Reading and Writing as a Creative Cycle: the Need for a Computational Model

Pablo Gervás
Instituto de Tecnología del Conocimiento
Universidad Complutense de Madrid
28040 Madrid, Spain
pgervas@sip.ucm.es

Carlos León
Facultad de Informática
Universidad Complutense de Madrid
28040 Madrid, Spain
cleon@fdi.ucm.es

Abstract

The field of computational narratology has produced many efforts aimed at generating narrative by computational means. In recent times, a number of such efforts have considered the task of modelling how a reader might consume the story. Whereas all these approaches are clearly different aspects of the task of generating narrative, so far the efforts to model them have occurred as separate and disjoint initiatives. There is an enormous potential for improvement if a way was found to combine results from these initiatives with one another. The present position paper provides a breakdown of the activity of creating stories into five stages that are conceptually different from a computational point of view and represent important aspects of the overall process as observed either in humans or in existing systems. These stages include a feedback loop that builds interpretations of an ongoing composition and provides feedback based on these to inform the composition process. This model provides a theoretical framework that can be employed first to understand how the various aspects of the task of generating narrative relate to one another, second to identify which of these aspects are being addressed by the different existing research efforts, and finally to point the way towards possible integrations of these aspects within progressively more complex systems.

Introduction

The field of computational narratology has been steadily growing over the recent years. There have been many effort aimed at analysing narrative in computational terms (Mani 2012), and generating narrative by computational means (Gervás 2009). With respect to computational creativity, the latter is more immediately relevant. Though it is possible to argue for a strong role for creativity in the understanding of narrative, this is less obvious than the role of creativity in the generation of narrative. This kind of argument has lead over the years to many research efforts that focus on generation of narrative to the detriment of the understanding of it. This is also supported by an argument of a different kind related to the perceived difficulty of narrative understanding from computational terms, and the lack of success of the efforts accumulated on that topic over the years. Yet it is also very clear to any seasoned reader or writer that the task of generating narrative is intrinsically bound to that of reading it. A writer writes to be read, and a writer aiming to succeed writes with the reactions of possible readers in mind. This point was originally argued in the field of narratology by authors such as Barthes (Barthes, Miller, and Howard 1975) and Ecco (Eco 1984), and in the field of automated storytelling by Paul Bailey (Bailey 1997) but it has taken a long time for the research community to act upon it. In recent times, a number of research efforts arising from an initial focus on narrative generation have started to consider the task of modelling how a reader might consume the story based on the plausible inferences that arise from a narrative discourse. From a technical perspective, these approaches are based on techniques used to obtain a plausible inference of causal and intentional relations in the discourse (Niehaus 2009; Cardona-Rivera et al. 2012; O'Neil 2013). These efforts arise from the need of generation processes to have access to some kind of feedback based on how the results of the construction process will be perceived by a potential reader. The pragmatic needs of research seem to require the implementation of at least some parts of this cycle between writing and reading that are intuitively evident to most people.

The present paper provides a breakdown of the activity of creating stories into five stages that are conceptually different from a computational point of view and represent important aspects of the overall process as observed either in humans or in existing systems. A fundamental hypothesis of the proposed breakdown is that, even though intended as a model of the composing task, it includes two additional processes concerned with modelling the task of interpretation. These processes are aimed at estimating the impression that a composition will make on an asumed interpreter, and they provide a feedback loop to improve the results of composition. This extension provides the means both for including a model of the reader in the composition process, and for explicitly representing evaluation features as part of the construction process. The proposed breakdown into five stages is analysed in terms of its relation to existing models of: creative endeavour from a computational point of view, the writing task from a cognitive perspective, and natural language generation as a set of tasks. The set of five stages is postulated as a possible model to understand how existing efforts in the field of story generation relate to one another and how future progress in the field might explore possible

interactions between them. To this end, a number of existing systems are reviewed in the light of the model.

Previous Work

The set of existing theoretical models or frameworks that may have a bearing on the task of story creation are reviewed in the following order. First, models of creative systems, then models of the writing task, and finally models of natural language generation.

Computational Models of Creativity

Wiggins (Wiggins 2006) takes up Boden's idea of creativity as search over conceptual spaces (Boden 2003) and presents a more detailed theoretical framework intended to allow detailed comparison, and hence better understanding, of systems which exhibit behaviour which would be called *creative* in humans. This framework describes an exploratory creative system in terms of a tuple of elements which include elements for defining a conceptual space as a distinct subset of the universe of possible objects, the rules that define a particular subset of that universe as a conceptual space, the rules for traversing that conceptual space, and an evaluation function for attributing value to particular points of the conceptual space reached in this manner.

The IDEA model (Colton, Charnley, and Pease 2011) assumes an (I)terative (D)evelopment-(E)xecution-(A)ppreciation cycle within which software is engineered and its behaviour is exposed to an audience. An important insight of this model is that the invention of measures of value is a fundamental part of the creative act. In the case of story generation this corresponds to developing models of reader response that can be used to provide feedback to the generation process.

Cognitive Accounts of Writing and Narrative Comprehension

Flower and Hayes (Flower and Hayes 1981) define a cognitive model of writing in terms of three basic process: planning, translating these ideas into text, and reviewing the result with a view to improving it. These three processes are said to operate interactively, guided by a monitor that activates one or the other as needed. The planning process involves generating ideas, but also setting goals that can later be taken into account by all the other processes. The translating process involves putting ideas into words, and implies dealing with the restrictions and resources presented by the language to be employed. The reviewing process involves evaluating the text produced so far and revising it in accordance to the result of the evaluation. Flower and Hayes' model is oriented towards models of communicative composition (such writing essays or functional texts), and it has little to say about narrative in particular. Nevertheless, a computational model of narrative would be better if it can be understood in terms compatible with this cognitive model.

