hyperrhiz: new media cultures

COMPUTATIONAL NARRATION OF INNER THOUGHT: MEMORY, REVERIE MACHINE

JICHEN ZHU

School of Visual Arts and Design
University of Central Florida

D. FOX HARRELL

Comparative Media Studies Program | Program in Writing and Humanistic Studies
Computer Science and Artificial Intelligence Laboratory
Massachussetts Institute of Technology

ABSTRACT

Human mental activities such as memories, reveries, and daydreams are important aspects of our experiences and literary productions. In the emerging area of computer-generated narratives, however, the subjects of preconscious and subconscious inner thought are rarely tackled because the subjective and affective natures of such phenomena are difficult, and perhaps not possible, to pin down in terms of traditional computer science and artificial intelligence (AI). This paper calls attention to underlying parallels and synergies between stream of consciousness literature, cognitive linguistics, and AI. In this historical moment, we propose that concerns of modernist writers regarding inner thought have been reinvigorated in light of these contemporary cognitive scientific developments regarding preconscious conceptualization. Specifically, we present our computational narrative project Memory, Reverie Machine (MRM), which explores how to generate stories about artificial memories, reveries, and daydreams with AI and conceptual blending techniques. Our work exemplifies a new form that can leverage AI technologies for expressive narrative without being burdened by their philosophical baggage or implicit aesthetic dictates.

1. INTRODUCTION

Tyrell: Commerce is our goal here at Tyrell. 'More human than human' is our motto. Rachael is an experiment, nothing more. We began to recognize in them strange obsession. After all they are emotional inexperienced with only a few years in which to store up the experiences which you and I take for granted. If we gift them the past we create a cushion

or pillow for their emotions and consequently we can control them better.

Deckard: Memories. You're talking about memories.

— *Blade Runner*, 1982 (Scott)

Narratives incorporating memories and daydreams are commonly deployed by writers to depict the inner world of characters. These mental phenomena, however, have not been fully explored by the electronic literature and computational narrative communities, both of which have been profoundly touched in recent years by advances in computer gaming. Until recently, many gaming-oriented works have focused on advancing the plot through object acquisition, combat, or puzzle solving, leaving behind a full range of psychological aspects of the characters.

One of the main reasons for this disconnect is that the fluidity and subjectivity of these inner experiences are difficult to pin down in objectivist computer science terms, for example, using classic artificial intelligence (AI) approaches. For this same reason, they are often portrayed in fiction as the final frontier between the rational AI-powered machines and the emotional mankind in popular culture, as in the excerpt above.

In this paper, we explore the generation and narration of characters' inner worlds through memories and daydreams. Our approach is motivated by parallels we have noticed in the concerns of stream of consciousness literature, concept generation in cognitive linguistics, and knowledge representation in artificial intelligence. We argue that the modernist writers' concerns of human subjectivity and pre-speech level consciousness are reinvigorated in the context of computational narrative by recent developments in cognitive linguistic theories and AI techniques. As the result of our interdisciplinary theorization, we introduce our computational narrative project *Memory, Reverie Machine* (MRM), which is geared towards generating textual exposition of artificial memories and daydreams that are triggered by and contrast with the world at hand. Informed by related literary antecedents, MRM extends the framework of Harrell's GRIOT system and generates affective and subjective inner world of its main character.

Acknowledging cognitive linguistics critiques of computational approaches to cognitive modeling (Evans et al., 2006), our work is not an attempt to reduce mental activities to a formal algorithmic process. We are inspired by the fact that algorithmic and knowledge engineering approaches themselves can be critically engaged (Agre, 1997) as well as expressive (Mateas, 2001, Wardrip-Fruin, 2008). We believe that the theory, content, and form we present below have collided in this work in a way that is charged at this historical moment where scientists of the mind provide a new lens on modernist literary concerns, while computation provides a new means for engaging both.

In the rest of the article, we first locate the nexus of our literary, cognitive, and computing concerns in **Section II**.We trace the historical lineage of intersections and divergences between stream of consciousness literature, AI, and cognitive science and highlight a convergence of concerns that we feel gives *Memory, Reverie Machine* (MRM) particular salience now. Next, **Section III** describes the creative framework of MRM. After acknowledging MRM's literary antecedents in print and electronic literature, we will present our theoretical framework that draws together computational approaches to meaning and narrative with the conceptual blending cognitive linguistic theory. Different components of MRM will be demonstrated through an example. Finally, we conclude the article in **Section IV** with reflections upon our results and possible future directions in the near and long terms.

2. HISTORICAL CONTEXT: SCIENTIFIC AND LITERARY EXPLORATION OF INNER-THOUGHT REINVIGORATED

"Stream of consciousness" is a psychological term that William James coined in his 1890 *text The Principles of Psychology*(James, 1890). The term was later applied to works by various modernist writers such as Dorothy Richardson, James Joyce, Virginia Woolf, and William Faulkner, indicating both their literary techniques and the genre itself. Beyond various formal experiments, stream of consciousness literature reflects a conceptual purpose — to use the internal thoughts as a primary way of depicting fictional characters. As Humphrey puts it in his study in *Stream of Consciousness in the Modern Novel*(Humphrey, 1954), the works under this genre replace the motivation and action of the "external man" with the psychic existence and functioning of the "internal man."

