Daydreaming with Intention: Scalable Blending-Based Imagining and Agency in Generative Interactive Narrative

Jichen Zhu and **D. Fox Harrell**

Digital Media | School of Literature, Communication, and Culture
Georgia Institute of Technology
686 Cherry Street
Atlanta, GA 30332, USA
jichen@gatech.edu, fox.harrell@lcc.gatech.edu

Abstract

This paper presents recent development of a blending-based system that generates interactive narrative output. Besides expanding the GRIOT system to dynamically reconfigure discourse, we introduce a novel construct called a "scale of intentionality." As a cultural production, the first application of these advances is an interactive narrative of a character's daydreamed memories, intertwined within a story world. Our approach is technically grounded in algebraic semiotics, conceptual blending theory, computational narrative research, and is informed by the notion of critical technical practice of AI research.

1. Introduction

In the domains of interactive narrative and gaming, player characters (PCs) are often presented as avatars responding entirely according to the will of the user. The emphasis on user agency impedes a reading of the characters as possessing intentional states. Non-player characters (NPCs), on the other hand, embrace system autonomy and are not subject to user control. Many characters fall into one or the other category. For instance, PCs may convey their boredom by foot-tapping or smile when receiving a power-up. We call attention to the grey area between PCs and NPCs as an expressive tool and as philosophical inquiry regarding AI.

Based on Harrell's existing work on the GRIOT system (Harrell 2007), this paper presents our preliminary results on constructing an interactive system that narrates autonomous daydreaming-like behaviors of a character controllable by the user. Our work explores the aesthetic and critical possibilities enabled by letting the status of the PC vary between user control and system control along several dimensions during narrative interaction or gameplay. The term "character" is used neutrally to denote what may be considered an avatar at some times and a narrated autonomous agent at other times (we use the term narrated because the character is not explicitly modeled but rather is embedded in narrative exposition).

Our goal is two-fold. (1) We explore blending-based creative imagining through the narration of daydreaming.

Conceptual blending is the human ability to dynamically integrate and generate new concepts and an algebraic formalization of this cognitive process serves as the basis for generating narrative content. (2) We highlight variable perception of our system's intentionality as a creative and expressive narrative tool. Our system invokes a scale between narrating highly user-controlled character behaviors directed by the user's actions and desires within the story world and highly autonomous ones that exhibit situated "aboutness" regarding the system's agency within its domain of operation. We call this a "scale of intentionality," a feature also useful as a design tool in applications such as educational "intelligent" tutoring software that balances creative user problem solving with system recommendation.

Our system introduces technical advances that (1) allow discourse structure to be reconfigured and extended dynamically in order to narrate the daydreams of a character, and (2) alter degree of user control over both narration of a character's actions and proportion of extended discourse ("daydreaming") content being generated, as a critical reflection on AI research and cultural narratives. Conceived as a critical technical practice (Agre 1997a, 1997b), and acknowledging cognitive linguistics critiques of computational approaches to cognitive modeling (Evans, Bergen, and Zinken 2006; Lakoff 1999), we view this work not as an attempt to reduce daydreaming or creativity to a formal algorithmic process, but as developing computational constraints based on empirical research on human conceptual blending and generating output, using narrative techniques, that can be perceived and narrated as daydreaming phenomena.

In this paper, we introduce several research approaches and areas of inquiry that influence our theoretical framework in Section 2. We describe the GRIOT system's functionality and recent extensions in Section 3, and provide discussion of how narration of daydreams and intentionality can play central roles in generating computational narrative works in Section 4. Section 5 provides sample output and explanation. We conclude the paper with a discussion of recent insights and future work in Section 6.

