaches a specific moment in the story, she triggers a change in the piece of music that is playing. The current state of the weighted model is used as a basis for determining which new piece of music is selected.

Adaptive musical changes, on the other hand, happen in real time at the inter-song level. Each piece of music is written in a modular fashion, so as to support the addition or subtraction of its constituent melodic and harmonic components depending on the state of the user model. In order to keep these elements simple enough to reliably control while maintaining a consistent musical coherence, we have chosen to limit adaptive musical changes to variations in the instrumentation and orchestration of the piece. This choice was inspired by the classic 1936 Prokofiev piece *Peter and the Wolf*, in which instrumentation choices are tied directly to the narrative and emotive qualities of the story which it accompanies. When the model is in the “happy” state, bright trumpets and horns surge to the foreground and carry the melody. When it is in its “scary” state, *marcato* violins take the lead. In the default state, a mournful solo english horn carries the tune. Together these changes of instrumentation and song choice combine to form a wide range of potential musical moments in *SSAU*.

We have taken a similar approach to the ambient audio elements. The story setting determines which particular ambient environment is selected, but the user model determines the mood of that environment. This is done by layering many separate tracks of audio. The modularity of the ambient sounds allows us to parameterize them and map them to different levels of the moods represented in the model. For example, if the “scary” mood is moderately weighted, then nighttime sounds, like crickets chirping and owls hooting become audible. If the reader further weights “scary”, then it begins to rain, and the wind begins to blow. If they continue to push the model towards scary then wolves begin to howl, babies cry, and the rain becomes a thunderstorm. Each of these elements is independently controlled by the state of the user model.

6.3 Visual Feedback

The field of Visual Arts considers color to be one of the most effective and direct means by which an artist may influence the mood or emotional state of the viewer. [22] [29] One of the primary reasons for this is that people have a universally immediate reaction to color without the need to rationalize how we are supposed to feel about it, thereby rendering it one of the most expressive elements of visual design. [22]

Table 1. Conover’s positive and negative color associations [10]

	Positive	Negative
Black	accomplished and worldly	desolate
Blue	secure and peaceful	depressing
Brown	dependable and logical	plain and boring
Yellow	happy and sunny	show-off
Gray	secure and calm	plain and colorless
Green	calm and natural	jealous
Red	power and excitement	aggressive
Pink	sweet and soft	femininity
Tan	calm and natural	ordinary

While there is not a general consensus as to why this phenomenon occurs, there is however, a “general assumption among behaviorists that many people have similar reactions when exposed to a specific color”. [29] For example, it is generally accepted that cool, dark colors have a depressive influence on most people, while warm, bright colors instill feelings of joy and happiness. As early as the 19th Century Goethe wrote, “We also experience a very warm and cozy impression with yellow” in reference to the feelings engendered in people through their phenomenological experiences with color. [8] As further stated by Chang et al. in *Example-Based Color Transformation of Image and Video Using Basic Color Categories*, the mood of an image can be controlled by managing its saturation, brightness and hue:

“For example, if a user changes the color of a forest to green-yellow, we may perceive the forest as being vivid. On the other hand, if it is changed to dark-olive-green, we may perceive it as being calm.” [7]

It is these principles upon which we designed the color component of the feedback systems for SSAU.



Fig. 2. Three moods in *Scarlet Skellern* from left to right: Happy, Melancholy, Scary

In SSAU, successive choices mapped to “happy” cause an incremental increase primarily in the yellow and, to a lesser degree, red color values creating an impression of a bright, sunny day. Choices mapped to “scary” increase blue values, while simultaneously darkening the image, creating an impression of encroaching darkness. No choice at all, or an even balance of choices mapped to both “happy” and “scary” result in a sustained state of “melancholy” which is represented using muted color values giving the impression of a rainy afternoon/twilight [Figure 2].

6.4 Changing Meanings in *Scarlet Skellern*

One of the largest authoring challenges in *Scarlet Skellern* has been constructing the static narrative content in such a way as to support multiple possible interpretations and emergent meanings without rendering it “generic” or “bland”. We have taken two complimentary approaches to this problem. The first has been to remain agnostic about which reading of the story is the “correct” one. The second approach has been to encode the conflict between the possible moods into the characters themselves.