Sharples (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. He 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. Apparently the human mind is incapable of addressing simultaneously these two tasks. 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.

Narrative comprehension involves progressive enrichment of the mental representation of a text beyond its surface form by adding information obtained via inference, until a situation model (representation of the fragment of the world that the story is about) is constructed (van Dijk and Kintsch 1983). A very relevant reference in this field is the work of (Trabasso, vand den Broek, and Suh 1989), who postulate comprehension as the construction of a causal network by the provision by the user of causal relations between the different events of a story. This network representation determines the overall unity and coherence of the story.

Natural Language Generation

The general process of text generation takes place in several stages, during which the conceptual input is progressively refined by adding information that will shape the final text (Reiter and Dale 2000). During the initial stages the concepts and messages that will appear in the final content are decided (*content determination*) and these messages are organised into a specific order and structure (*discourse planning*), and particular ways of describing each concept where it appears in the discourse plan are selected (*referring expression generation*). This results in a version of the discourse plan where the contents, the structure of the discourse, and the level of detail of each concept are already fixed. Although the overall process includes a number of additional stages (*aggregation*, *lexicalization* and *syntactic choice* - collectively referred to as *sentence planning* -, and *surface realization*) these will not be relevant for the purpose of the present paper, which remains focused at the level of discourse.

The ICTIVS model

At its most abstract level, the task of composing a narrative must be considered in the broader context of an act of communication (see Figure 1). The communication takes place as an exchange of a linear sequence of text that encodes a large and complex set of data that correspond to a set of events that take place over a volume of space time, possibly in simultaneous manner at more than one location. To convey this complexity as a linear sequence and recover it again at the other end of the communication process requires a process of condensing it first into a message and then expanding it again into a representation as close as possible to the original. There is a *composer*, in charge of composing a linear discourse from a conceptual source that may also

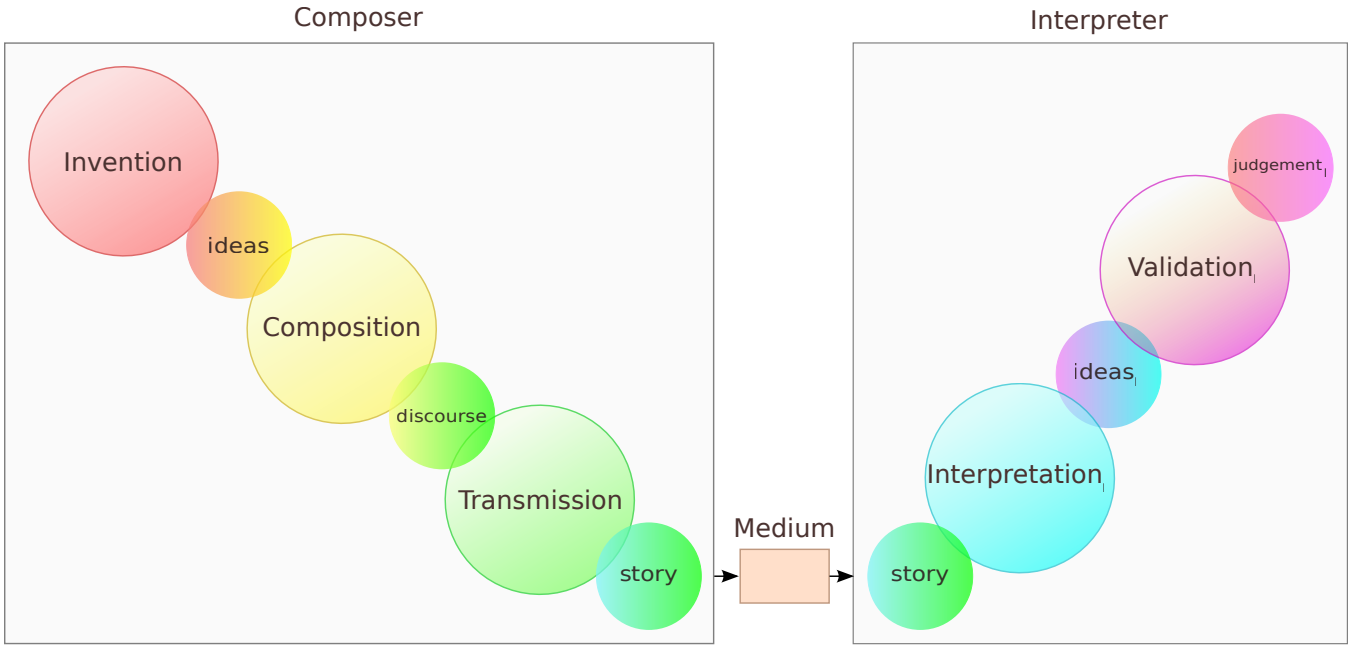

Figure 1: The traditional view of the communication process. Each big circle corresponds to an operation by one of the actors involved, whereas each small circle corresponds to the type of information conveyed from one to another. Note that $\texttt{ideas}_I$ recovered by the interpreter need not correspond faithfully to the `ideas` originally conceived by the composer.

have been produced by himself, and an *interpreter*, faced with the task of reconstructing a selected subset of the material in the conceptual source as an interpretation of the received narrative discourse. [1] The task of the composer involves four facets: the construction of the source material for the message as a conceptual representation, the selection of what subset of the conceptual source to convey, the linearization of that selection as a discourse, and the encoding of the message in a particular medium. The task of the interpreter involves a number of tasks concerned with the process of interpretation of the story into a conceptual representation, and validation of the corresponding content with respect to the criteria of the interpreter. The main hypothesis defended in this paper is that the composer also has the responsibility of ensuring that the discourse she produces is optimized to help the interpreter construct exactly the interpretation she desires to convey. To this end, the composer may need to resort to local models of the processes applied by the interpreter, used to produce copies of the conceptual interpretation and the validation that an interpreter might obtain by applying them. In consequence, the models of the interpretation process considered in this paper are not strictly concerned with the tasks carried out by the interpreter, but rather with how the outcomes of these tasks might best be modelled relying as much as possible on the resources and capabilities already available to the composer.

