Computational Approaches to Storytelling and Creativity

Pablo Gervás

■ *This article deals with computational approaches to storytelling, or the production of stories by computers, with a particular attention on the way human creativity is modeled or emulated in computational terms. Features relevant to creativity and to stories are analyzed, and existing systems are reviewed under the light of that analysis. The extent to which they implement the key features proposed in recent models of computational creativity is discussed. Limitations, avenues of future research, and expected trends are outlined.*

Creativity and storytelling are two ideas that one does not naturally associate with computers. Yet over the last few years there has been a surge of research efforts concerning the combination of both subjects. This article tries to shed light on these efforts. In carrying out this program, one is handicapped by the fact that, as words, both *creativity* and *storytelling* are severely lacking in the precision one expects of words to be used for intellectual endeavor. If a speaker were to mention either word in front of an audience, each person listening would probably come up with a different mental picture of what is intended. To avoid the risks that such vagueness might lead to, an initial effort is made here to restrict the endeavor to those aspects that have been modeled computationally in some model or system. The article then proceeds to review some of the research efforts that have addressed these problems from a computational point of view.

The Creative Process

Thinking of "creativity" evokes several ideas that seem to go together. It generally brings up the idea of someone generating something new. But it also has connotations that whatever is generated must be somewhat unexpected or different from what others might have produced. There is also an implicit restriction that what is generated satisfy some goal, though in many cases the particular goal implied is not altogether clear. The fact that someone is involved indicates we are reviewing an explicit action by some agent that we shall refer to as the *creator*. A sunset may generate a totally new combination of colors, possibly unexpected, but it would not be considered creative. The fact that something is generated indicates that the creative action must have a result, let us call it the *output*. In general terms this need not be a physical object (it may be an idea or a method), but something must result from the action that can be perceived or evaluated. The expectation of novelty has to be qualified further, because something may be new to some people and not to others. This introduces at least a second agent (usually many more are involved) that perceives or evaluates the result. Novelty must be considered from the relative point of view of the different agents involved. There may be more than one point of view to consider. For paintings, there may be critics and buyers; for movies there may be reviewers and filmgoers. This should include at least the creator and the particular person or persons for which the creative action is intended, which I shall refer to as the *audience*. Competing creators may also be considered. The idea of unexpectedness involves some expectations that get defeated. These expectations may arise from knowing what the creator has done before or from knowing what other creators have done before or from a generic canon established in some way for the type of output considered. Again, this must be relative to the perception of a particular person or set of persons.

The implicit understanding that the output meet some goal has sometimes been formulated in terms of requiring that the result be useful or have some value. The vagueness of these descriptions is at the root of the elusiveness of creativity as a concept. As a word, *creativity* seems to have been invented in the 19th century in an attempt to cover the different concepts of innovation that were accepted in art and science (Weiner 2000). In doing so, many relevant differences may have been swept under the carpet. The restrictions that an output must satisfy to be considered useful differ greatly between art and science. Among the design community, for instance, a design is expected to satisfy a set of requirements (usually involving a description of expected functionality). This partly determines its value or usefulness, but aesthetic considerations must also play a role. Within the arts, the restrictions applicable to a piece of music, a painting, or a novel may be radically different. Let us accept that the output of a creative process must satisfy some additional requirements related to usefulness, functionality, or value. These would have to be carefully specified in each particular case (and probably will give rise to altogether different creative phenomena).

If one is to consider the creative process as a productive activity, especially if one is to address it from an engineering point of view, it is worth considering its inputs and outputs and also whether feedback is contemplated.

To produce an output a creator works on a given material that we may consider as an input to the process. We want to take this factor into account when considering a particular creative process. The introduction of new material in the input may influence the perceived novelty of the output with little intervention from the creative process employed. Restricting the possible input to a given material may create particular expectations.

In a system that contemplates feedback, past outputs are taken into account when processing new inputs. The existence of feedback implies that the same input fed repetitively to the system may produce different outputs in each subsequent run.

Features of Interest in a Story

A story is a highly complex intellectual product that exercises a wide range of the cognitive abilities of humans, involving as it usually does perceptions of time and space, attribution of knowledge to particular characters, identifying character goals, validating character plans to achieve the goals, accepting plan failure in the face of obstacles, attributing feelings to characters, associating character intentions with feelings, developing empathy with characters, and including the underlying skill of natural language understanding. It may seem too ambitious to consider that a computer might address all these tasks successfully in solving a given problem. And yet a large number of these tasks have been the subject of intense study in AI over the last century.[1] This fact suggests that the parts may slowly be appearing to put together an interesting model of the tasks involved.

For a story to be satisfactory, many elements (characters, personality, knowledge that each character has, goals of each character, feelings or emotions of characters, dialogues between the characters, and so on) have to be combined in an intricate manner that is difficult to specify. This presents a serious problem for addressing creativity in the context of stories. The requirement that the output of a creative process meet some particular goal is, in the case of stories, very tightly coupled with the different ways in which these elements combine with one another to build the whole. In literary studies there have been many attempts to identify the features that give value to a story. If a consensus existed on what these might be, achieving them might be taken as the goal that the output of the creative process has to meet. There is currently no such consensus. Rather than attempt a single-handed reduction of centuries of literary theory in a few paragraphs, I will restrict the discussion in this article to those elements of stories that have at some stage been modeled in a computational storyteller. These are identified in a later section. Nevertheless, it is useful to outline some basic concepts of literary theory to guide the analysis: chronology, focalization, causality, and the distinction between *fabula* (what is told) and *discourse* (how it is told).

Labov's (1972) definition of *minimal narrative* as two states and a transition or movement between the two states provides a starting point to consider what the minimal model of a story might be. This captures the basic intuition that a story involves events happening over time. Time is therefore a basic element in stories. The simpler kind of stories that we can think of usually involve a series of events that are told in chronological order.

Narratologists, who specialize in the study of narrative, consider the concept of *focalization* as the way in which a narrator restricts what he is telling about a particular scene to what might have been perceived by someone present in that scene.[2] This may be one of the characters, if the scene is told in the first person, or the narrator himself as if he had been present (if the story is told in the third person). This has an interesting implication in the fact that, through focalization, narrative discourse (and thereby the structure of stories) is influenced by the perception of space: events that take place simultaneously in different locations that cannot be perceived at the same time (these may be different cities but also different neighboring rooms separated by a wall) usually require different narrative threads.

A crucial ingredient in understanding a story is causality relations, whereby an event *B* happens as a result of event *A* having occurred earlier. This corresponds to the intuition underlying the classic argument by E. M. Forster (1927), whereby "The king died and then the queen died" is less interesting as a story than "The king died and then the queen died of grief," on the grounds that the second example involves relations of causality between the events being told.

