

Slant: A Blackboard System to Generate Plot, Figuration, and Narrative Discourse Aspects of Stories

Nick Montfort

The Trope Tank, MIT
77 Mass Ave, 14N-233
Cambridge, MA 02139 USA
nickm@nickm.com

Rafael Pérez y Pérez

División de Ciencias de la
Comunicación y Diseño
Universidad Autónoma
Metropolitana, Cuajimalpa,
México D. F.
rperez@correo.cua.uam.mx

D. Fox Harrell

Imagination, Computation, &
Expression Laboratory, MIT
77 Mass Ave, 14N-207
Cambridge, MA 02139 USA
fox.harrell@mit.edu

Andrew Campana

Department of East Asian
Languages & Civilizations
Harvard University
Cambridge, MA 02138
USA
campana@fas.harvard.edu

Abstract

We introduce Slant, a system that integrates more than a decade of research into computational creativity, and specifically story generation, by connecting subsystems that deal with plot, figuration, and the narrative discourse using a blackboard. The process of integrating these systems highlights differences in the representation of story and has led to a better understanding of how story can be usefully abstracted. The plot generator MEXICA and a component of Curveship are used with little modification in Slant, while the figuration subsystem Fig-S and the template generator GRIOT-Gen, inspired by GRIOT, are also components. The development of the new subsystem Verso, which deals with genre, shows how different genres can be computationally modeled and applied to in-development stories to generate results that are surprising in terms of their connections and valuable in terms of their relationship to cultural questions. Example stories are discussed, as is the potential of the system to allow for broader collaboration, the empirical testing of how subsystems interrelate, and possible contributions in literary and artistic contexts.

Introduction

Slant is a system for creative story generation that integrates different types of expertise and creativity; the framework it provides also means that other systems, implementing other approaches to story generation, can be integrated into it in the future. The development of Slant has involved formalizing, reworking, and testing ideas about creative storytelling and what is important to writing stories—specifically, the poetics of figuration, the poetics of plot development, and the poetics of narrating. The system incorporates a new perspective on genre and integrates components from three existing systems: D. Fox Harrell’s GRIOT, Rafael Pérez y Pérez’s MEXICA, and Nick Montfort’s Curveship.

Story generation systems have not yet used an architecture of this sort to encapsulate different expertise and different aspects of creativity; nor have they incorporated major components that are based on existing, proven systems by different researchers.

Slant is a blackboard system in which different subsystems, each of them informed by and modeling humanistic theories, collaborate together, working incrementally to fully specify a story. An alternative, simpler process involves making decisions in a “pipeline,” in which one system offers, for instance, a plot and another system determines how the narrative discourse will be arranged. Although this system seems to be a poor model of human creativity, it is a reasonable first step toward a “blackboard” system. Two of the Slant collaborators previously developed such a pipelined system with two stages (Montfort and Pérez y Pérez 2008). The current project involves five major subsystems rather than two and uses a blackboard architecture, allowing any of the subsystems that work during the main phase of generation to augment the story representation at any point.

The generation of stories in Slant begins with minimal, partial proposals from a simple unit, the Seeder. In turn, the subsystems MEXICA, Verso, and Fig-S read and add to this set of proposals, each according to its focus. When the proposals are complete, the finished story specification is sent to GRIOT-Gen so conceptual blending can be applied to the relevant templates and then to the three-stage pipelined text generation component of Curveship. Curveship-Gen realizes a finished story in the form of a text file that can be read and considered by human readers.

This paper introduces the architecture of the system and describes the subsystems that build and realize stories together. It includes a discussion of what was learned by inte-

grating three different lines of research on story generation. Reflections are also offered on the first set of stories produced by the system, and some discussion of the potential of the system is included as well. Slant will undergo more refinement and development, but the work that has been done so far is of relevance to those working to implement large-scale computational creativity systems that integrate heterogeneous subsystems, to those developing representations of story and other creative representations, and to those working specifically in story generation.

Creativity and the Architecture of Slant

Boden holds that creativity involves the production of new, surprising, and valuable results (Boden 2004). In the case of story generation and other literary endeavors, being new involves not repeating what has been done before (by the system or in the wider culture); surprise often manifests itself in unusual juxtapositions that are effective, though one would not have guessed it; and value, rather than indicating that the story is of didactic or economic value, means that a story accomplishes some imaginative or poetic purpose—it connects in some way to cultural or psychological issues or questions and allows the reader to think about them in new ways. Stories that surprise readers by bringing unusual elements together and which provide for this sort of reflection, but which do so in the same way as existing stories, are not new. Stories that are innovative and could allow for reflection, but which do not involve unusual juxtapositions or connections, are not surprising. Stories that are fresh and involve unusual combinations of elements, but do not ultimately seem to have a point of any sort, are not of value.