Based on these ideas, an abstract model for covering these aspects of narrative has been created. It has been called ICTIVS (the name stands for INVENTION, COMPOSITION, TRANSMISSION, INTERPRETATION and VALIDATION of Stories). This model divides the communicative act of narration into five stages carried out by the composer as part of an iterative cycle. Figure 2 depicts this cycle as a refinement of the traditional view of the task of the composer, now extended with an explicit representation of the task of the interpreter. This model of the interpreter provides a feedback loop on the composition process that can be used for progressive refinement of the result. The ICTIVS model does not try to solve or study *how* each process is carried out from a social or psychological point of view, it rather identifies those stages that are important from the Artificial Intelligence point of view, and those that help to model the human behaviour in narratives.

- During the INVENTION stage, the narrative content is created, based on incomplete knowledge or from scratch. Characters, narrative objectives, places and events (the *ideas*) all emerge and get related, thus creating a complex set of facts that constitute the source for the story. These facts could be understood as the log of a simulation run on the set of characters. As in real life, events produced in this way may have happened simultaneously in physically separated locations, and constitute more of a cloud than a linear sequence, a volume characterised by 4 dimensional space-time coordinates.
- The COMPOSITION stage arranges all data from the previous stage (INVENTION) and outputs a *discourse*. Composing a discourse for the source content involves drawing a number of linear pathways through the volume of space

[1] In real life, the role of the composer is usually played by a writer and the role of interpreter by a reader, but in the present case a more generic formulation has been preferred for generality.

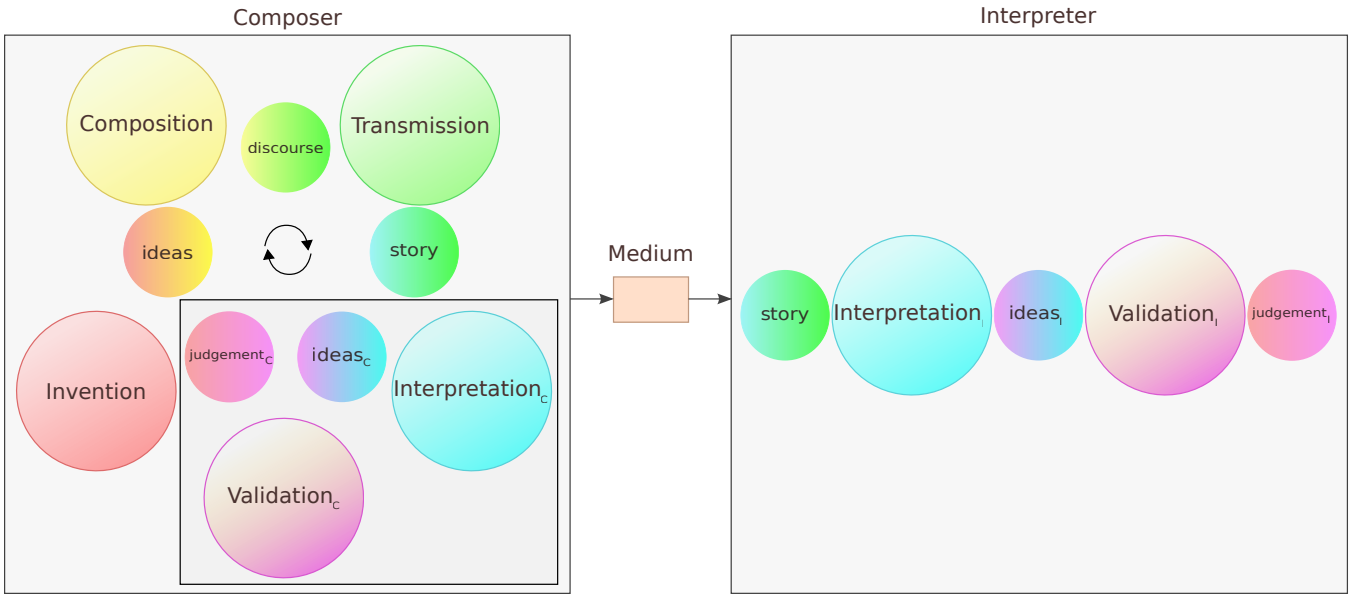

Figure 2: The ICTIVS model. It constitutes a model of the composing task. The picture includes a separate representation of the interpreter to capture two important ideas: that the proposed refinement is intended as a duplication of the interpretation task within the composer, and that the ideas ($\texttt{ideas}_C$) and the judgement ($\texttt{judgement}_C$) obtained by the composer may be different from those developed by the interpreter ($\texttt{ideas}_I$ and $\texttt{judgement}_I$, as a result of the fact that the procedures applied to obtain them are different ($\texttt{Interpretation}_C \neq \texttt{Interpretation}_I$ and $\texttt{Validation}_C \neq \texttt{Validation}_I$).

time produced by the invention stage. This type of linear pathway is sometimes referred to as a *narrative thread*. All the narrative threads deemed relevant from a given input (in truth a selection of all available ones or even a selection of fragments of the interesting parts of some of them) need to be combined together into a single linear discourse. As a result, this discourse is an ordered and filtered set of facts (properties, events, descriptions...) that are to be conveyed to the interpreter. Filtering involves considering the reader's common knowledge and inferential capabilities. Many concepts that the composer intends to convey may be omitted from the actual discourse if they can be considered to be known or obtainable via inference by the reader. It is also possible that the composer prefer to withhold particular items of information over particular stretches of the discourse, to create or enhance effects such as surprise, expectation, or suspense.