The character of Scarlet is our dark character. She is cynical, sarcastic, and morbid. The character of Errol is our lighthearted character. He is naïve, optimistic, and cheerful. Over the course of the story both characters express their conflicting perspectives through dialogue, set against the backdrops which are changing in response to the reader’s interactions. What happens, as a result of this, is that as the presentational context shifts, one character or the other’s perspective is privileged as the “correct view”. For instance, when Errol says: “Ahh, smell that sweet summer air. What a beautiful day for an adventure!” and the sky behind him is sunny and pleasant we accept what he says at face value. However, when that same line is juxtaposed against a dark gloomy

night sky there is a tension of meaning between the character's speech and our experience of the environment. When Scarlet responds to Errol with her customary cynicism, her perspective is either undermined or reinforced in an inverse relationship to Errol's. [Figure 3]

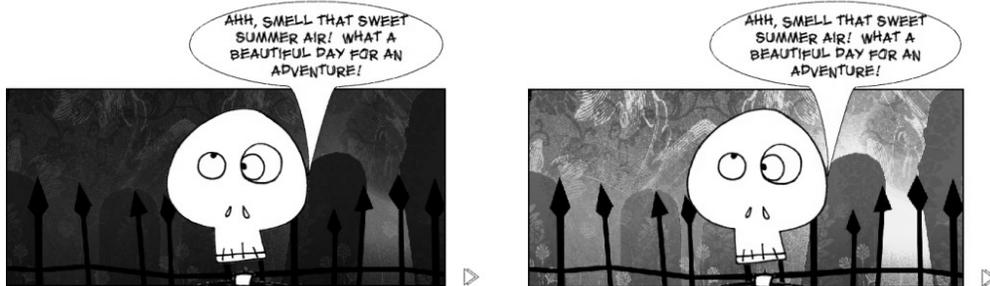


Fig. 3. Context changes meaning

By placing these two character's perspectives in a narrative tension with each other, we create a situation where shifting meanings in the story are a function of *dramatic irony*.

7 Conclusions and Future Research

While there are a number of directions that we can take this work, the next obvious step for this project is a comprehensive reader study.

With *Scarlet Skellern* we have attempted to demonstrate how it is possible to author a story that can simultaneously support adaptive meaning generation that responds to user interactions while also retaining the coherence and dramatic quality of a traditional linear narrative. We have argued that there is much potential for interactive stories that explicitly situate meaning as a function of user interpretation, and story presentation, in addition to plot sequence and character actions. We have also argued for further exploration of user modeling as a technique for personalizing computationally mediated narratives, and described several experimental heuristics for inferring meaning from user interactions. Finally, we have discussed the specifics of how *Scarlet Skellern and the Absent Urchins* uses these techniques to create emergent narrative meaning in the context of affective computing.

References

1. Aylett, R., Louchart, S., Dias, J., Paiva, A. and Vala, M.: FearNot! - An Experiment in Emergent Narrative. *IVA 2005: Intelligent Virtual Agents, 5th International Working Conference*, Kos, Greece. (2005)
2. Bakhtin, M.: Discourse in the Novel. in Holquist, M. ed. *The Dialogic Imagination: Four Essays*, University of Texas Press, Austin and London. (1981)
3. Barthes, R.: Death of the Author. in *Image Music Text*, Hill and Wang, New York. (1977)