According to many theorists, narrative has two components: what is told (what narrative is: its content, consisting of events, actions, time, and location) and the way it is told (how the narrative is told: arrangement, emphasis or deemphasis, magnification or diminution, of any of the elements of the content). These have been named different ways by different researchers, and there are also alternative analyses that postulate different subdivisions. Even between theories that agree on having just two levels of analysis there seem to be many subtleties that cast doubt on whether the same thing is meant by the different words. To avoid ambiguity, I will restrict my analysis here to two levels of conceptual representation of a story and refer to these as the *fabula* (what is told) and the *discourse* (the way it is told) for a given story. In some cases, to be able to describe the process leading to these elements, the concept of a *world* (a broader representation of events, of which the fabula is a selected subset) is useful.

Computational Creativity

Many efforts over the recent years that address the study of creativity from a computational point of view acknowledge as a predecessor the work of Margaret Boden (1990). Boden proposed that artificial intelligence ideas might help to understand creative thought. This idea was taken up by a number of artificial intelligence researchers and gave rise to a research line that attempts to model or reproduce creative thought in computer systems. Some of Boden's ideas have had great influence in later work. One important idea was the distinction between historical and psychological views of creativity. Historical creativity (H-creativity) involves the production of ideas that have not appeared before to any one else in all human history. Psychological creativity (P-creativity) involves the production by a given person of ideas that have not occurred before to that particular person. This distinction is important because it implies that, unless a computer program is given access to historical data (and generally provided with means for social interactions with other creators), it will only be capable of P-creativity. Another important contribution arose from Boden's application of the concepts of artificial intelligence to the understanding of creativity. Boden formulated the search of ideas in terms of search over a conceptual space. Such a conceptual space would be defined by a set of constructive rules. The strategies for traversing this conceptual space in search of ideas would also be encoded as a set of rules. Based on this model, Boden distinguished between exploratory creativity (the rules for traversing the space are simply applied in search for an idea that has not been found before) and transformational creativity (the rules for traversal are changed so that the search can now reach areas of the conceptual space that were not accessible before).

Sharples (1999) presents a description of writing understood as a problem-solving process where the writer is both a creative thinker and a designer of text. The account draws heavily on Boden's analysis. For Sharples, the universe of concepts that can be explored in the domain of writing could be established in a generative way by exhaustively applying the rules of grammar that define the set of well-formed sentences. The conceptual space on which a writer operates is a subset of this universe identified by a set of constraints that define what is appropriate to the task at hand.

Sharples also provides a description of how the typical writer alternates between the simple task of exploring the conceptual space defined by a given set of constraints and the more complex task of modifying such constraints to transform the conceptual space. Sharples proposes a cyclic process moving through two different phases: engagement and reflection. During the engagement phase the constraints are taken as given and the conceptual space defined by them is simply explored, progressively generating new material. During the reflection phase, the generated material is revised, and constraints may be transformed as a result of this revision.

Wiggins (2006) takes up Boden's idea of creativity as search over conceptual spaces and presents a more detailed theoretical framework that attempts to clarify the issue of whether creativity can be reduced to good old-fashioned AI search (GOFAI). By specifying formally the different elements involved (the universe of possible concepts, the rules that define a particular subset of that universe as a conceptual space, the rules for traversing that conceptual space), Wiggins points out that the rules for traversing a conceptual space may lead to elements in the universe but outside the definition of the conceptual space. This kind of behavior would be considered a design error in an AI system but is acceptable as a characterization of a creative situation, which need not include restrictions that the search space be explicitly defined or traversal functions be consistent and complete with respect to a defined search space. In fact, definitions of search space and traversal function in a creative setting are not only particular to a given creator and different from those used by others, but also constantly in flux. This description strengthens the argument for exploratory creativity as distinct from old-fashioned search. In a context where all the elements in the framework are in constant change, it would be at best risky to assume that the rules for traversal of the search space always produce elements within the conceptual space originally defined. As a corollary of this description, GOFAI search could be considered an instance of Wiggins's framework in which, by design, the rules for defining the conceptual space are used as rules for traversing it.

Ritchie (2007) addresses another important issue in the development of creative programs, that of evaluating when a program can be considered creative. He does this by outlining a set of empirical criteria to measure the creativity of the program in terms of its output. He makes it very clear that he is restricting his analysis to the questions of what factors are to be observed, and how these might relate to creativity, specifically stating that he does not intend to build a model of creativity. Ritchie's criteria are defined in terms of two observable properties of the results produced by the program: novelty (to what extent is the produced item dissimilar to existing examples of that genre) and quality (to what extent is the produced item a high-quality example of that genre). To measure these aspects, two *rating schemes* are introduced, which rate the typicality of a given item (item is typical) and its quality (item is good). Another important issue that affects the assessment of creativity in creative programs is the concept of *inspiring set*, the set of (usually highly valued) artifacts that the programmer is guided by when designing a creative program. Ritchie's criteria are phrased in terms of: what proportion of the results rates well according to each rating scheme, ratios between various subsets of the result (defined in terms of their ratings), and whether the elements in these sets were already present or not in the inspiring set.

Jennings (2008) introduced computationally plausible modeling of the fact that most human creativity takes place with the creator embedded in a broader society of other creators and critics, and that this context affects significantly the creation of new artifacts. To capture the way in which humans react to these constraints, Jennings defines the concept of *creative autonomy*, which requires that a system be able to evaluate its creations without consulting others, that it be able to adjust how it makes these evaluations without being explicitly told when or how to do so, and that these processes not be purely random. The model he proposes relates the evaluation of a system's creations to its perception of how other members of its social context are likely to evaluate them. Changes in how this evaluation is carried out may be triggered by the need to align personal evaluations with other members of the society or as a side effect of trying to justify past evaluations. Creative autonomy is therefore argued to emerge out of the interactions with multiple critics and creators, rather than from meditative isolation.

A Brief History of Storytelling Systems

There are currently many storytelling systems in existence. For this review, systems that generate classic sequential stories have been selected on the basis of whether they pioneered the introduction of a particularly significant feature (with special interest in those related to creativity). The amount of detail available for each system varies greatly, being limited to self-reported statements for older systems and much better covered with material on the Internet for more recent systems. In each case, only features distinguishing the system from others or particularly related to creativity are mentioned. Examples of story output are given for some systems where small enough significant fragments were available.