Decades have passed since modernist authors' initial experiments and many works associated with this literary experiment have entered the canon of "high" literature. Their approach to expressing aspects of human subjectivity and pre-speech consciousness nevertheless are still relevant to many recent technologies (AI), theories (cognitive linguistics), and forms (computational narrative). These younger developments, in their own ways, have taken the modernist writers' steps further in ways described below. Our goal of narrating memory, reverie, and daydreaming computationally in a new form of polymorphic fiction requires an in-depth

understanding of the synergy among these fields.

2.1 STREAM OF CONSCIOUSNESS LITERATURE AND AI

Stream of consciousness writing and AI may pose an unlikely match as a subject of comparative analysis. The two fields not only sprouted in different historical periods, but also reside in two separate communities. One was populated in the early twentieth century and is now associated with academic literary analysis more often than being seen as vibrant area for active creative production, whereas the other is a still on-going development in the techno-science sphere that underwent significant self-reevaluation after the so-called "AI-Winter" of the 1970s (Russell and Norvig, 2002). Beneath the obvious differences, however, are the similar overarching goals in their respective historical contexts and parallel roles that they both take on in their relationships to contemporaneous concerns.

First, stream of consciousness literature and AI speak to each other through a shared ambition. Humphrey observed that "[t]he attempt to create human consciousness in fiction is a modern attempt to analyze human nature (Humphrey, 1954)." If stream of consciousness writers sought their answers by portraying humans directly, the AI community pursued theirs by constructing the "other" — machines. AI practitioner Michael Mateas recently echoed that "AI is a way of exploring what it means to be human by building systems (Mateas, 2002)." These systems, built in attempt to resemble or surpass their human creators, have become our mirrors to reflect upon our identities as humans (Turkle, 1984).

Secondly, both fields rejected behaviorism in their respective historic periods, and turned their attentions to what happens internally in human mental activities as gateways to understanding "human nature." Prior to the turn of the twentieth century, fictional characters were typically represented by their external behaviors. Writers carefully crafted their actions, dialogues, and rational thoughts to create distinctive personas for their stories. What stream of consciousness writers were able to achieve, in comparison, was to create their characters mainly out of their psychological aspects, including their buzzing random thoughts and associative trails.

The scientific community from which AI grew out of in the 1950s, in parallel, was similarly dominated by behaviorism. The paradigm was based on the laws of stimulus-response and declared itself as the only legitimate scientific inquiry. Mental constructs such as knowledge, beliefs, goals and reasoning steps were dismissed as unscientific "folk psychology" (Russell and Norvig, 2002). Part of AI's contribution was to bring these scientific taboo back to the table by building powerful computational systems based on them. Like Newell and Simon's 1957 General Problem Solver (Newell et al., 1959), many research efforts have been poured into modeling human cognitive capabilities, including reasoning, planning, and learning.

It is worth pointing out some of the differences between the two fields that are relevant to our project. Although both look at cognitive phenomena, stream of consciousness writers and AI practitioners emphasize different stages of human consciousness. The term "consciousness" from the vantage point of modernist writers referred to "the whole area of mental processes, including especially the pre-speech levels." This was based on James' original psychological theory, in which "memories, thoughts, and feelings exist outside the primary consciousness" and, further, that they appeared, not as a chain, but as a stream, a flow (James, 1890). Early AI, on the other hand, regarded human rationality as the key to problem solving. Early practitioners in the field relied upon the rational and stable operations of our cognitive processes at the cost of the addressing the roles of the body, affect, and the uncontrollable stream of thoughts unmediated by logic and rationality.

Another difference between AI and stream of consciousness literature is the conflict between specificity and generalizability. Modernist writers such as Virginia Woolf believed that the important subject for an artist to express was her private vision of reality, that is life and subjectivity, is. Woolf's characters, such as Clarissa Dalloway, Mrs. Ramsay, and Lily Briscoe, all embodied her belief in the individual's constant search of meaning and identification (Humphrey, 1954). This individualistic approach contrasts strongly with AI's focus on generalizability, in which individual differences are often sacrificed for regularity and scalability. We reconcile these two stances by distancing our work from an attempt to reduce mental activities to uniform formal algorithmic processes. Instead, our project utilizes scientific computational methods, including logical/mathematical formalization, as a way to express our human search for meaning.