Taking value to indicate relevance within culture means that the value of a story is similar to what has been called, with regard to conversational stories of the sort that are uttered all the time by people, its “point” (Polanyi 1989). While the point of a story is understood in the context of a specific conversation, the ability of a story to have a point at all can be understood within the context of culture. Valuable stories are those that have a point to at least some readers when they encounter them in some context.

Beyond Boden’s three components of creativity, we also consider a higher level of creativity. Namely, the various cognitive processes for conceptualization that enable people to recognize and generate new, surprising, and valuable cultural content are forms of everyday creativity. Cognitive scientist Gilles Fauconnier has referred to these process of meaning construction as “backstage cognition” and asserts that backstage cognition includes specific phenomena such as “viewpoints and reference points, figure-ground/profile-bases/landmark-trajectory organization, metaphorical, analogical, and other mappings, idealized models, framing, construal, mental spaces, counterpart connections, roles, prototypes, metonymy, polysemy, conceptual blending, fictive motion, [and] force dynamics” (Fauconnier 1999).

These cognitive processes are especially important to note here because the notion of creativity informing Fig-S and GRIOT-Gen is based on a model of the creative backstage cognition phenomenon of metaphorical mapping, most prominently, but also mental spaces, counterpart connections, metaphor, analogy, and metonymy in the case of the GRIOT system that inspired them.

To succeed repeatedly and reliably at creativity, a story-telling system must have mechanisms relevant to each of these aspects of creativity. It must have some model of what has happened before to prompt novelty, somehow provide for stories that join aspects together in unusual and effective ways, and somehow provide for stories that relate to culture and have a point. The means of accomplishing these aspects of creativity do not have to be abstracted into separate components of a system, but they do need to be somehow realized by a creative system.

A simple way that systems can connect and to some extent collaborate involves organizing them in a pipeline. This can model a regimented assembly-line process or “waterfall” model in which each subsystem participates in one phase and interfaces only with the systems before and after it. For certain processes, this may be adequate, but for the nuanced process of creativity, which involves making interesting connections, the components of a system probably need to interact in a less constrained and unidirectional manner. This was the rationale for the blackboard architecture used in Slant.

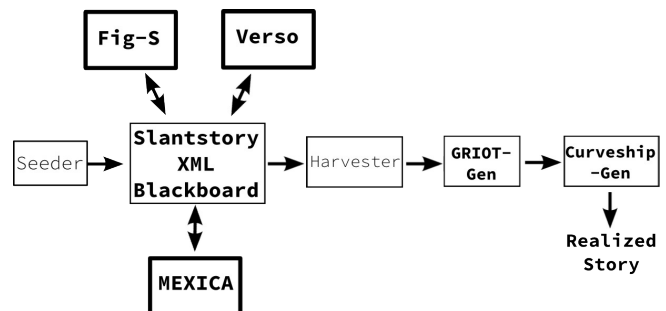


Figure 1: The architecture of Slant.

The Blackboard and Subsystems

In Slant, the three major story-building subsystems can write to and read from a blackboard representation of the story in progress. Currently, the systems function in practice much as a pipeline does, with each of the three subsystems augmenting the story representation once. The systems can influence each other “backwards” only via Verso examining the current plot and proposing a new action (not just a specification of narrative discourse, which is always proposed.) MEXICA can then incorporate that expanded plot into its next ER cycle that it uses to elaborate the plot. Although the interactions between subsystems are not intricate at this

point, the framework is in place for more elaborate black-board interaction in future versions of Slant.

Currently, MEXICA contributes an initial, partial plot – a minimal, random one will eventually be provided at the first step by the Seeder. Then, Verso assigns a genre and a specification of the narrative discourse, and MEXICA further elaborates the plot until it is complete. Verso may specify constraints on how the story is to be developed. For instance, it may specify that a particular character, who has been designated as the narrator of the story, should not die. MEXICA will respect these in elaborating the story. Finally, Fig-S determines what figuration will be used. Eventually, another system, the Harvester, will check to see if all aspects of the story are complete, allowing the subsystems to augment the story in several different orders. After the story representation is complete, it is realized. GRIOT-Gen determines how to realize figurative representations and Curveship-Gen does content selection, microplanning, and surface realization to produce the final text.