Klein's Novel Writer (1973)

The first storytelling system for which there is a record is the Novel Writer system developed by Sheldon Klein et al. (1973). Novel Writer created murder stories within the context of a weekend party. It relied on a microsimulation model where the behavior of individual characters and the events that transpired were governed by probabilistic rules that progressively changed the state of the simulated world (represented as a semantic network). The flow of the narrative arises from reports on the changing state of the world model. However, the sequence of scenes was hard wired in the code to correspond to the expected development of a weekend party, with the simulation only accounting for the interplay between the characters that fleshes out the plot. This sequence of scenes could be considered an instance of a primitive story grammar. A description of the world in which the story was to take place was provided as input. The particular murderer and victim depended on the character traits specified as input (with an additional random ingredient). The motives arise as a function of the events during the course of the story. The set of rules is highly constraining and allows for the construction of only one very specific type of story.

Novel Writer suggests a way to model focalization through the use of private semantic universes. This is only sketched in terms of how operating systems may allow the system to keep copies on disk of a private universe, load it into memory, treat it as the total universe while it is resident in memory, operate with it as required, and then store it back on disk—presumably to be able to manage a different private universe. No mention is made of how material from one private universe might interact with material from another private universe.

Here is an example story (this is just a single episode within a 2,100 word larger story):

> The day was Monday. The pleasant weather was sunny. Lady Buxley was in the park. James ran into Lady Buxley. James talked with Lady Buxley. Lady Buxley flirted with James. James invited Lady Buxley. James liked Lady Buxley. Lady Buxley liked James. Lady Buxley was with James in a hotel. James caressed Lady Buxley with passion. James was Lady Buxley's lover. Marion following saw the affair. Marion saw the affair. Marion was jealous.

Meehan's Talespin (1977)

Talespin (Meehan 1977) was a system that told stories about the lives of simple woodland creatures. Talespin combined forward chaining (from events to their consequences) and backward chaining (from desired outcomes expressed as goals that resulted from a previous event, to the particular events that will lead to the outcome). Goals could also be decomposed into subgoals during the backward-chaining mode.

In this way, Talespin introduced character goals as triggers for action. It also introduced the possibility of having more than one problem-solving character in the story (and it introduced separate goal lists for each of them). Another important problem introduced by Talespin was that of character perception (a concept of "noticing" is addressed), and maps of physical space are kept and used to compute the story. Complex relations between characters were modeled (competition, dominance, familiarity, affection, trust, deceit, and indebtedness). These relations acted as preconditions to some actions and as consequences to others. This constitutes a simple model of character motivation. Personality of characters was modeled in terms of degrees of kindness, vanity, honesty, and intelligence.

Meehan discusses what makes a story valid (existence of a problem, degree of difficulty in solving the problem, and nature or level of problem solved), but this seems to be an evaluation procedure external to the program.

An example Talespin story is given in the following extract:

> John Bear is somewhat hungry. John Bear wants to get some berries. John Bear wants to get near the blueberries. John Bear walks from a cave entrance to the bush by going through a pass through a valley through a meadow. John Bear takes the blueberries. John Bear eats the blueberries. The blueberries are gone. John Bear is not very hungry.

Dehn's Author (1981)

Author (Dehn 1981) was a program intended to simulate the author's mind as she makes up a story. Dehn claimed story worlds are developed by authors as a post hoc justification for events that the author has already decided have to be part of the story. In explaining this process, author goals are first mentioned. An author may have particular goals in mind when he sets out to write a story. But even if she does not, it is accepted that a number of metalevel goals drive or constrain the storytelling process. These concern issues such as ensuring that the story is consistent, that it is plausible, that characters be believable, that the attention of the reader is retained throughout the story. These may translate at a lower level into subgoals concerning situations into which the author wants to lead particular characters or the role that particular characters should play in the story. A story is understood as "the achievement of a complex web of author goals." These goals contribute to give the story its structure, guiding the construction process, but they are not visible in the final story. According to this model, a large part of the work of making up a story is the successive reformulation of author goals. This is captured by the concept of *conceptual reformulation*: initial idea gets reformulated into kernel episode, that into a succession of episodes, a characterization gets reformulated as an episode that illustrates it, a change in the relation between two characters gets reformulated as a dialogue that triggers that change. Some example high-level author goals are given: make the story plausible, make the story dramatic, illustrate key facts.

Author attempts to model an author's mind in the kind of knowledge that it applies in making up stories (facts about the story world already constructed but also memorable episodes, characters, and so on from the author's experience). It also tries to model the human mind in the way its knowledge is organized and how it is accessed by following theories of how human memory works (Kolodner 1980, Schank 1982).

Dehn considers story generation to be a process of creative reasoning, and as such it should capture two general characteristics: the degree to which the process is deliberate and the degree to which it is serendipitous. To model this, Dehn postulates two different metagoals: achieving the current narrative goal and finding better narrative goals to pursue. It is this second metagoal that guarantees the directed-serendipitous duality, allowing for changes in direction when unforeseen opportunities arise.

Lebowitz's Universe (1983)

Universe (Lebowitz 1983) modeled the generation of scripts for a succession of TV soap opera episodes (a large cast of characters play out multiple, simultaneous, overlapping stories that never end). Universe is the first storytelling system to devote special attention to the creation of characters. Complex data structures are presented to represent characters, and a simple algorithm is proposed to fill these in partly in an automatic way. But the bulk of characterization is left for the user to do by hand.

Universe is aimed at exploring extended story generation, a continuing serial rather than a story with a beginning and an end. It is in a first instance intended as a writer's aid, with additional hopes to later develop it into an autonomous storyteller. Universe addresses a question of procedure in making up a story over a fictional world: whether the world should be built first and then a plot to take place in it, or whether the plot should drive the construction of the world, with characters, locations, and objects being created as needed. Lebowitz declares himself in favor of the first option, which is why Universe includes facilities for creating characters independently of plot, in contrast to Dehn who favored the second.

The actual story-generation process of Universe (Lebowitz 1985) uses planlike units (plot fragments similar to the plot units of Lehnert [1981]) to generate plot outlines. Treatment of dialogue and low-level text generation are explicitly postponed to some later stage. Plot fragments provide narrative methods that achieve goals, but the goals considered here are not character goals, but author goals. This is intended to allow the system to lead characters into undertaking actions that they would not have chosen to do as independent agents (to make the story interesting, usually by giving rise to melodramatic conflicts). Actual system operation is similar to decompositional planning. The system keeps a precedence graph that records how the various pending author goals and plot fragments relate to each other and to events that have been told already. To plan the next stage of the plot, a goal with no missing preconditions is selected and expanded. Search is not depth first, so that the system may switch from expanding goals related with one branch of the story to expanding goals for a totally different one. When selecting plot fragments or characters to use in expansion, priority is given to those that achieve extra goals from among those pending.