2.2 STREAM OF CONSCIOUSNESS LITERATURE AND COGNITIVE LINGUISTICS

The pre-speech level of thought that was neglected by the AI community has been scrutinized again recently in a new field closely built, in part, upon AI: cognitive science. To contemporary cognitive linguists, such as Gilles Fauconnier, George Lakoff, Mark Johnson, and Mark Turner, this neglected land of consciousness holds the basis for our basic conceptual, and even literary thought (Fauconnier and Turner, 2002, Fauconnier, 1985, Lakoff and Johnson, 1980).

Language is only the tip of a spectacular cognitive iceberg, and when we engage in any language activity, be it mundane or artistically creative, we draw unconsciously on vast cognitive resources, call up innumerable models and frames, set up multiple connections, coordinate large arrays of information, and engage in creative mappings, transfers, and elaborations" (Fauconnier and Turner, 2002).

Beneath the tip of this iceberg is what Fauconnier calls "backstage cognition" (Fauconnier, 2000), defined as "the intricate mental work of interpretation and inference that takes place outside of consciousness" (Fauconnier and Turner, 2002). Thus, we could say that cognitive linguists cite phenomena that are even below the unarticulated thought phenomena explored by stream of consciousness authors — but at a level that still addresses conceptualization as opposed to perception, motor-action, or other pre-conscious cognitive phenomena. Fauconnier sites a range of results in cognitive science to support his conjecture that many cognitive phenomena are rooted in backstage cognition:

viewpoints and reference points, figure-ground / profile-base / landmark-trajectory organization, metaphorical, analogical, and other mappings, idealized models, framing, construal, mental spaces, counterpart connections, roles, prototypes, metonymy, polysemy, conceptual blending, fictive motion, force dynamics (Fauconnier, 2000).

It may be argued that one of the reasons that early AI largely confined itself to the territory of rationality is the extreme difficulty that the field ran into in its attempt to model common sense and contextual reasoning explicitly. These powerful, but for the most part invisible, operations are seen within the field of cognitive linguistics to be partially observable in the structure of our linguistic creations. The challenges posed by the cognitive linguistics enterprise offer the opportunity to revisit some of the compromises that AI made in its early stage.

2. 3 CHALLENGES OF ENGAGING LEGACY FORMS

In forging this bond between stream of consciousness literature, AI, and cognitive science, we are sensitive to a common trend in computing, that is, trying to prove its triumphs by taking on tasks that seem to be humanly creative in ways that are highly esteemed culturally. Examples include composing music like Mozart (Cope, 2001), playing chess better than a grandmaster (Hsu, 2002), or completing complex mathematical proofs (Lenat and Brown, 1984). Likewise, we are aware the literature is often viewed in academic disciplines as a succession of movements, tastes, trends, and techniques. Stream of consciousness literature holds a special status in the canon of English literature and therefore imposes unique challenges when used as an inspiration for this computational narrative project. When engaging what might be called legacy forms, authors are inevitably put in dialogue with the original works' historic contexts, social statuses, political agendas, etc. The following clarifies our project's position in relation to this set of modernist literature works.

Our aim in referencing stream of consciousness writing is not to inherit its highbrow status or to prove the effectiveness of the authors' system by tackling the "greats" of literature. Nor is it a sign of indifference to more recent modes of literary cultural production. The reason for engaging this particular literary form is that it offers critical insights into the AI practice and expressive platform to deploy cognitive science research results (as well as the admittedly desirable benefit of extended dialogue with the prose of Virginia Woolf — her work inspired early projects that were precursors to *MRM*). Engaging a legacy form in a different medium requires more than mere translation. The intent of our project is not to generate texts similar to texts by stream of consciousness writers that can "fool" the reader or pass a type of "stream of consciousness writing Turing test." Instead, we intend to establish a new aesthetic form to depict human conditions and subjective experiences in unique ways afforded by computational techniques.

3. MEMORY, REVERIE MACHINE: AN INTERACTIVE POLYMORPHIC NARRATIVE

Below, we describe the influences of forbearers in conjoining concerns of computing and literature, specifically relevant work in cognitive science, and foundational work by one of the authors for our current project. We then highlight the model proposed by *Memory, Reverie Machine* using an example comprised of actual system output.

3.1 LITERARY ANTECEDENTS AS CATALYSTS FOR INTEGRATING AI, COGNITIVE SCIENCE, AND STREAM OF CONSCIOUSNESS CONCERNS

Before we turn to fully to our new developments, we call attention to our antecedents. The goal of generative literature to stir human imagination through works that are different on each reading is not new. We are inspired by precursors in the area, not the least because for many of them each work represented a new literary form. An early relevant work is Raymond Queneau's 1961 "Cent Mille

Milliards de Poémes" ("One Hundred Thousand Billion Poems") (Queneau, 1961), originally published as a set of ten sonnets with interchangeable lines, but later made available in computer implementations. This work is relevant because of its exploration of the idea of writing as a combinatorial exploration of possibilities, which exemplifies the experimental literary group Oulipo's often whimsical use of mathematical ideas. This view is well explicated by another Oulipo member, Italo Calvino, in his essay/lecture "Cybernetics and Ghosts" (Calvino, 1982). Calvino claimed that writing was a combinatoric game and cited work such as Vladimir Propp's morphology of the folktale to support his thesis (Propp, 1968). Calvino stated "the operations of narrative, like those of mathematics, cannot differ all that much from one people to another, but what can be constructed on the basis of these elementary processes can present unlimited combinations, permutations, and transformations." In Calvino's novels such as "If on a winter's night a traveler" there was also a strong sense of narrative coherence and a concern for a careful balance between experimental form and meaningful expression, yet Calvino also carefully explicated the algorithmic generation of the novel's form (Calvino, 1995).