The MEXICA subsystem has the most explicit model of an aspect of creativity; it explicitly evaluates the novelty and interestingness of the component of story that it develops, the plot. Verso and Fig-S both aim to add surprise by combining conventional genres and metaphors in unusual ways. They do not currently measure how surprising their results are, but they embody techniques for choosing appropriate combinations that may be seen as creative by readers.

Foundational Systems

MEXICA. This system generates plots or frameworks for short stories about the Mexicas, the old inhabitants of what today is México city, also known as the Aztecs. MEXICA's process is based on the engagement/reflection cycle, a cognitive account of writing by Mike Sharples (Pérez y Pérez and Sharples 1999, 2001, 2004). During engagement the system focuses on generating sequences of actions driven by content and rhetorical constraints and avoids the use of explicit goals or predefined story-structures. During reflection MEXICA evaluates the novelty and interestingness of the material produced so far and verifies the coherence of the story (see also Pérez y Pérez et al. 2011).

The design of the system is based on structures known as Linguistic Representations of Actions (LIRAs), which are sets of actions that any character can perform in the story and whose consequences produce some change in the story-world context. There are two types of possible pre-conditions and postconditions in MEXICA: emotional links between characters and dramatic tensions in the story.

MEXICA is incorporated as the generator of plot. It generates plot in stages, allowing other systems to interact with the story representation as it does so. In the current system, it can be influenced by actions added to the story by Verso.

GRIOT. This is a system that is the basis for interactive and generative text and multimedia works using Harrell's Alloy algorithm for conceptual blending. These works include poetic, animated, and documentary systems that themselves produce different output each time they are run. While GRIOT allows authors to implement narrative and poetic structures (e.g., plots), a major contribution of GRIOT is its orientation toward the dynamic generation of content resulting from modeling aspects of figurative thought that can be described formally. That is, GRIOT allows authors to fix elements such as narrative structure while varying output in terms of theme, metaphor, emotional tone, and related types of what is here called "figuration" (results of figurative thought).

Rather than being based on a single knowledge base or ontology, as is the case with many classic AI systems, GRIOT creates blends between different ontologies (Harrell 2006, 2007). Indeed, a key feature of GRIOT is the ability of authors to construct subjective ontologies based in specific authorial worldviews, elements of which are then blended in a manner that maintains coherence based on several formal optimality principles inspired by a subset of those proposed by Gilles Fauconnier and Mark Turner (1999). This approach allows for novel, surprising, and valuable content to be generated that retains conceptual coherence. GRIOT, like MEXICA, has also been used to implement cultural forms of narrative that are not often privileged in computer science, in this case oral traditions of narrative from the African diaspora (Harrell 2007a). This is important because some forms of oral narrative have more in common with narratives in virtual worlds than the graphocentric (text-biased) forms of narrative privileged in most research in the field of narratology in literary studies.

The implemented GRIOT system, and experience with it, have informed the development of Fig-S, a component of Slant that proposes what types of figuration, mainly metaphor, will be used in telling the story. GRIOT also inspires GRIOT-Gen, the component that generates natural language representations for figuratively enriched versions of particular actions after the story representation is completely developed (see also Goguen and Harrell 2008).

Curveship. This is an interactive fiction system that provides a world model (of characters, objects, locations, and things that happen) while also modeling the narrative discourse, so that the narration and description of the simulated world can change (Montfort 2009, 2011). Curveship can tell events out of order, using flashback and other techniques, and can tell the story from the standpoint of particular characters and their perceptions and understandings. It is based on Genette's theories (Genette 1983) and incorporates other ideas from narratology. The architecture of Curveship draws on well-established techniques for simulating an IF world, separating these from the subsystem for narrating, which in-

cludes a standard three-stage natural language generation pipeline. To make use of the system, either for interactive fiction authoring or story generation, one specifies high-level narrative aspects; the system does appropriate content selection, works out grammatical specifics, and realizes the text with, for instance, proper verb formation.

Some world simulation abilities and the narrative text generation capabilities of Curveship are used directly in Slant in Curveship-Gen, the component that outputs the finished, realized story.

The Slantstory XML Format

Connecting different systems so that they can work together means establishing shared representations. For Slant, that representation is an XML format called Slantstory. It contains all of the information that is needed in the final steps to represent each action and realize the story, meaning that it must contain sufficiently granular information about the plot, the narrative discourse, and the types of conceptual blending that are to be done. This information is not only needed at the last stage, where the generation of text is done. It can also be read by the different subsystems during story generation, when the story is not yet complete, and can influence the next stage of story augmentation. Because of this, Slantstory is a format not only for representing entire, complete stories but also for representing partial stories, the composition of which is in progress. In the current implementation, subsystems can augment a story and declare it complete, but cannot revise or remove what has already been contributed.