An interesting point in Universe is that, being a story with no recognizable ending, the system alternates between planning a continuation for the plot so far and telling the accumulated extension of the plot since it last told something.

Of particular interest for creativity, Universe included a mechanism for automatically expanding its plot fragment library by creating new plot fragments. This was done by generalizing existing ones and then instantiating the resulting structure to build new ones. To avoid information loss, this process had to be guided by a causal analysis of the initial plot fragment, which the abstraction and instantiation operations had to respect. To ensure a certain flavor of plot fragments was maintained after generalization, only some of the features were generalized at a time. This process lead to the identification of fundamental author goals, such as "maintain romantic tension" and "keep the story moving." This would drive the story generation, but the only justification given for these goals is that they seem validated by experience with melodramatic stories.

Turner's Minstrel (1993)

Minstrel (Turner 1993) was a computer program that told stories about King Arthur and his knights of the round table. It was the first storytelling to address specifically issues of creativity (it is explicitly described in Turner's thesis as "a computer model of creativity and storytelling"). The program was started on a moral that was used as seed to build the story. Minstrel created stories about one-half to one page in length. According to its author, Minstrel could tell about 10 stories of this length, and it could create a number of shorter story scenes.

Minstrel uses building units consisting of goals and plans to satisfy them. These operate at two different levels: in terms of author goals and in terms of character goals. Story construction in Minstrel operates as a two-stage process involving a planning stage and a problem-solving stage. The planning stage operates on an author-level agenda that stores author-level goals. The process consumes the goals in the agenda by either breaking them into smaller author-level goals (which are stored in the agenda) or passing them on to the problem-solving stage (which attempts to solve them by adding the required ingredients to the story). Two operations of particular interest should be mentioned. One involves solving author-level goals, which takes the form of instantiating a set of partially complete character schemas. This is done by querying episodic memory with partially complete schemas and using the results to instantiate the corresponding author-level goal. Another is the way opportunistic goals are triggered. Each time a new scene is created, Minstrel revises it to check whether it provides an opportunity to apply one of the author-level goals that ensures consistency or introduces one of the desired literary motifs. This matches Dehn's second metagoal of looking for new author goals to follow.

The process of querying episodic memory is handled by *transform recall adapt methods* (TRAMs). Basic TRAMs just pass the query as it stands to episodic memory and return any matching schemas found. However, in cases of failure, more complex TRAMs operate by applying a basic modification to the input query, querying episodic memory with the resulting new query, and returning an adaptation of any results obtained by reversing the modification applied to the original query. TRAMs can be linked into a chain (leading to a chain of successive adaptation operations). This captures a concept very similar to the generalization procedure suggested by Lebowitz, and it is the basis for Turner's claim of creativity in Minstrel.

An example Minstrel story follows:

The Vengeful Princess

Once upon a time there was a Lady of the Court named Jennifer. Jennifer loved a knight named Grunfeld. Grunfeld loved Jennifer.

Jennifer wanted revenge on a lady of the court named Darlene because she had the berries which she picked in the woods and Jennifer wanted to have the berries. Jennifer wanted to scare Darlene. Jennifer wanted a dragon to move towards Darlene so that Darlene believed it would eat her. Jennifer wanted to appear to be a dragon so that a dragon would move towards Darlene. Jennifer drank a magic potion. Jennifer transformed into a dragon. A dragon moved towards Darlene. A dragon was near Darlene.

Grunfeld wanted to impress the king. Grunfled wanted to move towards the woods so that he could fight a dragon. Grunfeld moved towards the woods. Grunfeld was near the woods. Grunfeld fought a dragon. The dragon died. The dragon was Jennifer. Jennifer wanted to live. Jennifer tried to drink a magic potion but failed. Grunfeld was filled with grief.

Jennifer was buried in the woods. Grunfeld became a hermit.

MORAL: Deception is a weapon difficult to aim.

Pérez y Pérez's Mexica (1999)

Mexica (Pérez y Pérez 1999) was a computer model designed to study the creative process in writing in terms of the cycle of engagement and reflection (Sharples 1999). It was designed to generate short stories about the early inhabitants of Mexico.

Mexica relies on certain structures to represent its knowledge: a set of *story actions* (defined in terms of preconditions and postconditions) and a set of *previous stories* (stated in terms of story actions). Mexica stands out from other systems in that it actually builds its own set of schemas from the set of previous stories. A single type of knowledge structure, known as a *story-world context* (SWC), is used to represent these schemas. SWCs represent instances of contexts (described in terms of emotional links and tensions between existing characters) in which an action has appeared in a previous story, and they act like rules during the engagement phase: an action is added to the plot if a story-world context for that action can be found that matches the plot so far. It is important to note that SWCs (and not the definitions of action in terms of their preconditions) are used to find the next action to extend the plot. The reflection phase revises the plot so far, mainly checking it for coherence, novelty, and interest. The checks for novelty and interest involve comparing the plot so far with that of previous stories. If the story is too similar to some previous one, or if its measure of interest compares badly to previous stories, the system takes action by setting a guideline to be obeyed during engagement. These guidelines are a low-level equivalent of author goals, driving which types of action can be chosen from the set of possible candidates. The check for coherence is only carried out over the final version of the story, and it involves inserting into the text actions that convey explicitly either character goals or tensions between the characters that are necessary to understand the story. Unless they are explicitly added during this check, goals and tensions are not included in the discourse.

A Mexica story follows:

Jaguar knight was an inhabitant of the Great Tenochtitlan. Princess was an inhabitant of the Great Tenochtitlan. Jaguar knight was walking when Ehecatl (god of the wind) blew and an old tree collapsed injuring badly Jaguar knight. Princess went in search of some medical plants and cured Jaguar knight. As a result Jaguar knight was very grateful to Princess. Jaguar knight rewarded Princess with some cacauatl (cacao beans) and quetzalli (quetzal) feathers.

The Virtual Storyteller (2003)

The line of work initiated by Talespin, based on modeling the behavior of characters, has led to a specific branch of storytellers. Characters are implemented as autonomous intelligent agents that can choose their own actions informed by their internal states (including goals and emotions) and their perception of the environment. Narrative is understood to emerge from the interaction of these characters with one another. This guarantees coherent plots, but, as Dehn pointed out, lack of author goals implies they are not necessarily very interesting ones. However, it has been found very useful in the context of virtual environments, where the introduction of such agents injects a measure of narrative to an interactive setting.

