Article

Institute of Materials, Minerals & Mining

Interdisciplinary Science Reviews
2024, Vol. 49(2) 269–290

DOI: 10.1177/03080188241262862
journals.sagepub.com/home/ids
Sage

Storytelling systems: Progress towards generation informed by models of the reader

Pablo Gervás[1]

Abstract

The present paper attempts to identify specific points of contact between efforts to generate narrative in the field of artificial intelligence and the field of narratology. Such points of contact centre around the tasks of modelling processes related to the construction of narrative in computational terms. A number of existing research efforts to develop systems related to storytelling are reviewed. An analysis of the functionalities that they cover suggests that storytelling might be profitably studied in terms of at least four different sub-tasks: Fabula generation, narrative discourse composition, text generation, and narrative interpretation. An analysis from a computational point of view of the context in which narrative is used as a means of communication between author and reader is presented. This shows that parts of the cognitive processes involved are yet to be modelled, which suggests that the set of sub-tasks identified may continue to grow with further research.

Keywords

Storytelling systems, narratology, fabula generation, discourse composition, narrative interpretation

Introduction

At this point in time, artificial intelligence has been a scientific discipline for over 60 years. Over that time there has been opportunity for many attempts to explore its applicability to the generation of narrative. This involved artificial intelligence researchers deciding to address challenges in the world of narratology. To achieve that goal, these researchers were forced to extend significantly their knowledge of these fields. In more recent times, the advent of digital humanities (Terras, Nyhan and Vanhoutte 2016) has opened the gates to a flow of information in the other direction: Researchers in humanities having to expand their knowledge of computational methods and artificial intelligence solutions for handling text.

[1]Facultad de Informática, Universidad Complutense de Madrid, Madrid, Spain

Corresponding author:
Pablo Gervás, Facultad de Informática, Universidad Complutense de Madrid, Madrid 28040, Spain.
Email: pgervas@ucm.es

The fact that these flows of information exist across disciplines must be considered as a positive step for all the fields involved.

The present paper attempts to identify specific points of contact between efforts to generate narrative in the field of artificial intelligence and the field of narratology. Such points of contact centre around the tasks of modelling processes related to the construction of narrative in computational terms.

As artificial intelligence has sought to gain a purchase on the cognitive processes involved in generating narrative, it has increasingly turned to psychology as a source for understanding what it is that people do when confronted with a narrative. As various theories have been considered as sources for computational models, this search has opened a window onto how much of the impact of a narrative on its readers arises not from the text itself but rather from cognitive processes that the reader applies to the material she receives, and from the outcomes of such processes. This ties in with certain notions in postmodern narratology that suggest that a significant part of the value of a narrative is created by the reader in the process of reading it (Barthes 1967).

A significant part of the difficulties that have faced these efforts is related to the confusion prevailing as to how to describe the objects of study. Different disciplines use different terms to describe stories, consider them at different level of detail, and again use different terms to describe the various elements that make them up. This problem is addressed in Section 'The problem of mismatched terminology' of the paper.

Although it has never been a main stream topic in artificial intelligence, efforts at developing story generation have recurred over time. A number of relevant systems working along these lines are reviewed in Section 'Computational models for generating narrative'. The development of these story generation systems – as argued in the review – led to a growing familiarity with narratological concepts among AI researchers working on these problems. As a result, it slowly became clear that for computational models of narrative to succeed they needed to model not only the author but also the reader. Along these lines, many systems have been developed aiming to model the dynamics of the reader's reaction to the content of a narrative text. Existing systems that involve this type of reader model have been reviewed in Section 'Computational models of processing of narrative'. Section 'Reflections on computational models of cognitive tasks related to literature' presents some reflections on what the computational models of narrative generation built to date say about the nature of narrative, and Section 'Conclusions' gathers some relevant conclusions and outlines what further models might be explored in the future.

The problem of mismatched terminology

Stories are naturally occurring phenomena across cultures and languages. Every human being has an internalized idea of what a story is, and a myriad of terms to refer to it. In the context of each particular language, there is a different such set of terms. Whenever researchers have turned to study stories, they have made an effort to develop more specific terminology to help them understand what they were dealing with. The problem is that these efforts at defining terminology often have been local to specific disciplines. This section reviews some relations between the terminologies developed around stories in the fields of narratology and artificial intelligence.

Narratology

Narratologists have over time attempted to clarify the set of concepts relevant to understand the phenomena of stories. Unfortunately for the layman, these efforts have given rise to many distinctions that are phrased in terms with generic meanings in specific languages – *fabula/sjuzhet* from Russian

(Viktorovič 1965), *histoire*/*discours* from French (Todorov 1980), *story*/*plot* from English (Forster 1927) – but which are chosen to represent very specific narratological concepts that are not necessarily derived univocally from their meaning in those languages. For people outside the field, it is extremely difficult to tell whether these are all different concepts, the same concepts by different names, or whether the different names capture subtle nuances between the concepts. If the latter is the case, a detailed analysis of how the different nuances relate to storytelling systems is beyond the scope of this paper. Yet maybe a working relation between concepts in both fields can be established to help understand the points of contact.

A helpful midpoint might be in an initial distinction between a sequence of events and a discourse that orders and presents events, as suggested by Culler (2004). Even then the actual assignment of specific terms to each of these two concepts varies depending on the particular theorist. The sequence of events told more or less aligns with what is referred to as *fabula* or *histoire*, the discourse that presents the events in a particular order similarly aligns with what is referred to as *sjuzhet* or *discours*. To cover the level of detail used in existing storytelling systems, one may need to go a little further and consider the three levels proposed by Bal and van Boheemen (2009): Fabula, story and text. The events of a narrative in the chronological order they happened is referred to as the *fabula*. The events of a narrative in the chronological order in which they are told is referred to as the *story*. Bal refers to the actual rendering of the *story* as the *text*.

In this setting, there is yet a further distinction that is relevant to the analysis presented in this paper: The concept of *plot* as distinct from those already outlined. The distinction is subtle, but hinges on the fact that a plot goes beyond a simple sequence of events by virtue of the events in it being connected together by relations of causality (Forster 1927).

A final concept in narratology relevant to our endeavour is that of a *storyworld* (Ryan 2019). According to Ryan, a storyworld is an image of the world described in (the text for) a given story. She explicitly notes that a storyworld is different from the image of the world being represented in the story, because the world being represented may exist independently of the text, and if such is the case, it may differ from the one that is actually described in the story. All these distinctions will be useful later to analyse how similar concepts are handled in storytelling systems.

All of these concepts are relevant to the analysis of story generation systems. For a start, because it is very plausible that when a system developer decides to call his system a 'story' generation system, it would ill-advised to assume that the term 'story' in that context is a direct reference to one of these carefully defined terms proposed by narratologists. In that context, 'story' is used in the same way that lay people use it, as a generic reference to some type of narrative. Much the same can be said about 'storytelling': Even though this term features 'story', it refers to the act of telling, and therefore it includes development not just of 'story' in the narratological sense, but also a rendering of the story – the 'telling' – in some form. That form will possibly be text, but it may also be visual or audiovisual. This is important because many of the storytelling systems in existence follow this route.

More refined discussion of the relations will be addressed in the rest of the paper.