- Once a discourse has been composed, it can be rendered in a particular medium that can be consumed directly by the intended audience (whether a single interpreter or many). This stage has been called TRANSMISSION, as it involves the task of rendering the discourse in a given medium and making the medium available to the audience, but the part of the process we want to consider here is that of rendering, which involves constructive decisions and may be informed by reflection.
- The INTERPRETATION stage involves the reconstruction of the content of the message from the discourse for it. This process, when applied to a story received from an external source, constitutes the main task that an interpreter faces. Our stance in this paper is that an integral part of the task of the composer could be to apply a similar procedure to a recently composed discourse, with a view to obtaining feedback on how a hypothetical interpreter might view it. Whether from the discourse itself or from the medium produced to render it, the composer attempts to reconstruct the meaning as a user would to extract feedback on how the result of his composition task satisfies his communication goals.
- Over the reconstruction of the content of a story interpreted from a discourse, interpreters (and composers simulating the reaction of an interpreter) develop judgments on the medium, the discourse or the content of the story. This set of operations we refer to as the VALIDATION stage. As with interpretation, we consider that a composer may rely on a version of this stage to obtain feedback on how his output might be received by an interpreter.

The role of the INTERPRETATION stage is crucial even if the model is nominally restricted to the task of composition. According to the Flower and Hayes model of the writing task, linearization would occur as part of the translation subtask (converting ideas into text), followed by a number of cycles of reviewing and improving the result. The accumulated literature on modelling story generation indicates that this reviewing stage of discourse, based on an attempt at reconstructing the desired content from the discourse and a comparison between the resulting interpretation and the selected subset of the source material, is a fundamental ingredient of the broader context of the task of story generation. We therefore consider that a model of the task of story gen-

eration should include all of the five stages described to be considered complete.

One may be tempted to ascribe creativity within this model only to the INVENTION stage, on the grounds that it is there that new content is put together by combining more basic elements. However, there is also room for creativity in the COMPOSITION stage - to come up with new solutions for encoding a given content, possibly fulfilling additional goals in terms of surprise, suspense, while still meeting the communicative constraints - or the TRANSMISSION stage - to produce alternative novel and valuable renderings for a given discourse. During the INTERPRETATION stage a new instantiation of the narrative message is created. In some cases, the process of COMPOSITION reduces the content so drastically that the INTERPRETATION process requires some creative mechanisms to come up with enough material to make sense of the story. In those cases new *ideas* not considered by the writer may emerge during this stage. The resulting story is not necessarily equal to the story that the writer invented and transmitted. This point aligns very well with the observations of postmodern literary studies - arising from the work of (Barthes 1977) - along the lines that a text does not acquire its ultimate value until it has been interpreted by a particular reader, and that the role of the reader in this process must be valued in comparable terms to that of the writer. The VALIDATION process is particularly interesting in terms of creativity. In line with the insights arising from the IDEA model of Colton et al, a fundamental part of the creative act may be the invention of new measures of value. This would correspond to applying creativity at the VALIDATION stage, and it is a feature that has received little attention in the past in terms of computational creativity research. Finally, it is quite possible that creativity as perceived by external observers arise only as a result of a complex interaction between all these processes. This possibility strengthens the argument in favour of models of the composition task that captures all these aspects in a single framework.

The ICTIVS Model and Existing Related Frameworks

The ICTIVS model is compared to a number of existing frameworks for understanding related processes, of creativity, of the writing task, and of natural language generation.

ICTIVS and Models of Creativity

Processes in the INVENTION and COMPOSITION stages would correspond to what Wiggins in his framework defines as rules for traversing the conceptual space. These stages carry out the identification of new artifacts in the conceptual space of stories of the working domain. On the other hand, both the INTERPRETATION and the VALIDATION stages can be seen as ingredients in an evaluation function function in Wiggins' formalization. They both compose a process in which a story is received and judgments are formed. The TRANSMISSION stage is not explicitly addressed by Wiggins, as his model only considers the generation of creative artifacts.

Although Colton et al's IDEA model is formulated in the context of the development of creative software, its description of the process as an (I)terative (D)evelopment-(E)xecution-(A)ppreciation cycle is applicable to the task of generating a story. Under this view, INVENTION would correspond to Development, COMPOSITION and TRANSMISSION would correspond to Execution, and INTERPRETATION and VALIDATION would correspond to Appreciation.

ICTIVS and Cognitive Models of Writing

From a cognitive point of view, the set of stages that constitute the ICTIVS model aligns reasonably well with the processes described by Flower and Hayes. In terms of Flower and Hayes' model, the INVENTION stage would constitute specific operation of the planning process. The COMPOSITION stage might be considered partly within the planning process (as regards discourse planning decisions) and partly within the translating process (as regards sentence planning processes). The TRANSMISSION stage would fall directly within the translating process, including the particular "restrictions and resources presented by the language to be employed", as Flower and Hayes phrase it. The INTERPRETATION and VALIDATION stages would correspond to the reviewing process of Flower and Hayes' model. The possibility of considering different paths through the various stages of the model would correspond to enriching the model with interaction between the various processes as controlled by a monitor, which is an integral part of Flower and Hayes model.