The Virtual Story Teller (Theune et al. 2003) introduces a multiagent approach to story creation where a specific director agent is introduced to look after the plot. Each agent has its own knowledge base (representing what it knows about the world) and rules to govern its behavior. In particular, the director agent has basic knowledge about plot structure (that it must have a beginning, a middle, and a happy end) and exercises control over agents' actions in one of three ways: environmental (introduce new characters and objects), motivational (giving characters specific goals), and proscriptive (disallowing a character's intended action). The director has no prescriptive control (it cannot force characters to perform specific actions).

Theune et al. report that nonstructural rules are contemplated to measure issues such as surprise and "impressiveness."

The Virtual Storyteller includes a specific narrator agent, in charge of translating the system representation of states and events into natural language sentences. The development effort on the narrator seems to have focused on correct generation of pronouns to make the resulting text appear natural.

Riedl's Fabulist (2004)

Fabulist (Riedl 2004) was an architecture for automated story generation and presentation. The Fabulist architecture split the narrative generation process into three tiers: fabula generation, discourse generation, and media representation. The fabula-generation process used a planning approach to narrative generation. The intent-driven partial order causal link (IPOCL) planning algorithm simultaneously reasoned about causality and character intentionality and motivation in order to produce narrative sequences that are causally coherent (in the sense that they drive towards a conclusion) and have elements of character believability. Fabulist first generates a narrative plan that meets the outcome objective, ensuring all character actions and goals are justified by events within the narrative itself.

An example story produced by Fabulist is given in the extract. Inputs provided included a domain model describing the initial state of the story world and possible operations that can be enacted by characters, and an outcome state (Jasmine and Jafar are married, and the genie is dead). The plan for the corresponding story is shown in figure 1.

> There is a woman named Jasmine. There is a king named Jafar. This is a story about how King Jafar becomes married to Jasmine. There is a magic genie. This is also a story about how the genie dies.
>
> There is a magic lamp. There is a dragon. The dragon has the magic lamp. The genie is confined within the magic lamp.
>
> King Jafar is not married. Jasmine is very beautiful. King Jafar sees Jasmine and instantly falls in love with her. King Jafar wants to marry Jasmine. There is a brave knight named Aladdin. Aladdin is loyal to the death to King Jafar. King Jafar orders Aladdin to get the magic lamp for him. Aladdin wants King Jafar to have the magic lamp. Aladdin travels from the castle to the mountains. Aladdin slays the dragon. The dragon is dead. Aladdin takes the magic lamp from the dead body of the dragon. Aladdin travels from the mountains to the castle. Aladdin hands the magic lamp to King Jafar. The genie is in the magic lamp. King Jafar rubs the magic lamp and summons the genie out of it. The genie is not confined within the magic lamp. King Jafar controls the genie with the magic lamp. King Jafar uses the magic lamp to command the genie to make Jasmine love him. The genie wants Jasmine to be in love with King Jafar. The genie casts a spell on Jasmine making her fall in love with King Jafar. Jasmine is madly in love with King Jafar. Jasmine wants to marry King Jafar. The genie has a frightening appearance. The genie appears threatening to Aladdin. Aladdin wants the genie to die. Aladdin slays the genie. King Jafar and Jasmine wed in an extravagant ceremony.
>
> The genie is dead. King Jafar and Jasmine are married. The end.

Riedl's Fabulist has addressed specific creative solutions (Riedl and Young 2006) to ordinary planning problems by including specific features that allow the planner to apply modifications to the given input world in order to meet the required goal.

Montfort's Narrator in nn (2007)

In the nn system for interactive fiction (Montfort 2007) the user controls the main character of a story by introducing simple descriptions of what it should do, and the system responds with descriptions of the outcomes of the character's actions. Within nn, the Narrator module provides storytelling functionality, so that the user can ask to be "told" the story of the interaction so far. The Nar-