To declare a common representation for (both partial and complete) stories, an agreement had to be reached between different perspectives on what the elements of a story are, what is to be represented about each, and how granular the representation of each element is. The Slantstory DTD specifies five elements that occur within the root:

```
<!ELEMENT slantstory
(existents, actions, spin?, genre?, figuration?)>
```

A story cannot be complete without all five of these present, but only existents and actions are required at every stage of story development. The existents are of three types: locations, characters, and things. Actions each have a verb (which might be a phrase such as “try to flee”) and may have any or all of agent, direct object, and indirect object specified. The “instantaneous” tension level, or change in the tension associated with an action, is also represented there. The actions also have a unique ID number which indicates their chronological order in the story world, as in:

```
<action verb="cure" agent="virgin" direct="enemy"
indirect="curative plant" location="Texcoco Lake"
tension="0" id="42" />
```

One challenge in developing and using this blackboard representation involves the different models of existents and

actions that the three foundational systems use. Characters and locations, but nothing like props or “things,” are represented in MEXICA, while Curveship represents all three sorts of existents to provide the type of simulation that is typical in interactive fiction, where objects can typically be acquired, given to other characters, placed on surfaces and in containers, and so on. MEXICA was modified for use in Slant to produce appropriate representations of whatever things were mentioned in actions.

The representation of action was also not consistent between the foundational systems. Curveship has a typology of four actions: Configure (move some existent into, onto, out of, off, or to a different location), Modify (change the state of some existent), Sense (gain information about the world from sensing), and Behave (any other action, not resulting in any change of state in the world). Although they may be quite different, all actions are meant to correspond to a sentence with a single verb phrase when realized. MEXICA’s actions, on the other hand, are not categorized in this way and include many different sorts of representations. There are, for instance, complex actions such as `FAKED_STAB_INSTEAD_HURT_HIMSELF`, indications that an action was not taken such as `NOT_CURE`, and indications that a state is to be described at a certain point such as `WAS_BROTHER_OF`.

The first of these issues, the granularity of action, was handled by developing a mapping between MEXICA actions and Slantstory actions. A limitation of this approach is that actions cannot be inserted into the middle of a series of Slantstory actions that correspond to a single MEXICA action; this is enforced by giving the actions consecutive IDs, so that there is no room to add further actions. Ideally, however, other subsystems would be able to modify the Slantstory representation of actions in any way. The second of these issues brings up the interesting issue of disnarration (Prince 1988), that it is possible in a story to not only tell what has happened but to also tell what did not happen, and that doing so can have an interesting effect on the reader. Disnarration is not the representation of action, however, so it cannot be represented in a straightforward way in a list of actions, and should be handled elsewhere—in the spin element, for instance. Resolving the final issue related to stative information also requires further work, since the system should both represent facts about the story world (probably in the existents element) and when to mention them (probably in the spin element).

GRIOT transforms, for instance, the “agent” and “direct” attributes of an action into conceptual categories. While Slantstory uses a grammatical-sounding model of actions, with direct and indirect objects, Curveship can in fact realize sentences out of these where the agent is the direct object and the “direct” existent is the subject—when it realizes a sentence in the passive, for instance. So, both GRIOT and Curveship treat the seemingly grammatical attributes of ac-

tion in slightly different ways.

Furthermore, the templates that are used to represent sentences in Curveship, which is designed for narrative variation, are not well-designed for the generation of figurative text. Curveship's templates are set up to allow a slot for an agent, for example, which might eventually be filled with "the jaguar knight" "I" "he" or "you" depending upon how narrator and narratee are set and whether the noun phrase is pronominalized. Fig-S, however, may determine that the adjective "enflamed" should be used with this noun phrase because it will participate in the conventional metaphor LOVE IS FIRE. In this case, Curveship-Gen should generate either "the enflamed jaguar knight" "I, enflamed," "he, enflamed," or "you, enflamed." All the possibilities for combinations of figuration (not just the use of an adjective) and all the existing ways that Curveship can generate noun phrases need to be implemented in the next version of Slant.