Finally, a set of additional narratological concepts are relevant to the discussion on systems that consider focalization, as covered in Section 'Focalization and perspective'. *Perspective* refers to the way in which the representation of a story is influenced by the choice of narrator, whether in terms of their personality, values or the relative access they have to events in the storyworld (Niederhoff 2009). It is generally considered equivalent to point of view, but some narratologists distinguish a specific concept of focalization. *Focalization* (Genette, Lewin and Culler 1983) corresponds to the way in which the content of a narrative, at a particular point of the telling, is constrained to what might be perceived by a particular agent, known as the *focalizer*. When the narrative is focalized on a particular character, the reader only gets access to the events in the story world that the character is aware of. This trait is often exploited by authors of narrative texts to allow the reader to access only certain parts of the story at particular times. These concepts are fortunately used in their correct sense in the publications that report the corresponding systems, so now issue of alignment arises.

Artificial intelligence

The growth of artificial intelligence over the last decades has been exponential. In that process, many disciplines have arisen to deal with specific aspects of human intelligence. Unfortunately, each such discipline has often proceeded to define specific terminology to better describe its subject of study, and these efforts at definition have not always been carefully informed by prior definitions in related disciplines.

Two particular such disciplines propose definitions that need to be considered in this context: Natural language generation (NLG) and simulation-based storytelling.

Natural language generation. NLG is concerned with generating text from conceptual descriptions of content to be conveyed. The body of knowledge that emerged with it considers the task of generating a text to involve four distinct stages (Reiter and Dale 2000): *content determination* – which establishes what material should be conveyed –, *discourse planning* – which establishes how the selected material should be organized –, *sentence planning* – which decides on the form that the sentences used to convey the material should take – and *surface realization* – which determines how the sentence plans should be rendered as text according to the grammatical and typographical rules of a particular language. Each of these stages receives a specific type of representation and produces an enriched or transformed version of it. These intermediate representations constitute a set of concepts relevant to the understanding of these systems: an initial representation of the material available, from which content determination selects the *content* – a representation of the material to be conveyed – which discourse planning transforms into a *discourse* – a conceptual description of the order in which the material is to be conveyed – which sentence planning represents as a set of *sentence plans* – conceptual descriptions of the actual sentences to be used to convey the material – which surface realization converts into *text* – an actual rendering of the material as text in a particular language.

The relation between these concepts and the relevant concepts from narratology detailed above is discussed in section 'A proposal for cross-disciplinary basic terms'.

Agent-based systems. Agent-based systems (Shoham and Leyton-Brown 2008) define a set of agents and a set of rules governing their behaviour, assign specific goals to the agents, and run the resulting system to simulate an evolving community. Such systems have from early on been used to provide input material for story generation systems (Theune et al. 2003). This is done by building agents, rules of behaviour, and goals that are likely to lead to simulations that include interesting successions of events. It is important to note that such systems do not necessarily produce sequences of events, as many agents may be operating simultaneously over a large space that cannot be encompassed as a single scene. The simulation that results from such a setup is considered a *storyworld* (Ryan 2018). Efforts to produce a story from such a storyworld involve processes of selecting a subset of the events produced by the simulation to be included in the story, a process of establishing an appropriate order to tell them, and processes for building a text rendering of the resulting ordered sequence. The process of identifying valuable stories to tell from storyworlds is referred to as *storysifting* (Ryan 2018).

When considered in combination with the set of concepts from NLG, the storyworld produced by an agent-based simulation operates as the full set of material available to the process of content determination, and subsequent stages can be aligned easily with discourse planning, sentence planning and surface realization.

A proposal for cross-disciplinary basic terms

The analyses of the terminologies for these disciplines present no explicit matches, but potential alignments emerge.

Both narratology and agent-based simulation present a 'storyworld' term. However, the concepts referred to in each case differ. A storyworld in the sense proposed by Ryan is an image of the world described in a story, and she distinguighes it explicitly from the image of the world from which the events in the story are extracted, which exists independently of the text. The concept of storyworld considered as input for a story sifting process is precisely a representation of the world that precedes the story, which inspires the story, and from which the material included in the story has to be selected. Such a storyworld includes significant amounts of material that will not appear in the story. It is the input to a content determination stage – as defined in NLG – that will carry out the selection.

The concept of storyworld as proposed for story sifting is also distinct from the concept of fabula, in as much as it includes events that will not appear in the story. I propose a specific concept of *source storyworld*, which consists of all the events that happened in the world where the story occurred, regardless of whether they are included in the narrative or not. When one studies narratives as finished products, these events not included in the story are not known and possibly irrelevant, but they become important when one attempts to model the process by which narratives are constructed from a set of events observed in real life – or observed in a computer simulation. In such cases, the establishment of appropriate criteria for selection of which events to tell and which events to omit becomes a fundamental challenge for the system developer.

The narratological concept of *fabula/histoire* – the events of a narrative in the chronological order in which they happened - aligns reasonably well with the NLG concept of *content* – a representation of the material to be conveyed. There is however a small difference related to the issue of its chronological order. A source storyworld does not necessarily satisfy a strict chronological order. Events may occur simultaneously, have different durations or overlap in surprising ways. As a result, it is very unlikely that a fabula defined in this sense take the form of a sequence. I will therefore consider for the rest of the paper a definition of fabula as the subset of the source storyworld that is included in the story. There may be some kind of partial order among the events included in that subset, but they are not likely to be strictly ordered and they cannot be considered a sequence. The task of organizing the selected events into a linear sequence involves a transformation from fabula onto an instance of the narratological concept of *sjuzhet/discours/story*.

The narratological concept of *sjuzhet/discours/story* – the events of a narrative in the chronological order in which they are told – aligns reasonably well with the NLG concept of *discourse* – a conceptual description of the order in which the material is to be conveyed. As 'story' is consistently used in the artificial intelligence literature to refer to the broader concept as understood by lay people, I have chosen to refer to the*sjuzhet/discours/story* concept in this paper as *discourse*. This reduces the risk of confusion. It also aligns well with the distinction proposed by Ricoeur (1976) between the level of *discourse* – a sequence of sentences, where each sentence involves a predicate applied to some entities that need to be identified by the subject (and objects) of the sentence – and the level of language – the particular structure of the particular linguistic system – at which *text* occurs.

At this stage, the alignment options run short, in the sense that I find it difficult to identify narratological concepts equivalent to the NLG concept of a sentence plan. However, it is important to note that a representation of a story in terms of sentence plans differs from a discourse in that it includes decisions that are relevant to the rendering of the discourse as text in a specific language – such as choices of referring expressions, lexical elements, and syntactic constructions. A sentence plan is in this sense already phrased at the level of language – in Ricoeur's sense – even though it is not yet text. In this sense, it may be considered as a conceptual description of the text, or an intermediate stage between discourse and text.

The concepts of text in both disciplines can be considered to match.

For the rest of the paper, the following terminological convention has been followed: for all research efforts reviewed, the specific terms used in the original publication are reproduced verbatim, and wherever necessary the relation of such terms to the terms agreed above is made explicit to avoid confusion.

Computational models for generating narrative

Early efforts at automated story generation focused on the qualities of a story as a product. Different researchers identified different aspects of a story that their systems would attempt to model. As progress was made in the field, some of these aspects were adopted as important by subsequent researchers, so that later systems did include them, whereas in many cases aspects covered by an earlier system were dropped by later systems in favour of other aspects considered more relevant. The review sections of this paper are structured in terms of different aspects relevant to stories that have been considered in story generation systems.