As computer scientists and computational cultural producers in literature began producing electronic literary works, a foundation was laid for the natural integration of AI, cognitive science, and literary concerns. To trace this foundation it is instructive to revisit several important texts in electronic literature. Influential systems from the field of computer science include Meehan's Talespin (Meehan, 1981), Scott Turner's MINSTREL (Bringsjord and Ferrucci, 2000), and Bringsjord and Ferrucci's BRUTUS (Bringsjord and Ferrucci, 2000). It is useful to contrast the differences between Talespin as an example of an early system and BRUTUS as an example of a more recent system. Meehan's 1976 Talespin is perhaps the first computer story generation system. It produced simple animal fables, with the goal of exploring the creative potential of viewing narrative generation as a planning problem, in which agents select appropriate actions, solve problems within a simple simulated world, and output logs of their actions (which are not generally effective as interesting stories — though Noah Wardrip-Fruin has convincingly argued that they are interesting as processes (Wardrip-Fruin, 2009). Selmer Bringsjord and David Ferrucci's 2000 BRUTUS system aims to explore formalizations for generating stories about betrayal, with the goal of being "interesting" to human readers. While Talespin directly exposed a reader to the output of a planning algorithm, BRUTUS uses rich textual descriptions. In fact, the output of BRUTUS is nearly pre-determined by the system's pre-authored output. For example a university is an object implemented with both positive iconic features:

```
{clocktowers, brick,ivy, youth, architecture, books, knowledge, scholar, sports}
```

and negative iconic features:

```
{tests, competition, 'intellectual snobbery'}
```

As a goal, Bringsjord and Ferrucci seek a type of storytelling Turing test competence. Viewing a system as an intelligent author calls into question the role of low-grained pre-authored content in such work, Bringsjord and Ferruci narrowly avoid claims that the system actually authored the texts.

In contrast to those computer science based systems, William Chamberlain and Thomas Etter's dialogue based program Ractor, and Ractor's book, *The Policeman's Beard is Half Constructed* (1984), used syntactic text manipulation to support conversation with users having text input and poetic output. This was not intended as scientific research, but rather as entertainment, with humorous and clever output. Chamberlain described such output as being computer authored, exploiting the novelty of being "written by a computer program." Charles Hartman's 1996 work in automated poetry generation (Hartman, 1996) was presented as literary experimentation, but Hartman realized that it is better not to ask "whether a poet or a computer writes the poem, but what kinds of collaboration might be interesting." Hartman's work emphasizes how a computer can introduce "randomness, arbitrariness, and contingency" into poetry composition.

3.2 BLENDING-BASED CONCEPT GENERATION

In contrast to the notion of computational generativity, the human capacity to generate concepts and metaphors has been explored by cognitive scientists as the root of our literary mental processes. Conceptual blending theory, building upon Gilles Fauconnier's mental spaces theory (Fauconnier, 1985) and elaborating insights from metaphor theory (Lakoff and Johnson, 1980, Turner, 1996), describes the means by which concepts are integrated. In short, the theory describes how we arrive at new concepts through blending partial and temporary pieces of information. Most importantly, the theory proposes that conceptual blending processes occur uniformly in pre-conscious in everyday thought and in more complex abstract thought such as in literary arts or rhetoric. (Fauconnier and Turner, 2002) Though the empirical grounding for such findings is still being developed (Gibbs Jr., 2000), this turn toward backstage cognition reflects the trend in cognitive science away from formal, logical, and rational thought and toward context-driven, cultural, and embodied thought (Fauconnier, 2000).

Blending is considered a basic human cognitive operation, invisible and effortless, but pervasive and fundamental, for example in grammar and reasoning. (Fauconnier and Turner, 2002) Simple examples of blending in natural language are reflected in the mental processes triggered by words like "houseboat" and "roadkill," and phrases like "artificial life" and "computer virus." However, blending has been proposed as underlying the ability to daydream. Eliding his example a bit, Mark Turner writes (Turner, 2003):

> Consider the as yet unexplained human ability to conjure up mental stories that run counter to the story we actually inhabit. ...suppose you are actually boarding the plane to fly from San Francisco to Washington, D. C. You must be paying attention to the way that travel story goes, or you would not find your seat, stow your bag, and turn off your personal electronic devices. But all the while, you are thinking of surfing Windansea beach, and in that story, there is no San Francisco, no plane, no seat, no bag, no personal electronic devices, no sitting down, and nobody anywhere near you. Just you, the board, and the waves.