Verso: Augmenting a Story Based on Genre

Verso, like MEXICA and Fig-S, reads a Slantstory XML file from the blackboard and outputs an updated one. While MEXICA is focused on plot and Fig-S selects an appropriate domain for blending particular representations of action, Verso's operation is based on a model genre. This subsystem operates by:

1. Detecting particular aspects of the in-progress story (typically actions with particular verbs, although possibly series of actions or sets of characters) that indicate the story's suitability to a particular genre, for all known genres.
2. Selecting the genre that is most appropriate.
3. Updating the story using rules specific to that genre. The narrative discourse is always updated by specifying attributes of and elements within "spin." This determines elements such as the focalizer, narrator, time of narrating, rhetorical style, and beginning and/or ending phrases to frame the story. The update can also contribute new actions to the story, which can influence the way that MEXICA continues to develop the plot.

This procedure brings a model of genre awareness into Slant, but it is an unusual process from the standpoint of conventional human creativity. More often than not, an author chooses a genre and then writes or tells something within it, rather than begin with a partial story and finding a genre that suits it. The overall effect, however, is to introduce sensitivity to an important aspect of human creativity.

Verso's model does not seem completely aligned with the direction of genre studies in recent decades. This field has moved from a formalist definitional framework of genre to one that is semiotic, focusing in particular on the "rhetorical study of the generic actions of everyday readers and writers" (Devitt 2008). Recently, genre studies has deemphasized and argued against the idea of genres as distinct

categories with characteristic elements that identify them. Scholars now dispute the idea that characteristics can be identified and summed up to indicate the likelihood that a text is part of a certain genre. They note that few genres have true fundamental elements. Particularly in the case of literary genres (e.g. detective fiction, science fiction, horror, fantasy), even when there seem to be some core characteristics that all works within a category share, almost any "defining" characteristic could be countered by an example work which lacks that element but is still undeniably of that genre. Furthermore, a fundamental dilemma arises in the act of classification itself, the problem of "whether these units exist independently of the taxonomical scheme, or arise as a result of the attempt to classify" (Ryan 1981).

However, these recent concerns pertain most directly to scholarly and critical work; they do not bear upon the way genre is used in literary creativity. Sharp definitions of genre that are developed through writing practice have served many authors well, including Raymond Queneau, who used 99 different genres, modes, or styles to retell the same simple story in *Exercices in Style*. The problem of whether classification compels texts into categories is a problem for analysis, but it is a productive idea for literary creativity. Additionally, as Steve Neale has pointed out, "genres are instances of repetition and difference;" it is precisely through the differentiation from the established norms of a genre that a work can become part of it (Neale 1984). Verso, while making use of those "instances of repetition," also aims to effectively model the production of this necessary difference.

The genres that have been implemented so far are not literary, either in the sense of broad differentiations such as "prose" and "poetry," or in the sense of categories such as "romance," "cyberpunk," "noir," and so on. Instead, Verso uses a broader definition of what constitutes genre, one which includes categories that may very well be alternatively thought of as styles, modes, or even distinct media, and which relate to both fiction and non-fiction as well as to oral and written communication. In the introduction to *Writing Genres*, Devitt provides many examples of the influence of genre in our daily lives, including such wide-ranging categories as the joke, lecture, mystery novel, travel brochure, small talk, sales letter, and, most appropriately, the research paper (Devitt 2008). It is this broader conception of genre, rather than a strictly literary one, that Verso aims to model.

The genres implemented in Verso tend towards the stylistic rather than the thematic. In part due to the pre-existing capabilities of Curveship, and in part because of the domain in which MEXICA operates, the genres used are those that can be identified and produced through changes in the narrative discourse (focalization, time of narrating, order of events in the telling, etc.) rather than the story world domain (which could incorporate dragons, spaceships, magic, etc.).

A concrete example is provided by the "confession"

genre, which casts a story so that it sounds like it is being told to a priest at confession. To determine if this genre is applicable, the system checks to see if one or more actions are likely “sins” (robbing, killing, etc.) based on a list of these. Each “sin” raises the suitability of this genre. If “confession” is selected as the genre to use, the Slantstory XML representation is updated. A “sinner” is located—the agent of the last sinful action. This sinner is specified as the narrator (the “I” of the story). There is no narratee (or “you”), since we presume that the priest was not part of the events that were being told. The time of narrating is set to “after,” which results in past-tense narration, and the “hesitant” style is used, injecting “um” and “er” into the story as if the speaker were nervous and reticent. Finally, a conventional opening (“Forgive me, Father, for I have sinned. It has been a month since my last confession.”) and a conventional conclusion (“Ten Hail Marys? Thank you, Father.”) are added.