2. Theoretical Framework

2.1 Blending-based Formal Concept Generation

Conceptual blending theory builds upon Gilles Fauconnier's mental spaces theory (Fauconnier 1985) and elaborates insights from metaphor theory (Lakoff and Johnson 1980; Lakoff and Turner 1989). It describes the means by which concepts are integrated, guided by "uniform structural and dynamic principles" both unconsciously in everyday thought and in more complex abstract thought such as in literary arts or rhetoric (Fauconnier and Turner 2002). The basic elements of a conceptual integration network are (Grady, Oakley, and Coulson 1999):

1. Input Spaces (the conceptual spaces to be combined)
2. Cross-space mappings (links between analogous elements in different input spaces)
3. The Generic Space (a conceptual space mapped to both of the input spaces that describes shared structure between the input spaces)
4. The Blended Space (the space in which elements from the input spaces are integrated)

Algorithmic concept generation is accomplished using the Alloy computer program for conceptual blending. It uses Joseph Goguen's algebraic semiotic approach to formalize conceptual spaces from conceptual blending theory as algebraic theories with additional structure, morphisms between these structures as structure preserving mappings of the constituents of these conceptual spaces, and the executing of a conceptual blending algorithm upon such structures (Goguen 1998; Harrell 2007). The blending algorithm accepts an "input diagram" consisting of a generic space G, two input spaces I_1 and I_2, and morphisms $G \rightarrow I_1$ and $G \rightarrow I_2$. The algorithm outputs a blended space B, and two morphisms $I_1 \rightarrow B$ and $I_2 \rightarrow B$, integrating conceptual spaces based upon structural principles such as degrees of commutativity, typecasting, and preservation of sorts (types) and constructors from the input spaces in the blended space. A brief, but precise, description of these structural optimality principles can be found in (Goguen and Harrell 2004).

The algebraic semiotics version of blending implements only a subset of Fauconnier and Turner's constraining optimality principles that are structurally determined and amenable to formalization, as opposed to subjective and only humanly determinable. Thus, the Alloy algorithm omits many aspects of conceptual blending, but offers a contribution in that it is not proposed as a cognitive model, but rather an experimental formal tool for precisely representing and testing structural aspects of conceptual blending (Harrell 2007).

2.2 Computational Narrative

The approach to computational narrative taken here follows that articulated in (Harrell 2007) and exemplified by the GRIOT system described below in Section 3. This subsection, adapted from the abstract of (Harrell 2007), serves as a high level overview of this perspective, which emphasizes computational narrative works with the following characteristics: generative content, semantics-based interaction, reconfigurable narrative structure, and strong cognitive and socio-cultural grounding.

A system that can dynamically compose media elements (such as procedural computer graphics, digital video, or text) to result in new media elements can be said to generate content. GRIOT's generativity is enabled by blending-based concept generation as described above.

Semantics-based interaction means here that (1) media elements are structured according to the formalized meaning of their content, and (2) user interaction can affect content of a computational narrative in a way that produces new output that is "meaningfully" constrained by the system's author. More specifically, "meaning" in GRIOT indicates that the author has provided formal descriptions of domains and concepts to either annotate and select or generate media elements and subjective authorial intent.

Meaning can also be reconfigured at the level of narrative discourse. The formal structure of a computational narrative can be dynamically restructured, either according to user interaction, or upon execution of the system as in the case of narrative generation. Discourse structuring is accomplished using an automaton that allows an author to create grammars for narratives with repeating and nested discourse elements, and that accept and process user input. Appropriate discourse structuring helps to maintain causal coherence between generated blends.

Strong cognitive and socio-cultural grounding here implies that meaning is considered to be contextual, dynamic, and embodied. The formalizations used derive from, and respect, cognitive linguistics theories with such notions of meaning. Using semantically based approach, a cultural producer can implement a range of culturally specific or experimental narrative structures.

2.3 Intentionality and AI

Intentionality, described simply as "aboutness" (Agre 1997a; Dennett and Haugeland 1987) or as Searle preliminarily offers "that property of many mental states and events by which they are directed at or about or of objects and states of affairs in the world," (Searle 1983) is a topic of longstanding concern in philosophy. In artificial intelligence the concern of intentionality has taken on several specific forms, particularly regarding whether computer programs can exhibit that the phenomenon called intentionality as exhibited by humans. In this paper, we focus on the issue, as identified by Philip Agre, of the role of narrated intentionality within AI practice (Agre 1997a). Agre argues that narrating a system's functioning as intentional is essential in the practice of implementing AI programs that involve, for instance, the AI keywords "planning" or "learning." The narration of a computational

system's functioning, in part via invoking narratives of human mental phenomena such as of "planning," and "learning," or "creativity," is construed as constitutive of to AI practice at least as much in as algorithmic or knowledge engineering innovations.