This ability to construct an alternative, imaginative world in the mind in contrast with the real world at hand is presented as a case of a "double-scope story," a blend in which the two spaces to be integrated consist of quite differently structured concepts that Turner calls event stories.

Conceptual blending theory provides constitutive principles (comprising an idealized model and process) and governing principles (constraints determining which conceptual blends are more "optimal" than others for everyday thought) for how blends are constructed (Fauconnier and Turner, 2002). Some of these seem to be purely structural and thus computationally amenable. The meaning-driven generativity that contributes to the variability of output in our own work is based on a model of conceptual blending. In our computational narrative work described below, we use an algorithm for conceptual blending to aid in generating the affective disposition differently for each output. We believe that such a technique should also be useful to generate memories and reveries.

3.3 THE GRIOT SYSTEM

The GRIOT system is the foundation of the *Memory, Reverie Machine* project both in terms of technical implementation and our approach to computational narrative (Harrell, 2007a, Harrell, 2006). This subsection, adapted from the abstract of (Harrell, 2007b),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. It uses Joseph Goguen's theory of algebraic semiotic approach from computer science to formalize key aspects of conceptual blending theory (Goguen, 1998). Technical details can be found in (Goguen and Harrell, 2008). 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. In the section below we describe our new work that arises from the historical context presented above and the theoretical and creative framework just described.

The goal with GRIOT is quite different from passing a type of Turing test for creative competence, as described earlier this section. It is designed to provide a technical framework for humans to provide rich content and narrative systems created with GRIOT are meant as cultural products themselves (as opposed to instances of output of such poetic systems). The GRIOT system utilizes models for cognitive science, informed by the skepticism of logical AI modeling of human thought in the cognitive linguistics enterprise (Lakoff, 1999), toward expressive ends that are often literary. *Memory, Reverie Machine* quickens this interest in terms of subject matter and

interaction model as well. We explore the meaning and limitations of machinic thought through fiction, while critiquing outmoded notions of machinic thought theoretically and technically. Finally, we engage the notion of generating memories and reveries cascading sequences of remembered events, as reinvigorated by work such as conceptual blending theory from cognitive science and novel techniques for narrative interaction and generation that are informed by AI-based technologies. We demonstrate these points through an example.

3.4 FRAMEWORK FOR *MEMORY, REVERIE MACHINE*

So far, we have situated our work in a historical web where stream of consciousness literature, artificial intelligence discourse, and cognitive science research complement each other. This correlation converges with our aesthetic goal and calls for a new form of computational narrative, on the subject of daydreams, memories, and reveries.

In this subsection, we present *Memory, Reverie Machine*, a text-based computational narrative project through one of our early results (illustrated in **Figure 1** below). For this work it is important to distinguish between the project's various levels of technological and expressive investigation and production:

1. the system as an abstract model for how computational narratives can be made generative, extensible, and reconfigurable,
2. the system that generates the story,
3. the narration techniques developed,
4. the story content, and
5. each instance of output.

The emphases in this article are upon (1) and (2), which comprise our technological framework as inspired by the historical context and theoretical and creative frameworks above, and secondarily on (3), the narrative techniques to depict inner thoughts, and (4) the self-reflexive subject matter. (3) is heavily influenced by Virginia Woolf's stream of consciousness novel *Mrs. Dalloway* (Woolf, 2002 (1925). Below we highlight particular aspects of the system relevant to the technical and expressive goals involving the narration of inner-thought, complete technical information on GRIOT can be found in and of *Memory, Reverie Machine* (formerly the Daydreaming Machine) in (Zhu and Harrell, 2008).
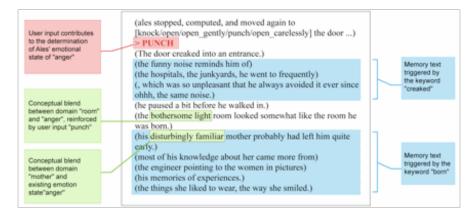

Figure 1: An example of output from *Memory, Reverie Machine*
(click for large version)

3.4.1 *Dynamic Narration of Affect Using the Alloy Conceptual Blending Algorithm*

Computationally, our system draws upon the GRIOT framework of (Harrell, 2007b), which identifies, formalizes, and implements an algorithm for structural aspects of conceptual blending with applications to computational narrative. Alloy is the primary generative component of the GRIOT system. The Alloy algorithm, when modestly applied in *Memory, Reverie Machine*, generates blends involving connecting current experiences of events, objects, and actors (Turner, 1996) to affective concepts determined by current state, in this case emotional state, of the protagonist (see next subsection). In the example in **Figure 1**, logical axioms selected from an ontology (semantically structured database) describing the concept "door" are blended with axioms describing affective concept "anger".