Story grammars

One early attempt to model the qualities of stories as a product was the development of *story grammars*. Pioneered by Rumelhart (1975) as a possible representation of the abstract structure of stories, story grammars focus on capturing the surface form of a story in terms of rules that allow parsing of the sequence of events in the story into a hierarchy of conceptual groupings of contiguous non-terminal symbols that represent concepts important to a given analysis of the story. Story grammars are useful to capture observable constraints on instances of narrative discourse, but their applicability is limited when one considers either fabula – which is not restricted to a particular order of presentation – or text – which requires an additional step of parsing the text onto some conceptual representation of the set of events narrated in it. Story grammars were taken up by a number of researchers both to understand particular types of story (Colby 1973) and for generating stories computationally (Pemberton 1989; Lang 1999).[1]

An interesting research effort is that of the BRUTUS system (Bringsjord and Ferrucci 1999), which explored a multi-layered solution generally based on the concept of story grammars. BRUTUS had different grammars operating at different levels of representation, with each one successively operating on a representation produced from the prior grammar until the level of text output was reached. The level of detail in the book does not allow for definite judgments – the set of grammars was never published – but that very lack of detail and the absence of examples beyond the single story that the book reports suggests that we are seeing a hierarchy of articulated construction templates that reverse engineer the structure of that story at different levels.

In relation to the terminology established in section 'A proposal for cross-disciplinary basic terms', systems that rely on story grammars to generate stories are actually generating discourse directly, with no explicit consideration of an associated fabula. When a story grammar allows for a flashback, it places it at the point of the story in which it is mentioned, not when it occurs. Nevertheless, readers manage to reconstruct the fabula in their minds as they read: the flashback is interpreted correctly as having happened earlier. The BRUTUS system combines story grammars at the level of discourse with grammars at the level of text to allow it to produce surprisingly rich texts.

Basic causality

The TALESPIN system (Meehan 1977) told simple stories about woodland creatures. It relied on a simple model of the mental state of the creatures – and what goals might be triggered by specific states, such as 'being hungry' – and a simple model of causality – such as 'eating will end hunger'. The system started from an initial situation for a particular character, established goals for them, and found a sequence of actions that would achieve the character's goals. The resulting sequence of actions was rendered as sentences in a text, and it constitutes the first instance of a computer-generated story.

The TALESPIN system was the first application of the artificial intelligence planning solutions to story generation. The solutions developed in artificial intelligence for automated planning (Nau, Ghallab and Traverso 2004) lend themselves very well to generate a sequence of actions connected by causality. These solutions are based on the concept of *planning operator*, which is an action that has *preconditions* – facts which need to hold so the action can be carried out – and *postconditions* – additional facts that will hold once the action is carried out. A *planner* is given a set of facts that describe an initial situation and a statement that describes a goal, and it will construct automatically a set of chained planning operators that connects the initial situation with the goal. A sequence of actions corresponding to those planning operators constitutes the output *plan*. Because the initial situation is usually made up of a number of facts, the output plan is very rarely a linear chain of actions. Instead, a set of causal chains is built that connect facts in the initial situation to the goal. These causal chains will normally converge on facts that lead to the goal, and eventually converge on the goal. The resulting plan is often in the form of a tree – it is actually a connected causal graph, but the difference is academic for the present argument – with the facts of the initial situation as leaves and the goal at the root. An additional step of establishing a linear order in which to present the actions of the plan as a sequence is required if the plan is to be told as a story.

The connected causal graph produced by the planner is an instance of fabula, in the sense that it contains the events in the story but these events are not in sequence, but rather appearing in causal chains that are branches of the graph. These branches eventually converge to the story outcome – the goal of the plan – but several events may be happening at the same time on different branches. This fabula requires a specific computational process to be transformed into the discourse sequence that is presented as the story. Many planning-based systems do not actually explain how this process takes place, but some do. In TALESPIN this step was included in the process of creating the sequence of sentences for the text. In contrast, others describe it explicitly. For instance, the Fabulist system Riedl (2004) presents an explicit architecture that includes: a story planner, a discourse planner and a media planner. The story planner produces a fabula, the discourse planner produces a discourse from the fabula, and the media planner produces a rendition of the discourse, which may be either as text or as a cinematic visualization.

Systems that generate stories based on planning are more elaborate than those that use story grammars: they include a representation of the fabula, the discourse and the text. Causal relations between the events in the story are captured explicitly in the intermediate representations employed by the system. In this sense, they include what might be considered a representation of plot in Forster's sense. If the process of constructing the discourse is elaborate enough, such causal relations may be explicitly mentioned in the resulting text for the story.

Theme and author goals

Early researchers in artificial intelligence acknowledged that a story was much more than a plan. The Minstrel system (Turner 1993), for instance, enriched the concept of a story as a sequence of actions connected by causality with at least two important additional ideas. First, it built its stories around small statements that acted as morals for the stories being built. The story was built by progressively transforming parts of the moral into related events that matched the draft, until a coherent story was developed that had the moral as a backbone. Second, it included explicit representation of a set of author-level goals, that is, goals that the author would be trying to satisfy at the same time as building the story. The general author-level goals considered by Minstrel were grouped around four different aspects: theme, drama, consistency and presentation. Although the Minstrel system was also generally based on planning, the search procedures used during automated planning were informed by these author-level goals.

The Minstrel system operates by progressively transforming a moral into a discourse for the story. It can therefore be considered to operate mainly at the level of discourse. However, the fact that it includes a representation in terms of causal graphs suggests that it represents explicitly both fabula and plot. In addition, it included a rich set of modules design to refine the discourse aiming for satisfaction of significant aesthetic goals. Relevant modules addressed issues such as suspense, tragedy, characterization and foreshadowing.

Character believability

The planning procedures in traditional planning solutions chain up the actions in the story based on whether they connect the initial situation with the desired goal, but they will not consider whether the set of actions attributed to a particular character constitute a coherent sequence for that character, or whether these actions match the goals of the state of mind that readers will attribute to that particular character based on the evolution of the story to that point. This runs the risk of compromising the believability of characters as perceived by readers of the story. This problem was addressed by Riedl (2004) in his Fabulist system. In Fabulist a specific adaptation of the planning solution guides the construction of a coherent chaining of events, but the adaptation is based on a model that represents explicitly goals for particular characters. This allows it to create stories in which characters acquire a set of goals, and these goals motivate the actions that these characters undertake as part of the story.

The Fabulist system describes among its motivations the need to develop systems capable of dynamically generating a narrative in support of an interaction with a user (Riedl and Bulitko 2013). In such system, user and system take turns at introducing actions into the narrative, and the final story emerges as an ordered compilation of the actions proposed by user and system. This type of interactive narrative system has been the focus of much research into computational modelling of narrative.

The architecture of Fabulist was already discussed on narratological terms in section 'Basic causality'. An additional point of interest is that, by representing explicitly the state of mind of characters and allowing it to inform the story generation process, Fabulist provides an initial modelling of possible worlds in the sense of Ryan (2019).

Conflict

Another aspect considered fundamental for a story that hopes to be interesting is the presence of conflict in some form, usually between characters, often between one character and the world around them. If one considers a classic adventure novel like 'Treasure Island' (Stevenson 1883), the plot[2] follows two opposing sets of characters – the treasure hunters and the pirates – as each set of characters attempts to acquire the treasure. Each plan proposed is thwarted by a counter-plan of the opposing set of characters. As a result, the novel presents an alternation of segments that each cover only the first few actions of a different plan.