In terms of Sharples' description of the writing task, it would be simple to say that INVENTION and COMPOSITION would correspond to the engagement phase, and that INTERPRETATION and VALIDATION would correspond to the reflection phase. However, Sharples' analysis indicates that the process of writing is far from being a simple cycle over such stages, and involves coming and going between them over a period of time, before the actual stage of TRANSMISSION is ever contemplated. In fact, it would probably be fair to say that there might be specific phases of engagement associated with INVENTION, combined with phases of reflection over whatever representation is achieved at that stage, followed by iterations of INVENTION and COMPOSITION engagements (with interspersed phases of reflection as INTERPRETATION and VALIDATION of the resulting discourse), followed by iterations of INVENTION, COMPOSITION and TRANSMISSION engagement (also combined with phases of reflection as above). Such a complex process would match the idea of heavy interaction between planning, translating and reviewing (in Flowers and Hayes terms), and should be considered corroboration of the need for a monitor module to govern how these interactions take place. This monitor would also be in charge of deciding when the final product is finally ready to be transmitted to the addressee, or generally made public.

The processes of progressive enrichment of the mental representation of a text beyond its surface form by adding information obtained via inference, as described by Van Dijk and Kintsch (van Dijk and Kintsch 1983) is the main component of the INTERPRETATION stage. This does indeed take

place when a reader attempts to comprehend a given text. However, the ICTIVS model considers this stage also to be a fundamental part of the process of creation applied by the writer. Much in the way described by Colton et al in their IDEA model, the process of creating a story is seen as an interactive cycle of production of a text (through processes of INVENTION, COMPOSITION and TRANSMISSION) followed by a process of appreciation (during INTERPRETATION and VALIDATION). The result of this appreciation process can then be fed back to the next iteration of the productive part of the cycle. Although the cycle is described in full, going all the way to the production of text before entering an appreciation phase, it is perfectly possible (and extremely plausible if considered in terms of how this task is addressed by humans) that appreciation in this sense may be applied much earlier in the cycle: for instance, once a process of INVENTION has taken place, whatever has been obtained, possibly a set of ideas represented conceptually, or a sketch of the fabula - in narratological terms - may be appreciated and the resulting information can be fed back to further processes of INVENTION. As INVENTION does not include a step of selection and encoding of information (these tasks concern the COMPOSITION stage) no stage of INTERPRETATION is required as part of this cycle, and feedback may be obtained by direct VALIDATION. A similar internal loop may occur involving COMPOSITION, with appreciation of the output of a COMPOSITION stage being submitted to appreciation even before entering a stage of TRANSMISSION. In this case, a process of INTERPRETATION may be required before VALIDATION can be applied.

Given that (Trabasso, vand den Broek, and Suh 1989) postulate the existence of a network of causal relations between the different events of a story as fundamental to determining the overall perception of its unity and coherence, it is very likely that VALIDATION of a story involve identification of an appropriate network of this nature. When VALIDATION is applied directly to the result of an INVENTION stage (fabula), it may consist simply of ensuring that such causal relations are present in the story. When applied to a narrative discourse, an intermediate stage of INTERPRETATION may be required to elicit a representation of such a network from the discourse.

ICTIVS and Natural Language Generation

At a first glance, with respect to the classic pipeline structure for natural language generation systems, the ICTIVS stage of INVENTION would correspond to the task of content determination, whereby a fabula is produced (content that may be told), with the discourse planning stage matching the COMPOSITION stage. However, there is a slight misalignment between the two models. The content determination stage of a NLG pipeline assumes all possible content to be present, and applies a selection process to establish what will be included in the communication under consideration. In contrast, the INVENTION stage is concerned with actual production of the content to be considered. In view of these, both content determination and discourse planning - as understood in NLG terms - can be considered as part of the COMPOSITION stage. In truth, all of the NLG pipeline could be considered as part of the COMPOSITION stage, with possibly only surface realization being included in the TRANSMISSION stage.

Grounding the ICTIVS Model in Existing Story Generation Systems

The applicability of the proposed model can be illustrated by using it to analyse existing efforts in story generation, with a view to recasting their apparent diversity into a homogeneous framework of understanding, and to better illustrate how they relate to the more complex aspects of narrative generation and to one another. A number of existing systems are discussed below. The selection is not meant to be exhaustive, and it has been designed to include examples of systems that cover different stages of the ICTIVS model.

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 MEXICAS (also wrongly known as Aztecs). *MEXICA* pioneered in the realm of automated storytellers the idea of a cycle of generation and evaluation, with the results of the evaluation being fed back to inform the generation process. In this case, the engagement cycle of *MEXICA* can be seen as a particular type of INVENTION process that directly produces a linear discourse. Over this discourse, the *MEXICA* system applies an instance of the VALIDATION stage, which is fed back into the generation process. In addition to this, *MEXICA* had a procedure for building from a set of known stories the knowledge structures called Story Contexts, which represented explicitly the emotional links and tensions between characters in the story. This process would correspond to an ICTIVS stage of INTERPRETATION. Finally, *MEXICA* provide a template-based procedure for rendering the final discourses as text. This would correspond to a stage of TRANSMISSION. There is very little in the operation of the system that might be considered an instance of COMPOSITION.

For ease of exposition, the reviewed systems are grouped into sets based on the stage that they devote most attention to.

Mostly Inventors

The *Virtual Storyteller* (Theune et al. 2003) introduces a multi-agent approach to story creation where stories are created by cooperating intelligent agents. 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. There is a specific director agent who has basic knowledge about plot structure and exercises control over agent's actions by: introducing new characters and objects, giving characters specific goals, or disallowing a character's intended action. There is also a specific narrator agent, in charge of translating the system representation of states and events into natural language sentences. In terms of the ICTIVS model, most of the operation of the *Virtual Storyteller* would correspond to a stage of

INVENTION, with very simple stages of COMPOSITION and TRANSMISSION encapsulated in the narrator agent.