The Alloy algorithm uses a set of formal optimality criteria to determine the most common-sense manner in which the concepts should be integrated (Harrell, 2007b). The result is a blended axiom or set of axioms that is then mapped to natural language output. For

example, the description of the door ranges from "distasteful wood-colored" to "irritatingly sturdy" or more depending on the concepts being blended. Since blending refers to the conceptual integration of multiple concepts, it is important to be clear that blending is not the mere concatenation of words to form compound phrases. In this case, compound phrases, some of the simplest indicators of conceptual blends, are the final result of an underlying process that is semantic, not lexical. It is not difficult to imagine that the phrase "inviting gateway" from the sample output in introduction comes from blending concepts of "door" and "happiness" and trying to articulate the outcome in the English language.

Constructing blends between objective and affective concepts allows us to achieve a balance between author determined plot and variable theme or emotional tenor. An artifact required by the plot can be depicted in various ways based on the character Ales's internal emotional state. The highly subjective description, in turn, portrays personality traits of the character, a recurrent technique in *Mrs. Dalloway*.

3.4.2 *The Emotion State Machine*

Actions taken by a character in a computational narrative, which are usually (but not exclusively) selected by a user, can guide building up of a profile that describes user's preferences, history of actions, and analysis of trends in those actions. A simple, but effective, step in this direction is tracking the emotional state of a character based upon actions that the character has taken. *Memory, Reverie Machine* allows user to directly influence the emotion state of Ales, and hence the selection of affective concepts for blends. She may choose among an array of pre-defined actions, such as seeing objects as "red", "yellow", "blue" or another color in the robot character's optic sensors, each connecting to a particular emotion. A keyword "red", for instance, may trigger an affective concept "anger". These emotional mappings are designed aesthetically by the authors to achieve narrative effects, not as an attempt at cognitively modeling emotion using a computer as has been the stated goal in multiple traditional AI projects.

A successful interactive narrative, however, requires a careful balance between the user's agency and author's intention. In our system, user's impact on the character's emotion is moderated by the emotion state machine component for the sake of narrative consistency. The state machine records Ales's current emotion based on the entire history of user input, instead of the most immediate one. It guarantees that changes of Ales's emotions will be gradual, even if user input oscillates between opposite emotions.

3.4.3 *Memory Structuring and Retrieval*

The GRIOT System is not limited to producing narrative discourse, indeed it has been used for various forms of poetry (Goguen and Harrell, 2005, Harrell, 2007a), and for the interactive, generative composition of visual imagery (Harrell and Chow, 2008). In the case of *Memory, Reverie Machine*, we seek to make output coherently extensible at runtime. For this project we allow the narrative to be punctuated with episodic remembered events and longer reveries of remembered experience. Again, this is not meant as an experiment in cognitive modeling or advanced algorithmic design, our goal is to demonstrate discourse that is meaningfully reconfigurable to serve an author's expressive goals in dialogue with a user's selected actions. In this case, each memory is annotated based on its subject matter and is retrieved when at least one subject item appears in the story line. In the example in **Figure 1**, Ales's unpleasant memory of hospitals and junkyards is triggered by the opening of a door through the mutual subject of a certain sound. The system also keeps track of the emotional tone of each memory and selects a memory only if it does not clash with the current emotion state. The example at the start of the **Introduction** also illustrates this feature.

We have described the major components of *Memory, Reverie Machine* that align with our tale of inner-thought and AI in historical concern through illustrative examples. To result in a longer form interactive and generative story, of course, requires expanding each component. In the next section we reflect upon our accomplishments and challenges and indicate both near and long term future directions for the project.

4. CONCLUSIONS AND FUTURE WORK

"Mrs. Dalloway said she would buy the flowers herself." And so Virginia Woolf started the journey in her 1925 novel *Mrs. Dalloway* (Woolf). She then took her readers in and out of the thoughts, the memories, and the emotions of her characters in such a fluid and seamless way that the boundary between the internal thoughts of a character and the external world, and between fact and fiction, collapsed. The reader no longer experienced the story as an objective observer; rather, she lived the story world as it was perceived by the senses of Woolf's characters.

The eventual goal of *Memory, Reverie Machine* is to revisit the expressive and humanistic concerns of inner subjectivity from the

vantage point of above-mentioned mental activities, as demonstrated in Woolf's work and many others'. Our approach is equipped with insights from empirical cognitive science discoveries and technical development of artificial intelligence. Though on one hand our work relates to interactive fiction in which the author's task is to author pre-written story elements and puzzles for a user to explore though selecting actions through entering noun verb combinations in text, we are also inspired by the generativity of our literary antecedents and our next step is to bolster the system's generativity and to reduce the use of pre-written templates. The blending algorithm can be modified to favor less optimal, but more interesting results, and the input to the blending algorithm can consist of more complex conceptual spaces, perhaps even to model the creation of double-scope stories. The emotional state machine may weight different input differently in its computation of Ales's emotion. Reveries can be implemented as a cascade of memories, triggering one another. Above all, the overarching plot needs to move beyond its current pre-written stage and needs to be generated while taking all these components into account.