This aspect of the narrative was considered the driving force of the narrative generation system developed by Ware (2014) in his Glaive system. Glaive extends the planning paradigm with a representation that considers these possible plans built by the characters – including the actions in them that never get executed – as a record of character intentions at some point in the story. The Glaive planner represents the full search space for the planning problem, including non-executed plans as possible worlds alongside those where the plans were executed. By comparing between elements in this set of possible worlds, the system can build stories featuring conflicts where actions from some plans thwart other plans. The Glaive system is tested experimentally in an interactive narrative set up, where a player tries to achieve a goal and the system coordinates the set of non-player characters to thwart the player's plans.

The explicit representation within the Glaive system of plans built by characters even if they are never executed constitutes an even richer modelling of possible worlds in the sense of Ryan (2019).

Focalization and perspective

More recent efforts have begun to consider – still in a tentative fashion – some elements arising from the additional complexity implied by larger texts, such as changes in perspective.

The system by Gervás (2014) develops a model of how stories may be created by sifting an exhaustive description of a world, identifying interesting subsets of the full set of events that happened in the world, and finding compelling ways of telling about those events as linear discourses. The model is instantiated by considering a chess game as an exhaustive description of a world, and mining the algebraic notation for the game for interesting stories involving particular pieces – rather than telling the story of the whole game or the whole set of pieces. The model allows different parts of the discourse being built to be focalized on particular characters – pieces, in the chess instance. This implies that the reader only sees the part of the game restricted to the field of vision of that piece at that point, which provides a model of focalization. The process of composing the narrative discourse needs to take into account how different choices of focalizer will affect which parts of the game are told and which are omitted. The system does include specific modules for sentence planning and surface realization. Relevant sentence planning operations include contextualising the description of events so that they are described in terms relative to the point of view of the perceiver – as in 'some time later a white knight appears from behind the pawn' – rather than use the absolute mode in which the fabula is described – in terms of positions of the board and number of turns elapsed – such as 'after two turns a knight arrives at position E4'.

A related system develops an interactive version of *The Merchant of Venice* (1598) (Shakespeare 2000) in which the user can experience (a version of) the story lived through the eyes of different characters (Porteous, Cavazza and Charles 2010). In this system, the story itself is generated according to the point of view of the character in terms of how they consider themselves. For the particular story chosen, a telling focalized on the character of Shylock may decide that Shylock is considered as a victim or as a ruthless money lender. This is done by means of a planning solution that consider different versions of the planning operators depending on the chosen point of view.

In terms of our terminology, in the system by Gervás, the chess game is the source storyworld, the set of events selected for the story is the fabula, the narrative discourse constructed by the system is a discourse. However, the system only determines the fabula by selection rather than construction. In the Merchant of Venice, in contrast, not just a different discourse but also a different fabula is constructed for each point of view. This is because the fabula is built using a different set of planning operators for each specific point of view. Gervás' system includes sentence planning and surface realization modules operate at the level of text. *The Merchant of Venice* system renders its discourses as cinematic animations with dialogues, therefore it includes a third level of representation in a particular medium.

Multi-plot stories

Another very important fact about narrative is that most narratives that we come in contact with in our daily life – novels, movies, TV series – involve a combination of more than one plot line. This phenomenon has only recently started receiving attention by the research community on computational narrative, but important insights have arisen from the effort.

The work of Fay (2014) proposes an algorithm for story generation that combines material from existing stories into new ones. It relies on the GENESIS system (Winston 2014) to extract from existing stories both a set of models for the characters involved and the subsets of events from the story in which each particular character is involved – referred to as a *plot thread*. Then it generates new stories by weaving together a selection of plot threads from different stories into a new story. The procedure for this is designed to ensure that the subplots for the different characters combined in the story are compatible, and that the resulting time line of events for the story is consistent. Fay allows for an operation of binding characters from different plot threads, based on the information available in their character models.

Porteous, Charles and Cavazza (2016) extend the plan-based paradigm for narrative generation to consider stories with several plot lines. Their solution receives as input the number of subplots desired and the frequency of switching between subplots. The algorithm selects a particular subplot to operate on, generates the required number of segments for it, then switches to another subplot, and it iterates until all subplots are closed. The system was tested by generating stories for a serial drama in the medical domain, with an interface that presents the story visually to the user as cinematic sequences of animated characters.

A different approach considers the use of evolutionary solutions for weaving together a set of plot lines (Gervás, Concepción and Méndez 2022). The solution they propose considers two different operations: one of *discourse planning* – which decides on the relative order of scenes from different subplots in the final sequence for the discourse – and one of *character fusion* – which may decide that two roles played by different characters in different subplots may be played by one single character in the final story. The procedure operates by building a population of story drafts and applying to them mutation and crossover operators – which revise the decisions in them – over a number of generations, selecting for each successive generation the drafts considered to be the fittest. Fitness is modelled by a set of metrics that capture consistency constraints at a semantic level – characters not being active after they are dead or before they are born – and desirable features at the discourse level – such as ensuring that there is an overarching plot for the whole story.

The type of narrative being constructed in these systems differs from the previous cases in that although there is a single fabula for the story, that fabula may now involve more than one plot. This simply means that the events in the fabula may not be a single connected causal graph but a number of graphs that are connected internally but not connected to one another. Each of these graphs constitutes a separate plot line. The various plot lines are not entirely independent of one another in that they may share some characters. This is often achieved by having principal characters in one plot play minor roles in another, or simply by having scenes from more than one plot take place at the same location, so that they participate in the perception of characters from other plots.

Fay (2014) operates at the level of discourse, though its character models, which do not appear in the final outcome, should really be considered a partial representation of the source story world. Because it relies on character presence rather than causality to establish the boundaries of each plot thread, its so called 'plot threads' may not entirely satisfy the definition of plot used here, which is based on Forster's requirement of causality. Similarly, the plot lines used by Gervás, Concepción and Méndez (2022) come predefined as plot templates, so there will be causal links tying together all the events in each plot line only if these plot templates themselves satisfy the requirement. In contrast, each of the subplots employed by Porteous, Charles and Cavazza (2016) is constructed as a part of the underlying fabula that is a connected causal graph built by a planner, therefore it satisfies the requirements for being considered a plot. The fabula for these stories is then a set of connected causal graphs. The relative ordering of spans from different subplots in the discourse is generated dynamically during the construction of the fabula, driven by the procedure for switching between subplots. As for the Merchant of Venice, the final story is represented cinematically and it allows for user intervention.