Our system exemplifies a special case of an interactive narrative program where output is seen as co-created by an author's data and procedural representations, a user's input, and a user's perception of the dialogic meaningful exchange between herself and the system. Interactive narrative systems are often described as being capable of "telling" stories. Characters in such systems are often said to have "goals" or "plans." The telling of the story is a real world event having to do with the user's construal of the system's performance. A character's behavior is also a matter of a user's construal of system behavior; however it is a behavior that exists in the story world of the narrative. The character is seen as a component embedded within a narrative, often one whose functioning can be perceived independently. In our system, however, the character is not explicitly modeled as an independent agent taking actions in a story world. The narration of the character's behavior is modeled as a part of the overall narration of events. Our unit of generation is the narrative clause, a sentential description of an event, not behaviors.

We call attention to the above because narrating intentionality in a system like ours can refer to either the effective generation of a narrative, or to the perceived functioning of a character within the narrative. In our case, when we discuss a perceived intentionality in our system, we play upon the conflation of these two different levels of perceived user and system behavior. The system may be said to generate a successful narrative regardless of whether or not the character is seen by a user as an avatar that does not exhibit self-directed behavior or as a character that seems to exhibit its own belief or desire driven actions within the story world.

The work described in this paper is meant as a critical reflection on the nature of narrated and perceived intentionality in expressive artificial intelligence related systems. Our modest current hope is to call critical attention to the slippage introduced by utilizing intentional terms to narrative user and system behavior and to use the effects of such narration toward expressive ends. Namely, we construct a system whose behavior is narrated using terms like "daydreaming," yet we also suggest that by conflating the two levels of intentionality (real vs. story world) we can vary the user's perception of intentionality by focusing on narrating, yet calling attention to the character as a proxy through which a user can interact in the story world. As we further refine the project conceptually and technically we hope that it can be used as a critique of the all too common AI pitfall of overpromising results of system behavior based upon describing algorithmic processes using terms whose everyday meanings when referring to human activities rely upon socio-cultural context and/or embodied experience.

2.4 Related work

We situate our work as related to the notion of expressive AI practice (Mateas 2001) in that we incorporate techniques inspired by, and providing a formal notation for, human cognitive process into the cultural production of narrative experiences. Phoebe Sengers's work (Sengers 1998) bridges Agre's theory and expressive AI practice by providing a culturally and physiologically based underpinning for autonomous agents design. She argues that narrative coherence supports a user's/viewer's construal of agent behaviors as intentionally comprehensible.

We are aware of relatively little research on computational models of daydreaming and most of those refer strongly to the psychological research of Singer (Singer 1975). Modeling daydreaming as planning, Erik Mueller's DAYDREAMER Program (Mueller 1990) focused in the domain of interpersonal relationships and common everyday occurrences. Our approach is different in that we do not propose to gain insights into the human cognitive processes of daydreaming through our program. Instead, we view our work as a narration of daydreams and rely on literature theory and other culture production as inspirations. Secondly, we do not find it necessary to view daydreaming as a goal-driven phenomenon such as planning, and do not agree with the reduction of daydreaming to planning and further from human planning to "planning" algorithms. Such techniques, while perhaps useful for generating output that exhibit some mechanical characteristics of human planning in particular contexts, are not our focus here.

A more recent work on visualizing computer generated daydreams with computational animations can be found in (Pérez, Sosa, and Lemaitre 2007). Also opposing to regarding daydreaming as a goal-driven phenomenon, Perez's model generates daydreams as a result of previous experiences and the interaction between characters. In contrast to this work, our eventual emphasis is on narration of daydreams and memories. Our major source of inspiration comes from stream-of-consciousness literature, particularly the novel *Mrs. Dalloway* (Woolf 2002 (1925)).

3. The GRIOT System

Our system draws upon the GRIOT framework of Harrell (Harrell 2007), which identifies, formalizes, and implements an algorithm for structural aspects of conceptual blending with applications to computational narrative. A condensed description of how GRIOT functions follows, details are available in (Harrell 2005). Section 3.2 describes recent extensions to the system that form the technical basis for the results presented in this paper.