Work such as ours speaks to several different audiences, each with their own demands, values, and registers. Elsewhere, we have described works in which characters or systems are seen as being about something in the word and having their own agency directed toward the world or story world as "intentional systems" (Zhu and Harrell). One criterion for systems to be accepted as AI systems, rather than only intentional systems is computational complexity, especially using complex algorithmic and knowledge engineering techniques accepted and explored within the computer science discipline of artificial intelligence. Our work contains many problems that would be construed as AI problems, such as generating discourse, generating conceptual blends, engineering a character's knowledge of the world, modeling the embodied narrative experience of exploring a world and more. However, we have focused on a developing an interdisciplinary theoretical and technical framework in which more fundamental questions about the nature of AI practice itself can be explored along with new techniques for support of expressive forms. So while we are not driven by the need for computational complexity, we see more generative algorithmic work as necessary to push forward our theoretical and expressive goals.

For literary artists and digital media arts practitioners, rich prose, genre experimentation, and interaction design will perhaps be seen as the ripest areas for future exploration. Allowing the user to input keyword, as in our current examples, can easily be bolstered so that a graphical user interface is what pushes information to the underlying discourse structuring and generation system. A user could perform meaningful actions such as opening doors, washing dishes, setting visual scanners, and more through a visual interface mechanism. Alternately, in a text-based system user input could certainly be more complex than simple keywords as in interactive fiction works or even more robust natural language based works such as the important expressive AI work *Façade* (Mateas, 2001).

Our goal, however, is not to acquiesce to the values system of any one component discipline. Rather, we intend to be driven by content and conceptual issues in our work. Tracing the historical lineage of ideas such as pre-speech thought and automated intelligent machines in both the computer science and literary fiction domains calls our attention to what is new, what is salient in light of new theory, and what are the anxieties and questions being considered in our current social world where computing and cultural expression move into greater and greater alignment. Hence, beyond technical or interaction design progress, we also seek to define a space in which the intersection of both of those concerns, and concerns for human experience and living, intersect. We hope that the path traced here, and its combination of historical and novel technical grounding, provides a step in that direction. Forwarding and communicating this interdisciplinary developing perspective is our most urgent future task.

APPENDIX: OUTPUT FROM MEMORY, REVERIE MACHINE

Ales the robot can choose how he sees his world. Blue, and he is in melancholy-world, with doors as "depressing entryways," yellow, and he is in happy-land, where each door is an "inviting gateway." Particular objects and events, such as entryways, gateways, or oil-changes trigger his memories. And those memories may trigger others, resulting in cascading sequences of memories to construct reveries. Below is a sample output from MRM, user input is formatted in bold.

```
(the color of the kitchen door today should be [yellow/blue])
> blue
(;) (ales said to himself, adjusting his visual scan system.)
(the room where he had his first encounter of tune-up and oil change had similar
depressing entryways.)
(the oil change left a sickly feeling in his gut;)
(he would rust like the tin man before enduring another.)
```

```
    (it amazes him that no one else complained about these things;)
    (they were too busy with the kitchen;)
    (he should probably start his daily routine now.)
```

BIBLIOGRAPHY

Agre, P. E. (1997) *Computation* Cambridge, U.K., Cambridge University Press.

Bringsjord, S. & Ferrucci, D. A. (2000) *Artificial Intelligence and Literary Creativity: Inside the Mind of BRUTUS, a Storytelling Machine,* Mahwah, NJ, Lawrence Erlbaum Associates.

Calvino, I. (1982) *The Uses of Literature,* San Diego, C.A., Harcourt Brace and Company.

Calvino, I. (1995) How I Wrote One of My Books. *Oulipo Laboratory: Texts from the Bibliotheque Oulipienne.* London, Atlas.

Cope, D. (2001) *Virtual Music,* Cambridge, MA, MIT Press.

Evans, V., Bergen, B. K. & Zinken, J. (2006) The Cognitive Linguistics Enterprise: An Overview. In Evans, V., Bergen, B. K. & 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. (2000) Methods and Generalizations. In Janssen, T. & Redeker, G. (Eds.) *Scope and Foundations of Cognitive Linguistics.* The Hague, Mouton De Gruyter.

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

Gibbs JR., R. W. (2000) Making good psychology out of blending theory. *Cognitive Linguistics,* 11**,** 347-358.

Goguen, J. (1998) An Introduction to Algebraic Semiotics, with Applications to User Interface Design. In Nehaniv, C. (Ed.) *Computation for Metaphors, Analogy, and Agents.* Yakamatsu, Japan.

Goguen, J. & Harrell, D. F. (2005) The Griot Sings Haibun. La Jolla.