Narrative prose

Callaway's AUTHOR architecture for rendering narative prose (Callaway 2002) addressed the problem of generating narrative prose for a given description of a narrative discourse. The AUTHOR system operates from a system that included both a representation of the set of events to tell (fabula) and a representation of the order in which to tell those events (referred to as a *narrative stream*, and corresponding closely to the concept of discourse discussed above). It also relied on additional inputs – required to address the tasks of sentence planning – such as an ontology where the entities involved in the story were represented. Callaway's solution relied on implementations of well-known sentence planning tasks from the field of NLG, such as referring expression generation, lexical and syntactic choice, and aggregation. The system addressed relevant narratological concepts such as time of narration – considered to establish appropriate tense for verbs - distance – that impacts on how dialogues are reported – and voice – driving choice of first or third person as required. The underlying technology was the FUF/SURGE unification-based functional systemic grammar (Elhadad 1993) that implemented the surface realization task – rendering the input into text – in specific languages. The AUTHOR architecture was instantiated in the STORYBOOK narrative prose generator. STORYBOOK accepted a discourse plan for the story of Little Red Riding Hood and it was capable of generating a range of different text variations for that story. These variations were systematically evaluated. However, as the evaluation did not go beyond the particular example of the Little Red Riding Hood story, it is likely that porting the solution to other domains may involve a costly effort in additional knowledge engineering of the required resources.

Computational models of processing of narrative

Many of the efforts reviewed in Section 'Computational models for generating narrative' had identified that, for the generated story to be valuable, the story generation system needed to have a representation of many aspects that were then not explicitly present at the surface level of the generated story, such as character goals and non-executed events from failed plans. These aspects are relevant because they are inferred naturally by the reader as she reads the story, and they play a role in the attribution of quality to a story. Over time, this type of consideration leads to a number of research efforts aimed at modelling computationally the response of a reader to a given narrative.

Emotions and tensions between characters

The Mexica system (Pérez y Pérez 1999) included a specific module for reading existing stories to create the materials from which new stories would be built. This module read an inspiring set of stories and built from them a set of knowledge structures that represented not just the events in the stories, but also an additional layer of representation that captured at each point in the story the emotions experienced by the characters and the tensions that arose between them. This additional layer of information was constructed based on a dictionary of story actions, which associate to each event preconditions and post-conditions in terms of character emotions and tensions between characters. The construction algorithm then chained up actions for new stories by taking into account patterns of co-occurrence of actions, emotions, and tensions in prior stories. The final stories generated by the system reproduced the sequence of actions obtained in this manner but generally avoided any mentions of the emotions and the tensions that had driven the construction process, leaving these to be inferred by the readers.

The Mexica system generally operates at the level of discourse, in the sense that it relies on knowledge resources in which the relative order of events in the inspiring stories is explicitly represented, and it constructs sequences of events in which the relative order is determined during construction based on these resources. There is no step in the representation that might be considered a fabula. However, the representation of affinities and tensions between must be considered a partial representation of the source storyworld, in the sense that it constitutes information about what happened in the world that is available to the author but not directly included in the story. It is important to note that this kind of information is usually inferred from the text by a reader, and therefore it may also be considered to be part of the storyworld in the sense of Ryan (1991) – as the representation of the world of the story that the reader builds from the text.

Comprehension

The work of Niehaus and Young (2014) expands on existing work on modelling of narrative with the aid of artificial intelligent planners and proposes an explicit computational model of a reader's comprehension process during reading. This model allows the system to predict which elements of the narrative are more likely to be foregrounded in the reader's mind at a given point in the story, and, based on those, what inferences the reader is likely to make. Niehaus uses this computational model to inform a process of generating narrative discourse by selecting content from an event log and organizing it into a linear sequence of statements. This is a slightly different approach to the generation of narrative in that it operates over an input that is already a plan-based representation of a story – as produced by some of the other systems reviewed above – and generates for it a discourse that is easier to comprehend by the reader.

Here we find again a system that receives as input a fabula and produces a discourse for it. The interesting feature of this system is that for a given fabula it can produce a number of different discourses, and it includes a module for selecting particular versions of the discourse that it considers will be easier to understand by a reader. The criteria that support these decisions are based on models of how readers process discourse as they read it.

Suspense

An important ingredient of narrative is the inducement of a feeling of suspense in the reader. This has been explored in computational terms in recent years, and this effort has led to a number of valuable insights on the nature of narrative and the elements that play a role in it.

Two different systems have relied on the definition of suspense proposed by Gerrig and Bernardo (1994), which assumes that the feeling of suspense experienced by a reader is related to the quality or quantity of possible avenues of escape available to the protagonist from some oncoming negative outcome.

The DRAMATIS system by O'Neill (2013) processed an input discourse and built an associated search space of possible outcomes to the situation as described at a given point in the story. The view constructed by the system of the possible outcomes open to a character was constrained by a model of what elements of the story are foregrounded in the reader's mind at a given point and the set of inferences that the reader can make from them. From these considerations, DRAMATIS produces a rating of suspense for each point of the story.

The DRAMATIS system is not a generator, it only processes a discourse and predicts a sequence of ratings of suspense for each point of the discourse. To achieve this, it constructs for each point of the discourse a forward projection of potential outcomes – restricted to the protagonist – which could be considered instances of possible storyworlds in the sense of Ryan (1991).

The Suspenser system (Cheong and Young 2015) processed a plan-based representation of a story and composed for it a discourse tailored to achieve a desired impression of suspense. To achieve this, the system first selects a subset of the events in the story that are considered fundamental for the story – these will be included in every version of the story – then considers which other events of the story might be added to maximize the perception of suspense. To inform the decisions involved, it applies a module that models the inferences that the reader might make using an additional planner to construct possible continuations to the situation. The measure of suspense at a given point is taken as the inverse of the number of possible positive outcomes found.

Suspenser involves complex handling of the narratological concepts we are looking at. In this case the representation of the story as a plan needs to be considered a source storyworld, in the sense that only part of it will be selected to make up the fabula of the story. Selection of which events in this source storyworld are included in the fabula is made based on criteria similar to those used by DRAMATIS: A planner is used to build potential continuations, which need to be considered possible storyworlds. Discourse is built for the resulting fabula by ordering the resulting events. In this way, Suspenser is operating on the story it builds at the levels of source storyworld, fabula, discourse and possible storyworlds that the reader might construct. It also includes a module to produce text from the discourse.

The IRIS (Intention Revision in Storytelling) story generation system (Fendt and Young 2017) relies on the definition of suspense used by the Suspenser system to develop stories together with a belief/desire/intention framework of intentionality for characters to build stories that create suspense from the structure of their events and their timing rather than from the way they are presented. It achieves this by developing the story in conjunction with an explicit representation of the plans of the protagonist. As soon as a plan for the protagonist is generated, the system draws upon a set of suspense action templates to introduce threats to the plan. This will force the protagonist to revise her intentions and replan.

The IRIS system exploits the same definition of suspense proposed in Suspenser but it actually uses it to inform the planner that is building the fabula. This allows it to introduce events that will threaten the protagonist's plan. To do this, it operates directly on the construction of the fabula, rather than selecting it out of a source story world. In addition, it includes explicit representation of all the other levels considered by Suspenser.

An alternative view on suspense is proposed by Doust (2015). This approach is based on a psychological model of narrative (Brewer and Lichtenstein 1982) that takes into account both the role of plan comprehension and story schemas. Doust postulates the concept of *narrative threads* as formal descriptions of a reader's expectations about what might happen next in a given story. These narrative threads tend to be motivated either by plan recognition or by identification of story schemas. The proposed model constructs the set of narrative threads triggered by the accumulated events at a given point in the story, and computes an estimate of the suspense perceived by a reader at that point based on the number of narrative threads still requiring completion and/or the number of threads that conflict with other active threads.