3.1 The GRIOT System

The author of a computational narrative or poetic work using GRIOT composes ontologies called theme domains and sets of clauses organized by discourse (usually narrative role) called phrase templates. Each theme domain contains sets of axioms consisting of binary relations between typed constants in the domain and the keywords that can activate it. Phrase templates store granular fragments of the story, including authored story components and variables called "wildcards" that can be replaced by generated content (then said to be "instantiated"). These wildcards are tokens representing where and how generated blending output can be incorporated, such as the choice of theme domains. Both daydream and non-daydream phrase templates are organized by discourse clause types. Theme domains and phrase templates are annotated by keywords that can evoke the domain, or describe the content of the phrase. User input, in the form of keywords, is used to select the conceptual space network from the set of domains.

An automaton called an "event structure machine" (or "discourse structure machine") specified by an author's input discourse structure configures the narrative by controlling the set and order of the phrase templates to instantiated. When a phrase template is selected, its wildcards will be replaced by blended results according to the specified blending constraints and then converted to natural language through mappings called "grammar morphisms." More specifically, the input diagram as described in Section 2.1 is passed to the Alloy conceptual blending algorithm, which generates a corresponding "output diagram," a blended conceptual space and morphisms to it, according to principles that produce "optimal" blends.

3.2 Extended System Framework

This project has required that several aspects of the GRIOT architecture be extended in several modest, but notable ways. These are extensions to the following aspects of GRIOT depicted in Figure 1:

1. Select Input Diagram: a finite state machine called the character state machine can select domains from which the input diagram is constructed from, in order to represent character states such as emotions.
2. Phrase Templates: annotated sets of phrase templates that can dynamically extend the discourse structure have been introduced, in order to incorporate varying degrees of daydreaming at run-time.
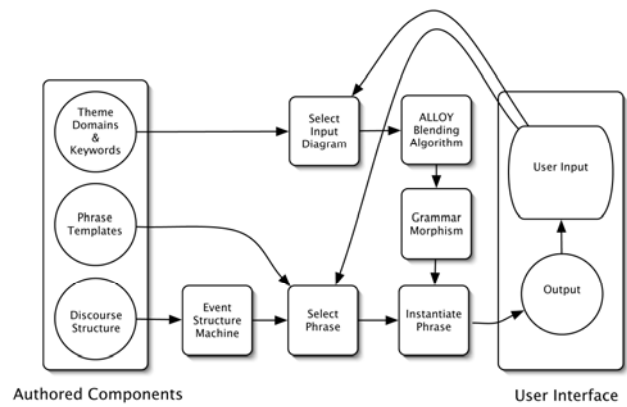3. Phase Templates: we can dynamically vary whether or not a phrase will accept user input.

Figure 1: The GRIOT Architecture

These extensions are meant to reinforce the consistency in the narrative and therefore further invite (or disinvite) intentional readings of the system. The character state machine mentioned as (1) allows blends to be generated based on a sequence of all previous user interaction and author specifications, instead of selecting domains based purely on the immediate previous user inputs, which may be inconsistent narratively. Currently, the user's input affects the character state machine, which in turn determines which domain is selected. This component has so far been used as an "emotion state machine" to influence the affective disposition incorporated into blends in a similar manner as in "Walking Blues Changes Undersea" (Harrell 2006). For example, emotion state machine indicates that the character is in a "`happy`" state and is looking at the "`room`" artifact, axioms are chosen from the "`happy`" and "`room`" domains to form an input diagram, Alloy outputs a blend, and this blend is mapped to a natural language phrase such as "`glossy painted room,`" as compared to an "`old-fashioned room`" that may have been generated if the emotional state had been "`reminiscence.`" This is not cognitive emotion modeling, but merely as a tool to exert influence on output phrases. (2) and (3) above provide the basis for varying the degrees of daydreaming output and the level of control and agency a user has through input at run-time.

4. Daydreaming and Intentionality

Daydreaming can be described as a creative and imaginative human capacity to construct and run imaginative concepts through the mind in parallel with perception and awareness of the real world. Such imaginative experiences can be described in recent cognitive science terms as cases of "mirror network" or "double-scope" conceptual blending, the mental processes of integrating and dynamically executing concepts with shared or clashing underlying frames (Fauconnier and Turner 2002; Turner 2003). Our daydreaming application currently focuses upon narration of a character's past memories in the midst of narrating present action and

perception. This application, discussed in Section 4.1, interactively generates sequences of blends of emotional and artifactual concepts for integration into narration of daydreaming-like experiences, based upon the framework of (Harrell 2007). The possible narrative and blended outputs are influenced by the perceived system intentionality, along a scale that we call the scale of intentionality. Section 4.2 provides details on the role of the scale of intentionality in this daydream narrative system.