The “confession” genre produces plausible and amusing results. Some of this has to do with the formulaic nature of the genre. As one reads additional confessions, the rigid, repetitive opening and conclusion can be amusing, because they model the ritualized interaction of confession. Read in this light, it is only more amusing that ten Hail Marys are always given for penance, whether the penitent tried to swipe something or committed a murder. Finally, because Spanish conquerors came to the Americas and imposed Catholicism on the natives, MEXICA-generated plots that are told in this genre can be read as a comment upon, or at least a provocation about, the colonial history of Mexico. Importantly, these two subsystems did not invent this juxtaposition of the MEXICA and Catholic ritual; rather, humans decided many years ago to develop a story generator about the Mexica and decided recently to develop a “confession” genre template. However, the subsystems’ collaboration as part of Slant involves automatically finding occasions when the juxtaposition of these two is particularly effective. Verso’s work and MEXICA’s work combine in Slant to provide more cultural resonance, to be more surprising and also to be more valuable by virtue of being thought-provoking.

In the current system 10 genres have been implemented: confession, diary, dream, fragments, hangover, joke, letter, memento, memoir, play-by-play, prophecy, and the default “standard” story. These take advantage of only a limited range of Curveship’s narrative variation capabilities. For instance, the focalization of a story can be varied, but we have not yet implemented genres that focalize stories based on particular characters; similarly, Curveship is already capable of narrating with flashbacks and making other more elaborate changes in order. There are now only two prose styles that are used, “excited” for play-by-play and “hesitant” for confession. It would also be straightforward to elaborate the Slantstory representation and to modify Curveship-Gen to allow for expression that better relates to a wider variety of genres. In discussions so far we have already listed more

than 100 genres, most of which we believe will be to some extent recognizable and applicable to the short stories produced by Slant.

Fig-S and GRIOT-Gen for Figuration

Fig-S reads a Slantstory XML file from the blackboard and updates it to include metaphorical content. Metaphor here can be understood as an asymmetrical conceptual blend in which all content from one domain called the “target space” is integrated with a subset of content from another called the “source space” (Grady, Oakley, and Coulson 1999). Fig-S currently implements ontologies representing several domains empirically identified as important in poetry such as “death” and “love” (Lakoff and Turner 1989) that can be used to generate metaphors such as REJECTION IS DEATH or ADMIRATION IS LOVE.

Fig-S begins by processing each of the actions from the Slantstory XML file to assess whether they will be replaced by metaphorical versions of the same action. Currently, there are two modes in which this processing can be done. If ONE-METAPHOR is set to true, then the Slantstory is analyzed to find which single source domain is appropriate to map onto the greatest number of actions in order to produce metaphors. Otherwise, each action will be analyzed individually in order to find an appropriate source domain to map onto it. The first mode typically results in more coherent output, the second mode typically results in a greater degree and variety of metaphorical output. As an example of an action that has been mapped onto by the source domain LOVE in order to produce a metaphorical action, the Slantstory action:

```
<action agent="virgin" direct="princess" id="61"
location="Texcoco Lake" tension="40" verb="get
jealous of" />
```

could be processed by Fig-S and added to the Slantstory as:

```
<figuration domain="fire">
<blend id="61" verb="get jealous of/burn for"
agent="virgin/burning one*agent"
direct="princess/hot*direct">
</figuration>
```

While Fig-S currently has implemented simple, metaphorical form of blending as a first step, it could be extended to use a more robust blending algorithm such as Alloy, or even to extend Alloy to result in even more novel, surprising, and/or culturally valued blends using an extended set of optimality principles.

GRIOT-Gen is used to produce specific output template from metaphorical actions in a Curveship-Gen format. For example, the metaphorical action above could be realized in a number of ways. The default produced by GRIOT-Gen, for a story in which neither virgin nor princess are narrator or narratee, would be structured as:

'61': 'the burning virgin [become/v] jealous-of the incendiary princess',

however, it can alternatively be structured as:

'61': '[@virgin/s] like burning [get/v] jealous of the incendiary [princess/o]',

if there is a preference for a simile-oriented style for the subject. It is also possible to use a “source-element/target-element” structure as in:

'61': 'the burning/virgin [get/v] jealous of and [burn/v] for the incendiary/princess'

to be very explicit about every element that has been integrated. GRIOT-Gen currently has multiple such exposition forms implemented and is easily extensible.