Like DRAMATIS, the model proposed by Doust (2015) does not generate stories, it is a model of how readers process the discourse for a story. In a similar fashion, it operates by building forward projections of the action for each point of the discourse, and establishing a rating for suspense for the corresponding point by applying specific metrics to the number of projections built. The projections are in this case not built by a planner but by matching the actions in the discourse to know patterns of behaviour.

Embedded narratives

A different important aspect of narrative is the possibility of including in a given story – the *frame story* – events where some character tells a story – the *embedded story*. Such stories told within stories give rise to distinction of different *narrative levels*: The telling of a new embedded story opens a new narrative level (Herman, Jahn and Ryan 2010). This phenomenon occurs with very high frequency in the stories that humans consume, but it has only very recently been addressed by computational research.

The need to have different levels of representation for each narrative level, and the need to account for the recursive nature of the embedding process by means of a stack of such representations was first proposed by Gervás (2021). A more recent development (Gervás 2022) defines a computational model of how to identify the relationships between the events referred to in different narrative levels: Sometimes the embedded story describes a completely different world, sometimes it introduces events in the current story world in the form of a flashback and sometimes it actually refers to events already told in the frame story. Based on that model it proposes simple mechanisms for establishing relative chronology for events about the same story world being referred to in different narrative levels. These efforts have only started to scratch the surface of the problem of computational treatment of embedded stories, which remains an open challenge.

This is yet another model that focuses on reading stories. In this case, it attempts to reconstruct the fabula underlying a given discourse, in that it aims to establish the real chronological order between events in the discourse. This is done by identifying events mentioned in embedded discourses and establishing their relative timing with respect to the event in the frame story. These systems work on inputs that are conceptual descriptions of discourse, and they do not include modules to process text onto discourse.

Possible worlds, fictional minds and tellability

The work of Berov (2021) describes a story generation system that includes many of the aspects mentioned to this point: specific character goals, possible words to represent outcomes of various – possibly conflicting – plans, and emotions. It does so by combining the theory of possible worlds by Ryan (1991) with the theory of how readers understand novels proposed by Palmer (2004) – he suggests it is primarily by following the functioning of the minds of characters in the novel in terms of storyworlds constructed by them from the reality available to them. In the process, Berov postulates the concept of a *narrative system*, which is 'a system whose dynamics, the movement of possible worlds, lead to the emergence of a plot', where plot is understood as a 'causal network of happenings, actions and mental events that is consistent with the narrative's discourse'. A narrative system as postulated by Berov includes: The Textual Actual World it constructs, the possible worlds considered by its characters and the set of potential plots that can be constructed from these elements. In this system, all the constituent parts of plot are defined in reference to the mental functioning of the characters.

This is an important extension beyond prior attempts at computational modelling of narrative: it represents not only characters but their thoughts. Prior systems had represented the salient view characters had of the world, character emotions or character goals and plans, as elements relevant to constructing believable stories. Berov defends the hypothesis that it is a fundamental goal of narrative to represent and communicate views on the minds of characters. He does so citing a considerable body of knowledge in narratology and psychology. His system includes modules that represent character personalities based on psychological theories of cognition. Based on this it presents an account of plot that captures character personalities, character emotions at particular points in the story, and affective causes and conditions for agents' plans.

Berov includes a computational definition of *tellability* as understood by Ryan (1991). This definition combines ideas of what makes a story worth telling and how to tell a story well. Tellability can then be used to select some options out of the set of possible plots for a given narrative system. By using this definition of tellability as fitness function, Berov develops an evolutionary solution that searches the space of possible plots from a given narrative system.

Reflections on computational models of cognitive tasks related to literature

The compilation of systems reviewed in Sections 'Computational models for generating narrative' and 'Computational models of processing of narrative' is not meant to be exhaustive. For a broader view on story telling systems the reader is invited to consult some of the existing surveys (Kybartas and Bidarra 2016; Gervás 2021). The selection of systems presented here is aimed to support a number of important reflections on the nature of the storytelling task from the computational point of view.

A subdivision of tasks in computational modeling of narrative

One of the main conclusions to be drawn from the review of artificial intelligence systems applied to narrative is the diversification that has occurred on the types of tasks addressed.

Early systems focused on the task of generating the outline of a story from scratch. They focus mainly on *fabula generation*. The TALESPIN (Meehan 1977), Fabulist (Riedl 2004) and Glaive (Ware 2014) systems all focus heavily on construction of fabula. Of the systems that generate multi-plot stories, the planning-based system by Porteous, Charles and Cavazza (2016) also includes substantial effort in fabula generation. All of these systems include an explicit representation of discourse as separate from the fabula they are building, and they include procedures for constructing the corresponding discourse from it. Most of them additionally include specific modules to transcribe the discourse to text, or sometimes into visual cinematics. Because of this, I prefer to consider them as systems for *narrative generation*, because even though the narratives they construct may be weak at the levels of discourse or text, they do explicitly represent all the levels and they include modules that implement the transitions between them.

A related but slightly different type of systems generate stories by building directly at the level of discourse, with no explicit representation of fabula. These are cases of *narrative discourse generation*. Systems of this type are those based on story grammars (Rumelhart 1975), the Mexica system (Pérez y Pérez 1999) and Minstrel (Turner 1993). Of the systems that generate multi-plot stories, those that operate from inputs that are already sequences of discourse (Fay 2014; Concepción Gervás and Méndez 2020) also count as discourse generators. All these systems mostly lack a representation at the level of fabula, but it is important to note that they do include modules for building text from the discourse.

At some point, a different kind of system appears, one that receives as input a conceptual description of the material to be included in a story – the fabula for the story – and generates a narrative discourse for telling it. This task we refer to as *composition of narrative discourse*. The solution for telling stories about a given chess game developed by Gervás (2014), the Suspenser system (Cheong and Young 2015), the generator for comprehensible discourse by Niehaus and Young (2014) and the interactive version of the Merchant of Venice (Porteous, Cavazza and Charles 2010) can all be considered systems for the composition of narrative discourse. The process of selecting a specific plot out of the set available from a given narrative system, as discussed by Berov (2021) also represents an instance of composition of narrative discourse. Systems for composition of narrative discourse focus on the transition between fabula and discourse. As such, they only address a particular stage of the process of narrative generation.

Finally, another type of system emerges, focused on generating a rendering as text of a given narrative discourse. We refer to this as *narrative text generation*. The STORYBOOK system Callaway (2002) for generating narrative prose from an input representation of the reference fabula and the desired discourse is an instance of narrative text generation. It is important to note that most fabula generation systems and system for the composition of narrative discourse include specific modules for rendering their outputs as text. Such modules may be considered baseline solutions for the task of narrative text generation.

As for cases that operate at more than one level, the Merchant of Venice system by Porteous, Cavazza and Charles (2010) is essentially a discourse composition system but it incorporates the ability to adapt the fabula to generate discourse with particular perspectives. It also includes a module for rendering the discourse as an interactive visual cinematic. The IRIS (Fendt 2014) system stands out because it includes significantly elaborate modules for both fabula generation and discourse composition.