Fabulist (Riedl and Young 2010) was an architecture for automated story generation and presentation. The *Fabulist* architecture split the narrative generation process into threetiers: fabula generation, discourse generation, and media representation. The fabula generation process used a planning approach to narrative generation and it would correspond to an ICTIVS stage of INVENTION. The discourse generation would correspond to an ICTIVS stage of COMPOSITION. The media representation would correspond to an ICTIVS stage of TRANSMISSION.

Inventors-Composers

MINSTREL (Turner 1992) was a computer program that told stories about King Arthur and his Knights of the Round Table. The program was started on a moral that was used as seed to build the story. Story construction in *MINSTREL* operates as a two-stage processes involving a planning stage and a problem-solving stage. At a high level of abstraction, the two processes described for MINSTREL seem to correspond to an amalgamation of the INVENTION and COMPOSITION stages.

BRUTUS (Bringsjord and Ferrucci 1999) was a program that wrote short stories about betrayal. The operation of *BRUTUS* involves three basic processes, carried out sequentially. First a thematic-frame is instantiated. Then a simulation-process is set in motion where characters attempt to achieve a set of pre-defined goals, thereby developing a plot. The process of converting the resulting plot into the final output is carried out by the application of a hierarchy of grammars (story grammars, paragraph grammars, sentence grammars) that define how the story is constructed as a sequence of paragraphs which are themselves sequences of sentences. Of these, the instantiation of the thematic frame and the simulation-process would correspond to an ICTIVS stage of INVENTION, the application of the hierarchy of grammars would blend together stages of COMPOSITION and TRANSMISSION.

Mostly Composers

There have been a number of systems developed that address the task of generating a discourse for a given set of events (León, Hassan, and Gervás 2007; Gervás 2012; Gervás 2013). These systems received as input a broad description of the set of events to consider and produce from it a conceptual representation of the discourse needed to tell them as a story. The main contributions of these systems correspond to implementations of an ICTIVS stage of COMPOSITION. Most of them include an additional stage of TRANSMISSION that renders the resulting discourses as text. In most cases these are intended for ease of evaluation, and little effort is invested in optimising the quality of the resulting texts.

In the *nn* system for interactive fiction (Montfort 2007) (now evolved into the *Curveship* system (Montfort 2009)) 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 Narrator module of *nn* was a pioneer among storytellers in that it addressed issues such: order of presentation in narrative and focalization, chronology, and appropriate treatment of tense depending on the relative ordering of speech time, reference time, and event time. In this case, the Narrator module of *nn* combines a very refined instance of a COMPOSITION stage, that deals with the issue of variation in the narrative form, and a much simpler instance of a TRANSMISSION module, which renders the resulting discourse as text.

Mostly Transmitters

STORYBOOK (Callaway and Lester 2002) produced multipage stories in the Little Red Riding Hood domain by relying on elaborate natural language generation tasks. Callaway's system is a realtime narrative prose generator that takes an instance of the presentational ordering desired for the text and an instance of the sum of the factual content that constitutes the story as input, and intelligently combines information found in the two and stylistic directives to produce narrative prose. In this sense, *STORYBOOK* can be said to be centred on the TRANSMISSION stage of the ICTIVS model. The process of devising the presentational ordering desired for the text from the sum of the factual content that constitutes the story would correspond to the COMPOSITION stage of the ICTIVS model. The task of developing the sum of the factual content that constitutes the story - not actually addressed by *STORYBOOK* - would correspond to the INVENTION stage of the ICTIVS model.

Inventors - Validators

Stella (León and Gervás 2011; León and Gervás 2012) performs story generation by traversing a conceptual space of partial world states based on narrative aspects. World states are generated as the result of non-deterministic interaction between characters and their environment. This generation is narrative agnostic, and an additional level built on top of the world evolution chooses the most promising ones in terms of their narrative features. *Stella* makes use of objective curves representing these features and selects world states whose characteristics match the ones represented by these curves. *Stella* is an example of INVENTION based on VALIDATION of internal states.

Composers-Interpreters

A significant example is the *INFER* system (Niehaus 2009), a narrative discourse generation system that employs an explicit computational model of a readers comprehension process during reading to select content from an event log with a view to creating discourses that satisfy comprehension criteria.

Mostly Interpreters

An example is *INDEXTER* (Cardona-Rivera et al. 2012), a cognitive framework which predicts the salience of previously experienced events in memory based on the current

event that the audience exposed to a narrative is experiencing. This system constitutes a model of the experience of the reader, and it involves a process of INTERPRETATION in the sense that it aims to model the online mental state of the audience which experiences the narrative. This requires progressive monitoring of the effect of each increment in the narrative on this model.

A Shortage of Validators

The VALIDATION stage of the ICTIVS model has not seen as many implementations over the years. There has been a significant research effort on the evaluation of results from story generators of various types but these consisted mostly on evaluations carried out by humans over results produced by generation systems. These efforts include: evaluating the effects of text choices on reader satisfaction (Callaway and Lester 2001), evaluating plots in terms of their acceptability and their novelty as perceived by users (Peinado and Gervás 2006), and development of specific frameworks for evaluating aspects of automatically generated narrative (Rowe et al. 2009).