Slant's First Stories

In the current system some spin (narrative discourse specification) is necessary, although it may simply involve the default settings, while figurative action representations are optional. To begin with, this amusing but flawed story was generated without figuration, but with contributions from MEXICA and Verso:

Forgive me, Father, for I have sinned. It has been a month since my last confession. An enemy slid. The enemy fell. The enemy injured himself. I located a curative plant. I cured the enemy with the curative plant. The tlatoani kidnapped me. The enemy sought the tlatoani. The enemy travelled. The enemy, um, looked. The enemy found the tlatoani. The enemy observed, uh, the tlatoani. The enemy drew a weapon. The enemy attacked the tlatoani. The enemy killed the tlatoani with a dagger. The enemy rescued me. The enemy entranced, uh, me. I became jealous of the enemy. I killed the enemy with the dagger. I killed myself, uh, with the dagger. Ten Hail Marys? Thank you, Father.

The “sinner” who narrates the story dies, a problem which can also crop up when the “diary” genre issued. Since Verso can assign the genre of the story before the plot is complete, there was initially no way that Verso be sure that the character it selects as narrator will not die. This requires an interaction between the genre-selecting system, Verso, and the plot-generating system, MEXICA. We implemented an additional set of constraints on how the plotting could be done which either require or prohibit that a certain tension, as defined in MEXICA, arise. One of these tensions is “actor dead,” letting Verso prohibit a narrator’s death.

A story with figuration follows. This one is generated without the constraint for a single conventional metaphor to be used (ONE-METAPHOR is false), so there is a colorful diversity of less consistent metaphors. The genre chosen is “play-by-play,” based on sports commentary, which may be a suitable one for the range of metaphor that is used:

This is Ehecatl, live from the scene. The cold-wind eagle knight is despising the icy jaguar knight! The cold-wind

jaguar knight is despising the chilling eagle knight! Yes, an eagle knight is fighting a jaguar knight! Look at this, the eagle knight is drawing a weapon! Look at this, the eagle knight is closing on the jaguar knight! The gardener eagle knight is wounding the weed jaguar knight! And now, the jaguar knight is bleeding! Yes, the consumed eagle-knight is panicking! And, eagle knight is hiding! Holy -- the snowflake slave is despising the chilling jaguar knight! The freezing-wind jaguar knight is despising the cold slave! And, yes, the cold-wind slave is detesting the chilling jaguar knight! A slave is curing the jaguar knight! And, the slave is returning to the city! And, the jaguar knight is suffering! The frozen jaguar knight is dying! Back to you!

MEXICA’s stative descriptions of characters could probably be mentioned more rapidly, or perhaps not at all, to keep the action going. This could be done with an existing facility in Slantstory for omitting actions when narrating. This story would also benefit from pronominalization, which Curve-ship-Gen is capable of but which would need to be either turned on for all stories or specified at an earlier stage.

Slant's Research Potential

We plan to further develop the system we have initiated to explore new ways that computational creativity researchers can collaborate, new models of storytelling that abstract different sorts of expertise and emphasis, and new ways to compare the importance of and interaction between different aspects of story. We intend that the system will be used for empirical studies of how people receive generated stories and will also be brought into literary and artistic contexts.

Using the Slantstory XML blackboard, many different subsystems can be developed for Slant, which will allow Slant to be run with any subset of them. For instance, if Verso is turned off so that the specification of the narrative discourse is not done by that subsystem, either a default narrative discourse specification could be used (as would be the case now, since Verso is the only subsystem that updates this aspect) or that specification can be built up by one or more other subsystems. This allows the effect of each subsystem, in the context of Slant overall, to be carefully examined. Readers of stories generated under different conditions could be asked not only to rank the outputs in terms of quality, but also to comment on what they thought about particular elements (such as characters) and high-level qualities (whether the story was funny, for instance, or whether it seemed plausible).

The project can also facilitate a broader collaboration between researchers of story generation. As long as researchers find the Slantstory XML representation adequate for their purpose, they can develop new subsystems that help to build stories based on other theories or concerns. For instance, a researcher interested in how creativity occurs in social contexts could model the process in a unit that reads from and writes to the blackboard and models social influ-

ence and awareness. As just discussed, this new system could be tried in many combinations with existing systems and the outputs could be compared. This would help to show not only the importance of social creativity as modeled in this particular subsystem, but also how creativity of this sort interacts with plot generation using the engagement-reflection cycle, figuration based on conventional metaphors, and awareness of genre.

We also anticipate that Slant will supply stories for exhibition and publication in arts contexts, and the functional system itself could be part of a digital media, electronic literature, or e-poetry exhibit. In this way, Slant can contribute to creative practice, and reactions and discussion in this context can help us further develop a system that relates to contemporary literary concerns.