4. Bizzocchi, J.: Games and Narrative: An Analytical Framework. in *Canadian Games Studies Association*. (2006)
5. Bordwell, D. and Thompson, K.: *Film Art: An Introduction*. The McGraw Hill Companies, Inc., New York. (1997)
6. Cavazza, M., Charles, F. and Mead, S.J.: Interacting with Virtual Characters in Interactive Storytelling. *AAMAS'02*, ACM, Bologna, Italy. (2002)
7. Chang, Y., Saito, S. and Nakajima, M.: Example-based Color Transformation for Image and Video. *Proceedings of the 3rd International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia*, ACM, Dunedin, New Zealand. (2005)
8. Clark, A.J.: On the Meaning of Color in Early Recollections. *Journal of Individual Psychology* 60(2) 141-154.
9. Conati, C. and Zhou, X.: Inferring User Goals from Personality and Behavior in a Causal Model of User Affect. in *IUI*, (Miami Florida, 2003), ACM.
10. Conover, T.: *Graphic Communications Today*. West Publishing Company, St. Paul, Minnesota. (1995)
11. DiPaola, S. and Arya, A.: Emotional Remapping of Music to Facial Animation. *Proceedings of the 2006 ACM SIGGRAPH symposium on Videogames*, ACM, Boston, Massachusetts. (2006)
12. Eco, U.: The Poetics of Open Work. in *The Open Work*, Harvard University Press, Cambridge. (1989)
13. Eladhari, M., Nieuwdorp, R. and Fridenfolk, M.: The Soundtrack of Your Mind: Mind Music-adaptive Audio for Game Characters. *Proceedings of the 2006 ACM SIGCHI International Conference on Advances in Computer Entertainment Technology*, ACM Press, Hollywood, California. (2006) 54.
14. El-Nasr, M.S.: An Interactive Narrative Architecture Based on Filmmaking Theory. *International Journal on Games and Simulation* 3(1) 29-36.
15. Fischer, G.: User Modeling in Human-Computer Interaction. *User Modeling and User-Adapted Interaction* 11 65-86.
16. Livingstone, S.R. and Brown, A.R.: Dynamic Response: Real-time Adaptation for Music Emotion. *Proceedings of the Second Australasian Conference on Interactive Entertainment, Creativity & Cognition Studios Press, Sydney, Australia*. (2005)
17. Magerko, B. and Laird, J.E.: Mediating the Tension between Plot and Interaction. in *Challenges in Game Artificial Intelligence*, AAAI Workshop Series 108-112.
18. McCloud, S.: *Making Comics: Storytelling Secrets of Comics, Manga and Graphic Novels*. Harper Collins, New York, NY, USA. (2006)
19. McCloud, S. and Martin, M.: *Understanding Comics: The Invisible Art*. Harper Collins, New York, NY. (1993)
20. Murray, J.: *Hamlet on the Holodeck: the Future of Narrative in Cyberspace*. The MIT Press, Cambridge, Massachusetts. (1997)
21. Nelson, M.J., Mateas, M., Roberts, D.L. and Isbell Jr, C.L.: Declarative Optimization-based Drama Management in Interactive Fiction. *IEEE Computer Graphics and Applications* 26(3) 32-41.
22. Ocvirk, O.G., Stinson, R.E., Wigg, P.R., Bone, R.O. and Cayton, D.L.: *Art Fundamentals: Theory & Practice*. WCB Brown & Benchmark, Madison, Wisconsin. (1994)

23. Pearce, C., Wardrip-Fruin, N. and Harrigan, P.: Towards A Game Theory of Game. in Anonymous ed. *First Person: New Media as Story, Performance, and Game*, MIT Press, Cambridge MA. (2004) 143-153.
24. Picard, R.W.: *Affective Computing*. The MIT Press, Cambridge. (1997)
25. Riedl, M.O. and Young, R.M.: From Linear Story Generation to Branching Story Graphs. *IEEE Computer Graphics and Applications* 26(3) 23-31.
26. Steiner, K.E. and Tomkins, J.: Narrative Event Adaptation in Virtual Environments. in *Proceedings of the 9th International Conference on Intelligent User Interfaces*, (Funchal, Madeira, Portugal, 2004), ACM Press, 46-53.
27. Suchman, L.A.: *Human-Machine Reconfigurations: Plans and Situated Actions 2nd Edition*. Cambridge University Press, New York, NY, USA. (2007)
28. Thompson, R.: *The Grammar of the Shot*. Elsevier, Oxford. (1998)
29. Wallschlaeger, C. and Busic-Snyder, C.: *Basic Visual Concepts and Principles*. Wm. C. Brown Publishers, Dubuque, Iowa. (1992)
30. Zimmerman, E.S.K.: *Rules of Play: Game Design Fundamentals*. MIT Press, Cambridge, MA; London, England. (2004)