Some existing systems (Pérez y Pérez 1999; Cheong 2007; Bae and Young 2008; Niehaus 2009; León and Gervás 2010) did include a specific module for validating their output as it is constructed. Of these, different systems focused on specific aspects, such as emotional tensions (Pérez y Pérez 1999), suspense (Cheong 2007), surprise (Bae and Young 2008), comprehensibility (Niehaus 2009) or conformance with a user given specification of the evolution over the story of particular parameters (León and Gervás 2012).

All these systems involve some type of cycle of construction of a candidate story (sometimes a partial draft rather than a complete one) and applying some function to validate this before continuing.

It is only in recent times that systems devoted specifically to validating properties of a narrative have been developed, such as the *DRAMATIS* model for evaluating suspense in narratives (O'Neil 2013), which includes a significant stage of interpretation to make validation possible.

Conclusions

The arguments presented in this paper suggest that the inclusion of explicit processes of interpretation and validation to inform and complement the task of constructing narratives is plausible in terms of existing models of the task in terms of human cognition. They also show how existing efforts at modelling various aspects of the story telling task have already addressed computational modelling of the various aspects that would be required to implement such inclusion. The proposed solution would achieve the integration within the computational model of the narrative construction of both a model of the reader and specific procedures for the evaluation of candidates results. This would address longstanding requirements on the storytelling task (Bailey 1997) and more recently voiced requirements on the improvement of scientific rigour in the evaluation of creative systems (Jordanous 2011).

However, it must be said that the ICTIVS model is not intended as a cognitively plausible model of the way humans deal with narratives. Instead, it is proposed as a conceptual framework that might help to understand the diversity of existing efforts in story generation, and how they relate to the more complex aspects of narrative generation and to one another. In this sense, the ICTIVS model is put forward as a rallying call for researchers in the fields of narrative modelling, story generation and computational creativity to start advancing along the difficult road of integrating together existing views and development efforts. The ICTIVS model may contribute to this task in two different ways. First, by naming and clarifying some of the subprocesses involved, it may allow future research efforts to focus on the less well explored aspects of the described cycle, which should help to enrich our overall understanding of the phenomenon. Second, by providing a simple framework for analysing existing systems in terms of a set of common elementary operations, it can help identify parts of existing systems that it might be useful to reuse in future developments or to combine with other existing ones. To this end, a conscious effort has been made to formulate the ICTIVS model at a purely conceptual level. To ensure compatibility with the broad variety of representations employed in existing systems, no detail is given of what specific representations might be considered for the data exchanged between different phases.

Progress along the lines of defining formal interfaces between the various stages is desirable in the long run, but it would require a thorough and detailed review of existing efforts in search of a consensus on possible representations for the various stages. The WHIM project, funded by the European Commission under call FP7-ICT-2013-10 with grant agreement number 611560, is a three year project that sets out to explore technologies for ideation, with a particular focus on the role that narrative generation might play in evaluating the quality of ideas. Among its objectives, it includes an effort to provide a workable specification of narrative oriented towards generation. It is envisaged that this effort will contribute to clarifying some of the details that have been glossed over in the present paper.

The effort invested so far in developing computational solutions aimed at achieving or improving computational generation of narrative has uncovered a number of different aspects to the basic phenomenon of telling a story. Whereas all these approaches are clearly different aspects of the task of generating narrative, so far the efforts to modelled them have occurred as separate and disjoint initiatives. There is an enormous potential for improvement if a way was found to combine results from these initiatives with one another. The model presented in this paper provides a theoretical framework that can be employed first to understand how these various aspects of the task of generating narrative relate to one another, second to identify which of these aspects are being addressed by the different frameworks, and finally to point the way towards possible integrations of these aspects within progressively more complex systems. Systems obtained in this way are more likely to be perceived as models of the human ability to generate stories.

A set of important insights arise from the application of the model to a selection of existing systems:

1. there are several distinct computational processes involved in the generation of a story: invention of the material to be used, composition of the material as a valuable linear discourse, transmission of this discourse using some medium
2. each one of these processes contributes some features to the final story that may be evaluated separately: on the material to be used one may evaluate coherence or originality, on the discourse issues such as comprehensibility, surprise, suspense, on the final medium grammaticality or fluency
3. some of the features arise only as an interaction between the processes and some require an intermediate process of interpretation to bring out to the fore this interaction between the underlying material and the discourse used to convey it

As a result, efforts at computational modelling must take into account the various processes, the interaction between them, and the need for a validation stage as an integral part of the process.

From the point of view of creativity, it is important to note that most existing efforts at story generation have focused on obtaining acceptable stories, with very little attention to the perceived creativity of the process. Even in cases such as (Turner 1992; Pérez y Pérez 1999) that declare an explicit interest in creativity, the actual implementation and evaluation process does not address issues that are considered fundamental in the emerging field of computational creativity, like novelty or sustained creativity. This is largely due to the inherent technical difficulties in achieving results that can be considered as acceptable stories, let alone creative ones. The creativity in story generation may arise from any of the processes involved and further creativity may arise from the interactions between them. Taking the argument above to the extreme, for story generators with an aspiration to being considered truly creative systems the validation stage must include specific solutions for measuring creativity related features beyond those that are elementary requirements of the story form.

Finally, two important ideas arise from the interaction between the proposed model and considerations on creativity. The first one is that creativity may be involved in many of the processes involved in this model, not just in that of inventing the content of a story. Composition and interpretation of stories may involve significant amounts of creativity. The creation of innovative procedures for evaluation or validation of stories may be considered a highly creative achievement. The second one is that a perception of creativity in a storytelling system may arise from the interaction between all these processes rather than be located in a particular one. This constitutes a strong argument in favour of attempting the implementation and study of models of story telling along the lines of the proposed model.