A final type of system receives as input a discourse – usually in the form of a succession of updates rather than as a single whole – and at each point returns a representation of conceptual constructs of some kind as might be produced in the mind of a reader during reading. Examples of these conceptual constructs may be plans attributed to characters, emotions, estimates of the subset of story entities foregrounded in the reader's mind or impressions of suspense perceived. The versions of these constructs built after the processing of a given update to the discourse are specific to that point in the processing and they will change with later updates. These systems or sub-modules constitute computational models of abilities applied by the reader. We will refer to this task as *narrative interpretation*. The GENESIS system (Winston 2014), the DRAMATIS system (O'Neill 2013), the model of narrative processing to identify suspense proposed by Doust (2015) and the models for processing embedded stories reviewed above (Gervás 2021, 2022) all constitute pure instances of narrative interpretation tasks. The Mexica system (Pérez y Pérez 1999) and the Suspenser system (Cheong and Young 2015) both include modules that address some aspect of the reading task, and therefore can be considered to include instances of narrative interpretation. The computational definition of tellability by Berov (2021) is also an instance of narrative interpretation.

The recognition of these five different tasks as separate challenges for computational models of narrative is a very important insight. However, it needs to be considered in conjunction with an awareness that, to address narrative capabilities as exercised by humans, it is often necessary to combine these tasks in different ways. Examples of this type of situations have already appeared in the set of systems reviewed: combining interpretation and discourse generation (Pérez y Pérez 1999; Fay 2014; Cheong and Young 2015), discourse composition and interpretation (Cheong and Young 2015; Niehaus and Young 2014; Berov 2021), combining interpretation, fabula generation, and discourse composition (Porteous, Cavazza and Charles 2010; Fendt 2014).

In terms of comparison with earlier efforts, the review presented here shows that there has been significant progress. An early paper describing collaboration between narratologists and researchers in artificial intelligence (Gervás et al. 2006) described as an important shortcoming of the story generator algorithms available at the time the fact that they did not allow representation of the differences between *histoire* – what we have been referring to as fabula – and *discours* – what we have been referring to as discourse. The present review shows that this is no longer the case. Most of the more recent systems allow explicit representation of both fabula and discourse, and although some focus more on generating fabula and some focus more on generating discourse, most are aware of the distinction and indeed include features that address the process of constructing one from the other, or even tailoring the story at both levels to achieve particular effects on the reader. In contrast, another important shortcoming noted at the time – namely that STORYBOOK was the only system that addressed generation of narrative prose at a reasonable level of complexity – still remains extant.

It is also important to consider that the set of subtasks of storytelling identified to this point is very probably not exhaustive. As more research is carried out, additional tasks are likely to appear. In fact, it is very probable that the shortcomings currently patent in the outcomes of storytelling systems when compared to human output relate closely to the fact that there are still important abilities that humans apply to the development of narrative that computational systems are not yet taking into consideration. This possibility merits some further analysis, presented in the section below.

A computational view on the context of narrative

If one considers the context in which narrative occurs in the wild, it is clear that approaches that focus exclusively on the narrative itself as a product are likely to miss out on many important characteristics of the phenomenon as a communication mechanism. Although we have become accustomed to the view of narrative as an object, in the way of a painting or a sculpture, it is important to consider that it originates as an act of communication between an author and a reader. From a computational point of view, the narrative itself is the result of an encoding process – the act of writing – that compresses a very complex construct that would be postulated to exist in the mind of the author into a very constrained format as a linear sequence of statements. The act of reading constitutes a corresponding process of decoding in which the linear sequence of statements is expanded into another very complex construct in the mind of the reader. This view is cited by Berov (2021) in support of several of his arguments.

It is important to note that this description needs to be hedged with a number of constraints. First, there is a huge difference in structural complexity between the information about the story held in the mind of the author and the information about the story actually encoded in the narrative. As the more elaborate computational models described above show, there may be a significant amount of information considered by the author that is not actually present explicitly in the narrative: character goals, character emotions, details of character plans that never get put into practice... Having observed this for the phenomena that have been treated computationally, it is extremely likely that similar differences occur for phenomena that have not been explicitly addressed yet in computational systems. Second, a similar difference holds between the information about the story encoded in the narrative and the information about the story that the reader reconstructs in her mind. The various models of narrative interpretation discussed above show how readers may extract via different types of inference very relevant information about the story that is not explicitly encoded in the narrative itself. The differences between the models – some consider plans but not emotions, some consider emotions but not plans – also show that it is perfectly possible to address the task with partial information and not realize that relevant information may be missing. This suggests that the features addressed to this point may be only the tip of the iceberg of what human readers really do process. Third, there is no guarantee that the constructs created by the reader during interpretation be faithful to the ones conceived by the author.

The concept of a narrative system as proposed by Berov (2021) is the closest computational system to these complex conceptual constructs that might represent the information about a narrative that passes through authors' or readers' minds during their corresponding interactions with it.

The actual existence of a full mental picture of the story as a unified construct persistent over time in the mind of the author is of course arguable. Because the narrative is produced sequentially – say, one page or even one sentence at a time – it is entirely possible that the full text come about without the author ever having a full picture of the complete story in their heads. In pragmatic terms, it would be sufficient that the author be capable of applying small tests for coherence within portions of the story and their partial recall of the structure of the whole, or larger tests assisted by memory aids such as plot outlines or index cards for characters. Such a view is argued by Dehn (1989) in her model of how the limitations of human memory influence and participate in the process of writing. The obvious need for these operations suggests a further task of *revision of narrative* should be considered (Gervás and León 2016).

In a similar fashion, these limitations in the amount of information we can remember clearly play a role in the way a reader processes a text. The models considered by O'Neill (2013) and Niehaus and Young (2014) both rely on the concept of salience, which arises from the recognition that readers cannot keep in their minds a complete picture of everything they have been told in a story. Instead, they operate with a restricted subset of what they have considered more relevant, or simply with a fragment that contains only the most recent information.

To further complicate the situation, we need to take into account that both author and reader are agents capable of complex cognition, and that both operate with a clear view of this very description of the process. Therefore both will apply all their cognitive potential to the task. This will very likely lead to the author having in mind a mental model of how a reader would process the narrative, and the reader having in mind a mental model of how the author produced it. The consideration of such mental models can certainly be exploited to further refine computational models of the processes of reading or writing. Several of the story generation systems reviewed above – indeed, all the more refined ones – include computational models of the reader to be used in support of the generation process. To my knowledge, there are still no systems of narrative interpretation that include a model of the author. The explicit consideration of the peculiarities of the author – personality, background, ideology – has clearly been a subject of interest in narratology for many years, so this is still an open avenue for computational exploration.

The advent of neural network storytelling

The review presented in this paper has made no mention of existing efforts to apply neural networks to storytelling tasks. There has indeed been great interest from the neural computing community to address the generation of stories. The recent advent of ChatGPT[3] appears to have achieved outstanding results in terms of generating correct stories on demand. The number of avenues of exploitation for this type of functionality has blossomed (Calderwood et al. 2020). However, the very nature of the neural network approach, strictly based on training from large volumes of data with little chance for the researcher to intervene in the process, implies that these efforts have had absolutely no input from the existing body of knowledge on narratology or any other discipline. It is for this reason that they are excluded from the reviews in this paper.

Conclusions

The review of existing storytelling systems shows that progress in the field has explored many different aspects that play a role in narrative. However, most systems tend to address a particular aspect and sidestep the others, even when these other aspects have already been proven to be relevant in some other system. This is in part justified by the generic engineering principle for addressing complex problems that suggests smaller sub-problems should be isolated and solved first. However, the principle requires that once the smaller sub-problems have been solved attempts are made to combine them into solutions for the complex problem. This second step has yet to be addressed in most cases. Part of the problem here may be that most of these efforts have historically arisen as doctoral theses, and the current view on requirements for a thesis by most universities proscribe the consideration of large scale complex problems. A very interesting example is the work of Berov (2021), which brings together a plethora of concepts from narratology and techniques from artificial intelligence.