Acknowledgements

Thanks to Clara Fernandez-Vara and Ayse Gursoy for their discussions of genre and of early ideas about Slant.

References

- Boden, M.A. 2004. *The Creative Mind: Myths and Mechanisms*. 2nd Ed. London and New York: Routledge.
- Devitt, A.J. 2008. *Writing Genres*. Carbondale: Southern Illinois University Press.
- Fauconnier, G. 1999. "Methods and Generalizations." In *Cognitive Linguistics, Foundations, Scope, and Methodology*, ed. T. Janssen and G. Redeker, 95–127. The Hague: Mouton De Gruyter: 96.
- Fauconnier, Gilles, and Turner, M. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.
- Goguen, J., and Harrell, D.F. 2008. Style, computation, and conceptual blending. In Argamon, S., and Dubnov, S., eds., *The Structure of Style: Algorithmic Approaches to Understanding Manner and Meaning*. Berlin: Springer-Verlag. 291–316.
- Grady, J. E., Oakley, T., and Coulson, S. 1999. Blending and Metaphor. In *Metaphor in Cognitive Linguistics*, ed. Gerard Steen and Ray Gibbs, 101–124. Amsterdam: John Benjamins.
- Genette, G. 1983. *Narrative discourse: An essay in method*. Cornell University Press.
- Harrell, D.F. 2006. Walking blues changes undersea: Imaginative narrative in interactive poetry generation with the GRIOT system. In *Proceeding of the AAAI 2006 Workshop in Computational Aesthetics: Artificial Intelligence Approaches to Happiness and Beauty*, 61–69. AAAI Press.
- Harrell, D.F. 2007. GRIOT's tales of haints and seraphs: A computational narrative generation system. In Wardrip-Fruin, N., and P. Harrigan, eds., *Second Person: Role-Playing and Story in Games and Playable Media*. Cambridge, MA: MIT Press, 2007. 177–182.
- Harrell, D. F. 2007a. "Cultural Roots for Computing: The Case of African Diasporic Orature and Computational Narrative in the GRIOT System," *Fibreculture Journal*, Vol. 11, http://journal.fibreculture.org/issue11/issue11_harrell.html
- Lakoff, G. and Turner, M. 1989. *More than Cool Reason—A Field Guide to Poetic Metaphor*. Chicago: University of Chicago Press.
- Montfort, N. 2009. Curveship: An interactive fiction system for interactive narrating. In *Proceedings of the NAACL HLT Workshop on Computational Approaches to Linguistic Creativity*, 55–62.
- Montfort, N. 2011. Curveship: Adding control of narrative style. In *Proceedings of the Second International Conference on Computational Creativity*, 163.
- Montfort, N., and Pérez y Pérez, R. 2008. Integrating a plot generator and an automatic narrator to create and tell stories. In *Proceedings of the 5th International Joint Workshop on Computational Creativity*. http://nickm.com/if/mexica-nn_ijwcc08.pdf
- Neale, S. 1980. *Genre*. London: British Film Institute.
- Pérez y Pérez, R., Ortiz, O., Luna, W. A., Negrete, S., Peñaloza, E., Castellanos, V., and Ávila, R. 2011. A system for evaluating novelty in computer generated narratives. In *Proceedings of the Second International Conference on Computational Creativity*, 63–68.
- Pérez y Pérez, R., and Sharples, M. 1999. MEXICA: A computational model of the process of creative writing. In *Proceedings of the AISB Symposium on Creative Language: Humour and Stories*, 46–51.
- Pérez y Pérez, R., and Sharples, M. 2001. MEXICA: A computer model of a cognitive account of creative writing. *Journal of Experimental and Theoretical Artificial Intelligence*. 13(2): 119–139.
- Pérez y Pérez, R., and Sharples, M. 2004. Three computer-based models of storytelling: BRUTUS, MINSTREL and MEXICA. *Knowledge Based Systems Journal*. 17(1): 15–29.
- Polanyi, L. 1989. *Telling the American Story: A Structural and Cultural Analysis of Conversational Storytelling*. Cambridge, MA: The MIT Press.
- Prince, G. 1988. The disnarrated. *Style* 22(1): 1–8.
- Ryan, M-L. 1981. The why, what and how of generic taxonomy. *Poetics* 10: 109–126.
- Ryan, M-L. 1991. *Possible worlds, artificial intelligence, and narrative theory*. Bloomington: Indiana University Press.