Goguen, J. & Harrell, D. F. (2008) Style as a Choice of Blending Principles. In Argamon, S., Burns, K. & Dubnov, S. (Eds.) *The Structure of Style: Algorithmic Approaches to Understanding Manner and Meaning.* Berlin, Springer.

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

Harrell, D. F. (2007a) GRIOT's Tales of Haints and Seraphs: A Computational Narrative Generation System,. In Wardrip-Fruin, N. & Harrigan, P. (Eds.) *Second Person: Role-Playing and Story in Games and Playable Media.* Cambridge, MA, MIT Press.

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

Harrell, D. F. & Chow, K. K. N. (2008) Generative Visual Renku: Linked Poetry Generation with the GRIOT System. *Electronic Literature Organization Conference.* Vancouver, WA.

Hartman, C. O. (1996) *Virtual Muse: Experiments in Computer Poetry,* Hanover, CT, Wesleyan.

Hsu, F.-H. (2002) *Behind Deep Blue: Building the Computer that Defeated the World Chess Champion,* Princeton, NJ, Princeton University Press.

Humphrey, R. (1954) *Stream of Consciousness in the Modern Novel,* Berkeley and os Angeles, University of California Press.

James, W. (1890) *The Principles of Psychology,* New York, Henry Holt And Company.

Lakoff, G. (1999) *Philosophy in the Flesh: The Embodied mind and its challenge to western philosophy,* New York, Basic Book.

Lakoff, G. & Johnson, M. (1980) *Metaphors We Live By,* Chicago, University of Chicago Press.

Lenat, D. B. & Brown, J. S. (1984) Why AM and EURISKO appear to work. *Artificial Intelligence,* 23**,** 269-294.

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

Mateas, M. (2002) Interactive Drama, Art, and Artificial Intelligence. *Computer Science.* Pittsburgh PA, Carnegie Mellon University.

Meehan, J. (1981) TALE-SPIN. In Riesbeck, C. & Shank, R. (Eds.) *Inside Computer Understanding: Five Programs Plus Miniatures.* New Haven, CT, Laurence Erlbaum Associates.

Newell, A., Shaw, J. C. & Simon, H. A. (1959) Report on a general problem-solving program. *Proceedings of the International Conference on Information Processing.*

Propp, V. (1968) *Morphology of the Folktake,* Austin, TX, University of Texas Press.

Queneau, R. (1961) *Cent émes,* Paris, France, Gallimard.

RACTER (1984) *The Policeman's Beard Is Half Constructed,* New York, Warner Books.

Russell, S. & Norvig, P. (2002) *Artificial Intelligence: A Modern Approach*, Prentice Hall.

Turkle, S. (1984) *The Second Self: Computers and the Human Spirit,* New York, Simon and Schuster.

Turner, M. (1996) *The Literary Mind: The Origins of Thought and Language,* New York; Oxford, Oxford UP.

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

Wardrip-Fruin, N. (2008) Expressive Processing: An Experiment in Blog-Based Peer Review. *Grand Text Auto.*

Wardrip-Fruin, N. (2009) *Expressive Processing: Digital Fictions, Computer Games, and Software Studies,* Cambridge, MIT Press.

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

Zhu, J. & Harrell, D. F. (2008) Daydreaming with Intention: Scalable Blending-Based Imagining and Agency in Generative Interactive Narrative. *AAAI 2008 Spring Symposium on Creative Intelligent Systems.* Stanford, CA, AAAI Press.

D. Fox Harrell, Ph.D., is a researcher exploring the relationship between imaginative cognition and computation. He is an Associate Professor of Digital Media at MIT in the Comparative Media Studies Program, Program in Writing and Humanistic Studies, and in the Computer Science and Artificial Intelligence Lab (CSAIL). His research involves developing new forms of computational narrative, gaming, social networking, and related technical-cultural media based in computer science, cognitive science, and digital media arts. He received an NSF CAREER Award from the Human-Centered Computing cluster for his project "Computing for Advanced Identity Representation." Harrell holds a Ph.D. in Computer Science and Cognitive Science from the University of California, San Diego. He earned a Master's degree in Interactive Telecommunication from New York University's Tisch School of the Arts. He also earned a B.F.A. in Art, a B.S. in Logic and Computation, and minor in Computer Science at Carnegie Mellon University. He is completing a book for the MIT Press called *Phantasmal Media: An Approach to Imagination, Computation, and Expression*.

Jichen Zhu is currently an assistant professor of Digital Media in School of Visual Arts and Design at University of Central Florida. She is also the director of the Procedural Expression Lab. Her work focuses on developing humanistic and interpretive theoretical framework of computational technology, particularly artificial intelligence (AI), and constructing AI-based cultural artifacts. Her current research areas include digital humanities, software studies, computational narrative, and serious games. Zhu received a Ph.D. in Digital Media from Georgia Tech. She also holds a M.S. in Computer Science from Georgia Tech, a Master of Entertainment Technology from Carnegie Mellon University, and a B.S. in Architecture from McGill University, Canada.