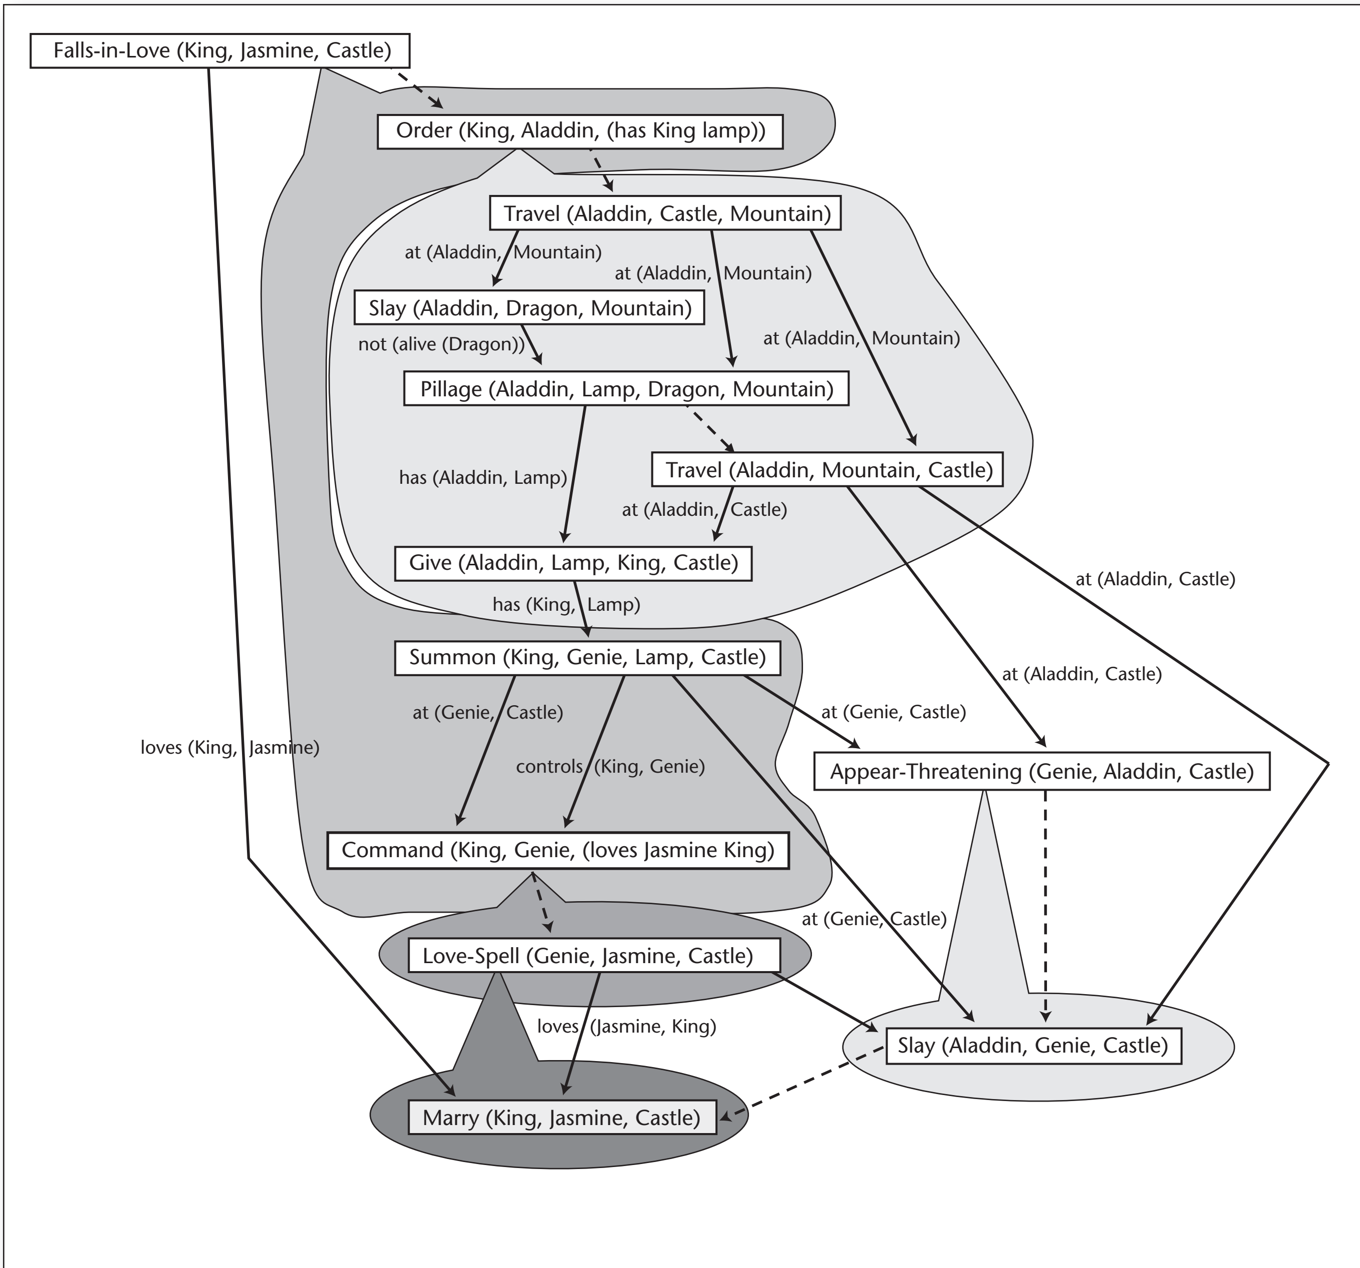

Courtesy, Mark Riedl.

Figure 1. Plan for a Fabulist Story.

rator module of nn addresses important issues in storytelling that had not been addressed by previous systems: order of presentation in narrative and focalization. Instead of telling events always in chronological order, the nn Narrator allows various alternative possibilities: flashbacks, flash-forwards, interleaving of events from two different time periods, telling events back to front. It also captures appropriate treatment of tense depending on the relative ordering of speech time, reference time, and event time. Focalization is handled by the use of different *focalizer worlds* within the system. Aside from the actual world of the interactive fiction system, nn maintains additional separate worlds representing the individual perspectives and beliefs of different characters. These can be used to achieve correct treatment of focalization (telling the story from the point of view of specific characters).

Discussion

The storytelling systems that have been reviewed can now be discussed in terms of the basic criteria set out in the earlier sections of the article.

The Creative Process

With respect to the creative process, I have identified several points: who is the creator, what is the output, who is the audience (especially in terms of the evaluation of novelty of the output), expectations and whether the output is unexpected in some way, whether the output meets some goal, what are the inputs, whether feedback is being contemplated.

The simplest starting point is to establish what it is that should be considered as output of the storytelling processes. Ideally, this should be stories in natural language (what I have defined as discourse). However, most systems concentrate their effort in producing fabula, a conceptual description of what should be told. We could be faced with two different products, involving two different processes (generating a fabula and generating a discourse). In these terms, the different systems reviewed present different characteristics. For instance, the nn system receives the fabula as a combination of the preprogrammed game file and the successive interactions by the player, and its role is to generate the corresponding discourse. On the other hand, several of the systems reviewed focus on generating a world simulation (a succession of world states, as in Novel Writer, Virtual Storyteller; a succession of episodes as in Author; a succession of story contexts as in Mexica; or a succession of actions in the form of a plan as in Talespin, Universe, Minstrel, and Fabulist). These world simulations are then rendered as a discourse by more or less complex methods. As a first approximation we may consider these world simulations as fabula. We can separate the reviewed systems into storytellers (nn) and story inventors (all the rest, though all of these do also to a certain extent tell the stories they invent).

Inventing Stories

In most of the systems reviewed there is very little consideration of the role played by the audience in the storytelling process. Fabulist does model audience perception of the believability of characters. Mexica uses its set of previous stories as a reference when checking output for novelty. In a way, this set of previous stories may be considered as a model of what the audience already knows, with the evaluation of novelty acting as a feedback loop over it. It also suggests that Mexica is programmed to ensure at least P-creativity (in Boden's terms), whereas the other systems have no guarantee of either P- or H-creativity (results after the first run may stop being novel as very similar stories to initial outputs are obtained).

In general terms, there is no modeling of what the active process of consuming a story might be like. Paul Bailey's work (1999) started to address the issue of story generation from the reader's perspective, arguing that existing storytellers were based either on author models, on story models, or on world models, and that a different kind of storyteller, based on a model of the reader, was required. It would be very positive if more storytellers undertook such a research program.

Some systems include a model of authoring (Author, Universe, Minstrel, Mexica) and others rely more on world simulations (Novel Writer, Talespin, Fabulist, Virtual Story Teller). Those that include a model of authoring have a higher potential for providing insights on how humans address the tasks being modeled. For instance, the review unearthed a conflict between Dehn's approach (story worlds built post hoc to justify decisions taken at the discourse level) and Lebowitz's (world should be built first to guide discourse). This seems to be a choice that human writers also face. Mexica's model of authoring is based on an existing account of how humans approach the task of writing (Sharples 1999).

With respect to unexpectedness, these storytellers prefer to play it safe. Although expectations are not modeled explicitly as such in any of these systems, many of them include guiding constraints in their code that ensure certain basic expectations are met by the resulting stories (the overall structure of the weekend party plot in Novel Writer, generic author goals in Author and Minstrel, specification of optimal tension arc in Mexica, basic knowledge about plot structure in Virtual Storyteller). Rather than experimental storytellers, trying to break genre conventions and still obtain something valuable, we clearly have storytellers that try to follow genre conventions very strictly. In terms of Ritchie's analysis, developers seem to have focused much more on obtaining results that rate highly on typicality ("looks somewhat like a typical story") rather than unexpectedness. These expectations can also be interpreted as the set of constraints that define the conceptual space.