Acknowledgments

This paper has been partially supported by the projects WHIM 611560 and PROSECCO 600653 funded by the European Commission, Framework Program 7, the ICT theme, and the Future and Emerging Technologies FET program.

References

Bae, B.-C., and Young, R. M. 2008. A use of flashback and foreshadowing for surprise arousal in narrative using a plan-based approach. In *Proc. ICIDS 2008*.

Bailey, P. 1997. *A Reader-based Model of Story Generation: Ph.D. Thesis Proposal*. DAI discussion paper. Edinburgh University.

Barthes, R.; Miller, R.; and Howard, R. 1975. *S/Z: An Essay*. Farrar, Straus and Giroux.

Barthes, R. 1977. The death of the author. In *Image, Music, Text*. New York, NY, USA: Hill and Wang.

Boden, M. 2003. *Creative Mind: Myths and Mechanisms*. New York, NY, 10001: Routledge.

Bringsjord, S., and Ferrucci, D. 1999. *Artificial Intelligence and Literary Creativity: Inside the mind of Brutus, a Story-Telling Machine*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Callaway, C. B., and Lester, J. C. 2001. Evaluating the effects of natural language generation on reader satisfaction. In *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 164–169.

Callaway, C. B., and Lester, J. C. 2002. Narrative prose generation. *Artificial Intelligence* 139(2):213–252.

Cardona-Rivera, R.; Cassell, B.; Ware, S.; and Young, R. 2012. Indexter: A computational model of the event-indexing situation model for characterizing narratives. In *Proc. Workshop on Computational Models of Narrative 2012*.

Cheong, Y. 2007. *A Computational Model of Narrative Generationfor Suspense*. Ph.D. Dissertation, North Carolina State University, Rayleigh, North Carolina, USA.

Colton, S.; Charnley, J.; and Pease, A. 2011. Computational creativity theory: The face and idea descriptive models. In *In 2nd International Conference on Computational Creativity*.

Eco, U. 1984. *The Role of the Reader: Explorations in the Semiotics of Texts*. A Midland book. Indiana University Press.

Flower, L., and Hayes, J. 1981. A cognitive process theory of writing. *College Composition and Communication* 32(4):365–387.

Gervás, P. 2009. Computational approaches to storytelling and creativity. *AI Magazine* 30(3):49–63.

Gervás, P. 2012. From the fleece of fact to narrative yarns: a computational model of narrative composition. In *Proc. Workshop on Computational Models of Narrative 2012*.

Gervás, P. 2013. Stories from games: Content and focalization selection in narrative composition. In *First Spanish Symposium on Digital Entertainment, SEED 2013*.

Jordanous, A. 2011. Evaluating evaluation: Assessing progress in computational creativity research. *Proceedings of the Second International Conference on Computational Creativity (ICCC 2011)* 102–107.

León, C., and Gervás, P. 2010. The Role of Evaluation-Driven rejection in the Successful Exploration of a Conceptual Space of Stories. *Minds and Machines* 20(4):615–634.

León, C., and Gervás, P. 2011. A top-down design methodology based on causality and chronology for developing assisted story generation systems. In *Proceedings of the 8th ACM conference on Creativity and cognition*, C&C '11, 363–364. New York, NY, USA: ACM.

León, C., and Gervás, P. 2012. Prototyping the use of plot curves to guide story generation. In *Third Workshop on Computational Models of Narrative, 2012 Language Resources and Evaluation Conference (LREC'2012)*.

León, C.; Hassan, S.; and Gervás, P. 2007. From the event log of a social simulation to narrative discourse: Content planning in story generation. In Olivier, P., and Kray, C., eds., *Conference of the Artificial and Ambient Intelligence*, 402409.

Mani, I. 2012. *Computational Modeling of Narrative*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.

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

Montfort, N. 2009. Curveship: An interactive fiction system for interactive narrating. In *Proceedings of the Workshop on Computational Approaches to Linguistic Creativity*, CALC '09, 55–62. Stroudsburg, PA, USA: Association for Computational Linguistics.

Niehaus, J. 2009. *Cognitive Models of Discourse Comprehension for Narrative Generation*. Ph.D. Dissertation, North Carolina State University, Rayleigh, North Carolina, USA.

O'Neil, B. 2013. *A Computational Model of Suspense for the Augmentation of Intelligent Story Generation*. Ph.D. Dissertation, Georgia Institute of Technology, Atlanta, Georgia, USA.

Peinado, F., and Gervás, P. 2006. Evaluation of automatic generation of basic stories. *New Generation Computing, Computational Paradigms and Computational Intelligence. Special issue: Computational Creativity* 24(3):289–302.

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

Reiter, E., and Dale, R. 2000. *Building Natural Language Generation Systems*. Cambridge University Press.

Riedl, M., and Young, M. 2010. Narrative planning: Balancing plot and character. *J. Artif. Intell. Res. (JAIR)* 39:217–268.

Rowe, J. P.; McQuiggan, S. W.; Robison, J. L.; Marcey, D. R.; and Lester, J. C. 2009. Storyeval: An empirical evaluation framework for narrative generation. In *AAAI Spring Symposium*.

Sharples, M. 1999. *How We Write*. 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.

Trabasso, T.; vand den Broek, P.; and Suh, S. 1989. Logical necessity and transitivity of causal relations in stories. *Discourse Processes* 12:1–25.

Turner, S. 1992. *MINSTREL: A Computer Model of Creativity and Storytelling*. Ph.D. Dissertation, University of California at Los Angeles, Los Angeles, CA, USA.

van Dijk, T. A., and Kintsch, W. 1983. *Strategies of Discourse Comprehension*. New York: Academic Press.

Wiggins, G. 2006. A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* 19(7).