4.1 Daydreams

Daydreams in our system are represented as sets of them domains and phrase templates, as the main narrative. They are considered as substructures because their contents are episodic and they can be dynamically integrated into the main narrative. We use "main" narrative here to refer to distinguish the narration of the current story world from the daydreaming output.

If the daydreaming mode is turned on, the system will search for matches between the current phrase template set and the set of daydream discourse structures, each of which is also annotated with related lists of concepts. Upon finding the appropriate daydream based on content correspondence, the event structure machine inserts the daydream sequence after the phrase template to be output. For instance, the "`funny-noise`" triggered by the "`open-door`" action can cause the insertion of a series of phrases organized by a discourse structure for an artifact triggered daydream. In expressive terms this means that the system will narrate the character's daydreaming about another time it heard a similar noise before returning to the main narrative describing the result of the present action of opening the door.

4.2 Scale of Intentionality

Daydreaming mode is triggered depending on the value the "scale of intentionality." While the scale itself is at present a simple numerical variable, the introduction of the concept is based on our critical play regarding the role of intentional narration in AI practice. As exemplified in section 5, its value affects the system output along three dimensions:

1. the proportion of daydreams output relative to the main narrative,
2. the proportion of automatically selected character actions relative to actions selected by user input,
3. the proportion of generated subjective description output (mapped from blends generated by Alloy) relative to pre-scripted phrases objectively describing the results of character actions, the phrases are also organized into subsets of varying levels of descriptive exposition.

5. Results

As a first application of the concepts above, we designed the micro-narrative (a small work focused on causal and temporal coherence, but not larger grained narrative structures such as plot) of a robot named Ales opening a door and entering a room. The ontologies for this micro-narrative currently are comprised of 10 domains (4 artifact domains and 6 emotion domains), each of which consists of 4 to 11 axioms. The discourse structure has 10 clause types, each subdivided into three levels of description. Corresponding to each level of the different clause types is a set of up to four phrase templates. At the time of this writing the micro-narrative application contains 47 main narrative templates and 17 daydream templates.

In this experiment, the value of the scale of intentionality was pre-determined and remained the same throughout the narrative. We ran the experiments with three different levels of scale of intentionality. For the sake of clarity, we have formatted the output text as follows. Input to the system (from a user unless otherwise indicated as described below) is bolded, blending generated text is italicized, and daydreams are underlined. Each clause resulting from an instantiated phrase template is enclosed in parentheses, and all recognizable commands are surrounded by square brackets. Finally, autonomously generated input surrounded by asterisks as well as being bolded.

Example 1 demonstrates the system behavior and its narrative output when the scale of intentionality is "low." In this scenario, the character follows user commands completely and the system output contains a low level of description. In this case, the character acts as transparently as an avatar, its emotional state does not have any effect on the main narrative, and no daydreams are generated.

Example 1:
```
(ales approached the door. in his visual scan
it was painted [red/yellow/blue/grey/brown]
...)
-> red
(.)
(he raised his mechanical arm to
[knock/open/open_gently/punch/open_carelessly]
the door ...)
-> punch
(it opened. )
(ales entered.)
(ales started [examining/glancing/looking]
around ...)
-> glancing
(.)
END
```

With an intermediate level on the scale of intentionality, as in Example 2, the narration of the character's thoughts and actions are both more subjectively depicted (via integration of emotion-artifact blends) and more extensive

(via incorporation of daydreams). The mappings from user input to emotional states are predetermined by an author, for instance, "red" is mapped into the "anger" state. For instance, the "disturbingly familiar mother" description arises from a natural language mapping of a blend resulting integrating the concept of "mother" and the concept of the emotion "anger." We must reiterate that our goal of consistently narrate emotions by selecting correlated domains to blend. Daydreaming output is also introduced, in this example it is triggered taking action on the "door" artifact.