An interesting exercise might be to try to correlate this set of expectations with what Ritchie calls the inspiring set for each system: a set of stories that inspired the system developer. It seems clear that these would be weekend party murder mysteries for Novel Writer, La Fontaine's fables for Talespin, King Arthur's tale for Minstrel, TV soap operas for Uni-


Once upon a time...

verse, and the set of previous stories for Mexica. As some of Ritchie's criteria evaluate how creative a system is in terms of whether its output is significantly different from the inspiring set, many of these system are likely to rate badly on those.

There seems to be no in-built procedure for determining the quality (as captured by Ritchie's second rating function) of the stories produced. Though some systems (Mexica) have reported external evaluations of the quality of the stories by human readers, in general it is fair to say that none of these systems meet the first criteria proposed by Jennings for creative autonomy (systems are capable of evaluating their own output).

With respect to the inputs, some systems receive as input an initial configuration of the world and a desired outcome (Fabulist), some receive only an initial configuration of the world (Novel Writer, Talespin, Universe, Virtual Storyteller, though in most of these the expectations described above should be considered as additional inputs), some receive an initial state and a moral to use as a guideline (Minstrel), and some receive just a set of author goals (Author).

Finally, some developers have consciously tried to put in mechanisms that could be interpreted as attempts to achieve what Boden calls transformational creativity. Dehn describes a process of creative reasoning, part deliberate and directed and part serendipitous, capable of identifying unforeseen opportunities during planning and deviating from its original goals to follow them. This would surely meet the requirements for being considered creative from a computational point of view. The actual operation leading to this behavior would have to be considered in more detail to evaluate its potential. Universe does indeed include a method for generalizing plot fragments that could lead to the production of new material, thereby qualifying as creative. Turner centered his development of Minstrel on the specific point of how to "find and use old knowledge in new ways to form a novel and useful solution to a problem." His TRAMs constitute an example of a possible computational solution to that problem. Riedl's extension of Fabulist (Riedl and Young 2006) that allows the planner to modify the given input world in order to meet required goals constitutes an example of how a small transformation of the search space (that goes beyond simply exploring it) may result in more creative outcomes. This particular mechanism introduced a measure of novelty in otherwise fairly deterministic plans. It also fits in with Boden's and Sharples's general hypothesis that transforming the constraints is a valuable operation in creative systems. The original report on Novel Writer, although it never mentioned creativity explicitly, included claims to cognitive functionality that would surely be considered computational creativity of the first order. I refer to the description of how the natural language metacompiling capability would "permit a character to develop new behavior patterns as a function of his experiences during the course of a simulation" (Klein et al. 1973). The report is unclear as to whether this claim had been substantiated or tested in any way. It is features such as these that may carry some of these systems beyond what Wiggins refers to as GOFAI.

Telling Stories

If the processes for inventing stories in the reviewed systems rated low in terms of creativity, the rating obtained by processes for telling stories is even sadder. The challenge of how to tell a story has received very little attention in general, and it is mostly tagged on as a final stage to systems that concentrate on inventing stories. The nn system is a notable exception in that it involves a significant effort to model computationally some of the basic elements contained in Genette's work on narrative discourse (Genette 1980): relative order of presentation and focalization. However, all the systems that tell the stories they invent do in fact include default solutions to many of the technical challenges involved in telling a story. In most cases, the representation of the fabula that is produced already has the form of a linear sequence. This would only be true for the simplest fabulas, with complex examples taking the form of a cloud of events taking place over a continuum of time and space. For all the systems that allow multiple characters (all the ones reviewed do), the process of telling the story involves working out which character to follow as the story is told and how/when to go back to tell what has been happening to other characters in the meantime. Some systems include features to represent the different views that individual characters have of the story.

The grandfather of them all, Novel Writer mentions how this representation of different focalized views of the world could be achieved, indicating awareness of its importance in storytelling. Talespin models character perception and physical space in terms of maps, which is crucial for computing what each character sees of the story. In the nn system, specific representations of this information are stored in the form of focalizer worlds. The Virtual Storyteller allows each agent to hold its very own knowledge base of what is happening in the world. However, most systems pay very little attention to describing these decisions, making it very difficult to consider whether any creativity may be involved.

In general terms, the task of telling stories has been explored very little and still requires much work. Although it has a great potential for creative solutions (as exemplified by human storytellers), I am not aware of any attempt to model computationally creative approaches to the task of telling stories.

Conclusions

The review of existing storytelling systems shows that this line of research has experienced considerable growth over the years. Although it has never been a popular research topic, nonetheless it has received sustained attention over the years by a dedicated community of researchers. In recent years the number of systems developed has increased significantly.

The body of work resulting from these efforts has identified a significant number of relevant issues in storytelling. Successive systems have identified particular elements in stories that play a role in the process of generation. Although a consensus is yet to be found, all of these insights may be applied successfully in future systems. Part of the problem may be that researchers in the field find it difficult to benefit from one another's systems due to the significant differences in story representation chosen for each of them.

Inventing a story and telling a story are different processes involving different inputs, different outputs, and different guiding criteria. Although they often interact heavily, some effort should be made to study them separately to identify the subtasks and intermediate representations that each one requires. This should make it easier to identify the way in which they interact when they do.

The task of modeling the role of the audience, not just as a set of constraints taken into account by the storyteller, but as an active process in itself, has been given little consideration in existing systems. This oversight should be corrected in the future, as modern literary theory is progressively attributing more importance to the role of the reader in evaluating literary texts.

With respect to whether they address recent trends in computational creativity, most systems are concerned with telling stories that are recognized as typical of the particular genre. There is very little emphasis on achieving novel stories or innovation of any kind in the way particular story elements are treated. Because the issue of what should be valued in a story is unclear, research implementations tend to sidestep it, generally omitting systematic evaluation in favor of the presentation of hand-picked star examples of system output as means of system validation. An effort should be made to adopt evaluation practices introducing some measurement of novelty and quality. Of all the systems reviewed only one (Mexica) included internal mechanisms to address the need for novelty. Concepts recently identified as important in the attribution of creativity to a system, such as existence in a social context of critics and creators, the ability to evaluate its own outcomes, or the ability to change the way it evaluates, are poorly represented among the reviewed systems. However, existing storytelling systems were mostly developed before the models for computational creativity appeared in recent years. Under the circumstances, it is not surprising that they perform poorly.

Future work in the field should address these issues as means of advancing toward more creative systems.

Mexica is also a pioneer in modeling an operation model of how humans reason specifically during writing, rather than adopting an existing computational model such as forward chaining of rules, planning, or case-based reasoning. This is considered a positive characteristic, and more complex models of specific human cognitive abilities may be used in the future to enrich the field.

The extent to which a system is considered creative should evaluate the ability to produce a significant number of different stories, and these stories should be significantly different from one another. To achieve this, systems should consider operating on a number of significantly different inputs (as human creators do). The selection of appropriate inputs to a creative process might be a more important ingredient in the perception of creativity in the output than has been allowed so far. After all, in many cases of striking human creativity, the significant innovation has lain in the adoption of unexpected inputs for well-known processes (as when designing new useful products by recycling old ones).

Acknowledgements

This research is partially supported by the Ministerio de Educación y Ciencia (TIN2006-14433-C02-01), Universidad Complutense de Madrid and Dirección General de Universidades e Investigación de la Comunidad de Madrid (CCG07-UCM/TIC 2803). The final version of this article has been considerably improved thanks to anonymous comments received during the review process.

Notes

1. Some have only appeared as interesting fields of study in recent times.

2. Focalization is also described as *point of view*, or *perspective*, but these terms were considered ambiguous. The term *focalization* was introduced by Genette (1980), and it has been preferred since.

References

Bailey, P. 1999. Searching for Storiness: Story-Generation from a Reader's Perspective. In *Narrative Intelligence: Papers from the AAAI Fall Symposium*, ed. M. Mateas and P.

Sengers. Technical Report FS-99-01. Menlo Park, CA: AAAI Press.

Boden, M. 1990. *The Creative Mind*. London: Weidenfeld and Nicholson.

Dehn, N. 1981. Story Generation after Tale-Spin. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence,* 16–18. Los Altos, CA: William Kaufmann, Inc.

Forster, E. M. 1927. *Aspects of the Novel*. New York: Harcourt.

Genette, G. 1980. *Narrative Discourse: An Essay in Method*. Ithaca, NY: Cornell University Press.

Jennings, K. E. 2008. Developing Creativity: Artificial Barriers in Artificial Intelligence. Paper presented at the Fifth International Joint Workshop on Computational Creativity, Ciudad Universitaria, Facultad de Informatica, Madrid, Spain, 17–19 September.

Klein, S.; Aeschliman, J. F.; Balsiger, D.; Converse, S. L.; Court, C.; Foster, M.; Lao, R.; Oakley, J. D.; and Smith, J. 1973. Automatic Novel Writing: A Status Report. Technical Report 186, Computer Science Department, The University of Wisconsin, Madison, Wisconsin.

Kolodner, J. L. 1980. Retrieval and Organizational Strategies in Conceptual Memory: A Computer Model. Ph.D. dissertation, Yale University, New Haven, CT.

Labov, W. 1972. *Language in the Inner City: Studies in the Black English Vernacular.* Philadelphia: University of Pennsylvania Press.

Lebowitz, M. 1985. Story-Telling as Planning and Learning. *Poetics* 14: 483–502.

Lebowitz, M. 1983. Story-Telling as Planning and Learning. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence,* volume 1. Los Altos, CA: William Kaufmann, Inc.

Lehnert, W. G. 1981. Plot Units and Narrative Summarization. *Cognitive Science* 5(4): 293–332.

Meehan, J. R. 1977. Tale-Spin, an Interactive Program That Writes Stories. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence,* 91–98. Los Altos, CA: William Kaufmann, Inc.

Montfort, N. 2007. Generating Narrative Variation in Interactive Fiction. Ph.D. dissertation, University of Pennsylvania, Philadelphia, PA.

Pérez y Pérez, R. 1999. MEXICA: A Computer Model of Creativity in Writing. Ph.D. dissertation, The University of Sussex, Falmer, UK.

Riedl, M. 2004. Narrative Planning: Balancing Plot and Character. Ph.D. dissertation, Department of Computer Science, North Carolina State University, Raleigh, NC.

Riedl, M., and Young, R. 2006. Story Planning as Exploratory Creativity: Techniques for Expanding the Narrative Search Space. New Generation Computing, Computational Paradigms and Computational Intelligence. Special Issue: *Computational Creativity* 24(3): 303–323.

Ritchie, G. 2007. Some Empirical Criteria for Attributing Creativity to a Computer Program. *Minds and Machines* 17(1): 67– 99.

Schank, R. C. 1982. *Dynamic Memory: A Theory of Reminding and Learning in Computers and People*. Cambridge, UK: Cambridge University.

Sharples, M. 1999. *How We Write: Writing as Creative Design*. London: Routledge.

Theune, M.; Faas, E.; Nijholt, A.; and Heylen, D. 2003. The Virtual Storyteller: Story Creation by Intelligent Agents. In *Proceedings of the Technologies for Interactive Digital Storytelling and Entertainment (TIDSE) Conference,* 204–215. Berlin: Springer.

Turner, S. R. 1993. *Minstrel: A Computer Model of Creativity and Storytelling*. Ph.D. dissertation, University of California at Los Angeles, Los Angeles, CA.

Weiner, R. 2000. *Creativity and Beyond: Cultures, Values, and Change*. Albany, NY: State University of New York Press.

Wiggins, G. 2006. Searching for Computational Creativity. New Generation Computing, Computational Paradigms and Computational Intelligence. Special Issue: *Computational Creativity* 24(3): 209–222.

Pablo Gervás works as associate professor (profesor titular de universidad) at the Departamento de Ingeniería del Software e Inteligencia Artificial, Facultad de Informática, Universidad Complutense de Madrid. He is the director of the NIL research group (nil.fdi.ucm.es) and also of the Instituto de Tecnología del Conocimiento (www.itc.ucm.es). In recent years, Gervás has taken part in the organization of several scientific meetings on topics related to creativity: AISB'03 Symposium on Artificial Intelligence and Creativity in Arts and Science; Workshop on Language Resources for Linguistic Creativity, LREC 2004; Computational Creativity Workshop, ECCBR 2004; Computational Creativity Workshop, IJCAI 2005; and Fourth International Joint Workshop on Computational Creativity, IJWCC 2008. He was founding member of the Computational Creativity Working Group (WG4) of initiative COST 282: "Knowledge Discovery in Science and Technology," funded by the European Commission. His research is on computational creativity, processing natural language input, generating natural language output, building resources for related tasks, and generating stories. In the area of creative text generation, he has done work on automatically generating metaphors, formal poetry, fairy tales, and short films.