Example 2:
```
("a door!" ales processed, his optics sensed
the color as [red/yellow/blue/grey/brown] ...)
-> red
(. distasteful wood-colored , he thought, just
like the one in mother's bedroom.)
(The room where he had his first encounter of
tune-up and oil change had similar doors.)
(the oil change left a sickly feeling in his
gut)
(he would rust like the tin man before enduring
another.)
(ales stopped, computed, and moved again to
[knock/open/open_gently/punch/open_carelessly]
the door ...)
-> punch
(The door creaked into an entrance.)
(the funny noise reminds him of)
(the hospitals, the junkyards, he went to
frequently)
(, which was so unpleasant that he always
avoided it ever since, ohhh, the same noise.)
(he paused a bit before he walked in.)
(the bothersome light room looked somewhat like
the room he was born.)
(his disturbingly familiar mother probably had
left him quite early.)
(most of his knowledge about her came more
from)
(the engineer pointing to the women in
pictures)
(his memories of experiences.)
(the things she liked to wear, the way she
smiled.)
(ales couldn't help
[examining/glancing/looking] around ...)
-> glancing
(. there was no one inside; nor was there much
furniture, except)
(a couple pieces of furniture.)
END
```

At the highest level of intentionality, the character "takes over the control" from the user and provides input based on its emotion state. In the following Example 3, we set the initial state of the character to "reminiscence" and the character acts in the story world autonomously based on this state: it "decides" to see the door as "blue" through its optical lens, opens the door "gently," etc; it also finds the door "dusty brown" and the room "gloomy white", all of which (we feel) helps the narrative to convey the sense of reminiscence and nostalgia. This emotion consistency is also carried through the daydream narratives.

Example 3:
```
(at the end of a long corridor a door loomed.
ales looked up; he decided to see it as
[red/yellow/blue/grey/brown] ...)
*BLUE*
(. it was dusty wood , almost exactly like the
one in his mother's bedroom.)
(the room where he had his first encounter of
tune-up and oil change had similar doors.)
(the oil change left a sickly feeling in his
gut)
(he would rust like the tin man before enduring
another.)
(he raised his steel arm to
[knock/open/open_gently/punch/open_carelessly]
the door...)
*OPEN_GENTLY*
(the door responded as doors do.)
(the funny noise reminds him of)
(the hospitals, the junkyards, he went to
frequently)
(, which was so unpleasant that he always
avoided it ever since, ohhh, the same noise.)
(he paused a bit before he walked in.)
(the gloomy white room looked somewhat like the
room he was born.)
(his unforgettably familiar mother probably had
left him quite early.)
(he got to know her mainly through)
(reading about her in the books)
(than his own memories.)
(the things she liked to wear, the way she
smiled.)
(ales couldn't help
[examining/glancing/looking] around ...)
*EXAMINING*
(. no one was inside; no things either except)
(a large window.)
END
```

The preliminary results above illustrate how the conceptual blending algorithm can be used in enhancing generativity and subjective description in computational narrative, and how the notion of a scale of intentionality and system agency may be used expressively. The subjective nature of creative processes such as daydreaming makes it difficult to establish objective criteria for evaluating the effectiveness of systems like ours. From our preliminary results, we believe that our generated sequences of blends mostly made narrative sense. Reinforcing emotion consistency, for example, is one of the ways to achieve this. We have not made much progress toward a model for evaluation as of yet, however we are aware of the need a cross-disciplinary approach and that the "success" of expressive work is often a matter of "interpretation" with respect to intertextually related works as opposed to a quantifiable metric for success or failure. The output of the work should be closely read in the tradition of literary theory and should be analyzed for meaningful difference between iterations.

6. Conclusion and Future Work

This paper has presented our initial experiments on generating computational narratives with daydreaming output as a novel feature. We have introduced the concept of a scale of intentionality, inspired by philosophy and critique of AI, as an expressive way to narrate user and system agency and to explore characters that can vary between operating as avatars and narrated agents. In the results section, we presented early achievements as promising initial steps toward our goals. However, the cross-disciplinary and subjective nature of our work requires more nuanced evaluation criteria.

As a future step, we shall perform more robust evaluation or interpretation of imaginative narrative works such as the ones presented in this paper. We have been analyzing stream-of-consciousness literature to inform the dynamic generation analogy-based narrative discourse structures. We also intend to expand work on daydream retrieval and generation. For instance, daydreams can be cascaded to evoke further daydreaming. Finally, we plan to improve our grammar morphisms to achieve more effective natural language output.

References

Agre, P. E. 1997a. *Computation and Human Experience*. Cambridge, U.K.: Cambridge University Press.

Agre, P. E. 1997b. Toward a Critical Technical Practice: Lessons Learned in Trying to Reform Ai. In Bowker, G. C.,Star, S. L.,Turner, W.,Gasser, L., and Bowker, G. eds., *Social Science, Technical Systems, and Cooperative Work: Beyond the Great Divide*: Lawrence Erlbaum Associates. 131-158.

Dennett, D. C., and Haugeland, J. 1987. Intentionality. In Gregory, R. L. ed., *The Oxford Companion to the Mind*. Oxford: Oxford University Press.

Evans, V., Bergen, B. K., and Zinken, J. 2006. The Cognitive Linguistics Enterprise: An Overview. In Evans, V.,Bergen, B. K., and Zinken, J. eds., *The Cognitive Linguistics Reader*. London, U.K.: Equinox Press.

Fauconnier, G. 1985. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge: MIT Press/Bradford Books.

Fauconnier, G., and Turner, M. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.

Goguen, J. 1998. *An Introduction to Algebraic Semiotics, with Applications to User Interface Design.* Paper presented at the Computation for Metaphors, Analogy, and Agents, Yakamatsu, Japan.

Goguen, J., and Harrell, D. F. 2004. *Style as a Choice of Blending Principles.* Paper presented at the Style and Meaning in Language, Art, Music and Design, Proceedings of the 2004 AAAI Fall Symposium, Washington D.C.

Grady, J. E., Oakley, T., and Coulson, S. 1999. Blending and Metaphor. In Steen, G., and Gibbs, R. eds., *Metaphor in Cognitive Linguistics*. Amsterdam: John Benjamins.

Harrell, D. F. 2005. *Shades of Computational Evocation and Meaning: The GRIOT System and Improvisational Poetry Generation.* Paper presented at the 6th Digital Arts and Culture Conference, Copenhagen, Denmark.

Harrell, D. F. 2006. *Walking Blues Changes Undersea: Imaginative Narrative in Interactive Poetry Generation with the GRIOT System.* Paper presented at the AAAI 2006 Workshop in Computational Aesthetics: Artificial Intelligence Approaches to Happiness and Beauty, Boston, MA.

Harrell, D. F. 2007. *Theory and Technology for Computational Narrative: An Approach to Generative and Interactive Narrative with Bases in Algebraic Semiotics and Cognitive Linguistics.* Ph.D. diss., Dept. of Computer Science and Engineering, University of California, San Diego, La Jolla.

Lakoff, G. 1999. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Philosophy*. New York: Basic Book.

Lakoff, G., and Johnson, M. 1980. *Metaphors We Live By*. Chicago, IL: University of Chicago Press.

Lakoff, G., and Turner, M. 1989. *More Than Cool Reason -- a Field Guide to Poetic Metaphor*. Chicago, IL: University of Chicago Press.

Mateas, M. 2001. Expressive AI: A Hybrid Art and Science Practice. *Leonardo, 34*(2): 147-153.

Mueller, E. T. 1990. *Daydreaming in Humans and Machines: A Computer Model of the Stream of Thought*. Norwood, NJ 07648: Ablex Publishing Corporation.

Pérez, R. P. Y., Sosa, R., and Lemaitre, C. 2007. *A Computer Model for Visual-Daydreaming*. Paper presented at the AAAI Fall Symposium on Intelligent Narrative Technologies.

Searle, J. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.

Sengers, P. 1998. *Anti-Boxology: Agent Design in Cultural Context.* , Carnegie Mellon University, Pittsburgh.

Singer, J. L. 1975. *The Inner World of Daydreaming*. New York: Harper and Row.

Turner, M. 2003. Double-Scope Stories. In Herman, D. ed., *Narrative Theory and the Cognitive Sciences*. Stanford, CA: CSLI Publications.

Woolf, V. 2002 (1925). *Mrs. Dalloway*: Harcourt