The review also shows four different types of system functionality. Fabula generation receives some specification of what story is desired and generate the set of events that will appear in the story, described here as the fabula. Discourse composition receives some data structure that describes the events to be included in the story – the fabula – and generate a discourse to tell the story. Text generation receives a data structure that describes the narrative stream – the discourse – and generate a rendition of it as text. Narrative interpretation receives a representation of narrative discourse and generate, for each point in the discourse, a data structure that describe some aspect relevant to the interpretation of the discourse. The use of these more specific terms to describe what a system does allows us to identify that most existing systems involve complex combinations of these functionalities.

An important conclusion that derives from the compilation of systems discussed is that, although it is true that there are no systems that integrate all the levels involved in construction of a complete narrative, solutions do exist for many of the parts. Many of the systems reviewed were never intended as generators of complete narratives, but as fabula generators, discourse composers, or prose realizers. The possibility of combining such systems as modules in a more complex narrative generator exists at the theoretical level.

In terms of specific aspects, it is commendable that the planning-based approach to narrative has managed to address so many different relevant aspects, and that later research efforts often consider aspects addressed by prior solutions. The consideration of emotions in story generation systems, although pioneered in the 1990s, has not been considered by any later systems. The issue of perspective has seen some progress, but there is still a need to explore the way human authors often exploit differences in knowledge about the story world between different characters or between the characters and the audience to achieve narrative effects. Generation of multi-plot stories has been addressed at a high level of abstraction, yet the existing solutions would greatly benefit from enrichment with the various aspects already tested on single plot stories. The small steps of progress being made on the processing of embedded stories should lead to richer models of narrative interpretation in the future. When they do, narrative generation solutions should consider how they might introduce embedded stories in their outputs. Since embedded stories are often used by human authors to manage the flow of knowledge about the story between the characters, systems that consider perspective might be well served by a solution that can handle the use of embedded stories.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The authors received the following financial support for the research, authorship and/or publication of this article: This paper has been partially funded by the project CANTOR: Automated Composition of Personal Narratives as an aid for Occupational Therapy based on Reminescence, Grant. No. PID2019-108927RB-I00 (Spanish Ministry of Science and Innovation) and the ADARVE (Análisis de Datos de Realidad Virtual para Emergencias Radiológicas) Project funded by the Spanish Consejo de Seguridad Nuclear (CSN), Grant Ref. SUBV-20/2021.

ORCID iD

Pablo Gervás https://orcid.org/0000-0003-4906-9837

Notes

1. For more detailed analysis of this line of research see (Gervás, 2021).
2. 'Plot' in the basic sense used by lay people.
3. https://openai.com/blog/chatgpt/.

References

Bal, M., and C. van Boheemen. 2009. *Narratology: Introduction to the Theory of Narrative*. Toronto: University of Toronto Press. ISBN 9780802096319.

Barthes, R. 1967. "The Death of the Author." *Apsen* 5–6. doi: 10.1515/JLT.2007.008.

Berov, L. 2021. "From Narratology to Computational Story Composition and Back–An Exploratory Study in Generative Modeling." PhD Thesis, University of Osnabrueck.

Brewer, W. F., and E. H. Lichtenstein. 1982. *Stories Are to Entertain: A Structural-affect Theory of Stories*. Technical Report No. 265. Center for the Study of Reading. Urbana: University of Illinois. https://eric.ed.gov/?id=ED222854

Bringsjord, S., and D. A. Ferrucci. 1999. *Artificial Intelligence and Literary Creativity? Inside the Mind of BRUTUS, a Storytelling Machine*. Mahwah: Lawrence Erlbaum Associates.

Calderwood, A., V. Qiu, K. I. Gero, and L. B. Chilton. 2020. "How Novelists Use Generative Language Models: An Exploratory User Study." In: *HAI-GEN+ user2agent@ IUI*.

Callaway, C. B. 2002. "Narrative Prose Generation." *Artificial Intelligence* 139 (2): 213–252. http://dx.doi.org/10.1016/S0004-3702(02)00230-8.

Cheong, Y. G., and R. M. Young. 2015. "Suspenser: A Story Generation System for Suspense." *IEEE Transactions on Computational Intelligence and Artificial Intelligence in Games* 7 (1): 39–52.

Colby, B. N. 1973. "A Partial Grammar of Eskimo Folktales1." *American Anthropologist* 75 (3): 645–662.

Concepción, E., P. Gervás, and G. Méndez. 2020. "Exploring Baselines for Combining Full Plots Into Multiple-plot Stories." *New Generation Computing* 1–41. doi: 10.1007/s00354-020-00115-x. https://link.springer.com/article/10.1007%2Fs00354-020-00115-x.

Culler, J. 2004. "Story and Discourse in the Analysis of Narrative." *Narrative Theory: Critical Concepts in Literary and Cultural Studies* 1: 117–131.

Dehn, N. 1989. "Computer Story-Writing: The Role of Reconstructive and Dynamic Memory." PhD Thesis, New Haven: Yale University.

Doust, R. 2015. "A Domain-Independent Model of Suspense in Narrative." PhD Thesis, Milton Keynes: The Open University.

Elhadad, M. 1993. "Using Argumentation to Control Lexical Choice: A Functional Unification Implementation." PhD Thesis, Columbia University, New York, NY, USA. UMI Order No. GAX93-33756.

Fay, M. P. 2014. "Driving Story Generation with Learnable Character Models." PhD Thesis, Cambridge/Mass.: Massachusetts Institute of Technology.

Fendt, M. 2014. "Leveraging Intention Revision in Narrative Planning to Create Suspenseful Stories." PhD Thesis, Raleigh: North Carolina State University.

Fendt, M. W., and R. M. Young. 2017. "Leveraging Intention Revision in Narrative Planning to Create Suspenseful Stories." *IEEE Trans. Comput. Intell. AI Games* 9 (4): 381–392. doi: 10.1109/TCIAIG.2016.2593977.

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

Genette, G., J. Lewin, and J. Culler. 1983. *Narrative Discourse: An Essay in Method*. Ithaca: Cornell University Press.

Gerrig, R. J., and A. B. I. Bernardo. 1994. "Readers As Problem-solvers in the Experience of Suspense." *Poetics* 22 (6): 459–472.

Gervás, P. 2014. "Composing Narrative Discourse for Stories of Many Characters: a Case Study Over a Chess Game." *Literary and Linguistic Computing* 29 (4): 511–531.

Gervás, P. 2021. "Computational Models of Narrative Creativity." In: *Artificial Intelligence and the Arts*, edited by P. Machado, J. Romero, and G. Greenfield. Cham: Springer.

Gervás, P. 2021. "A Model of Interpretation of Embedded Stories." In: *Text2Story: 4th International Workshop on Narrative Extraction from Texts*. CEUR Workshop Proceedings, Lucca: CEUR Workshop Proceedings.

Gervás, P. 2022. "A Discourse Interpretation Engine Sensitive to Truth Revisions in a Story." In: *Tenth Annual Conference on Advances in Cognitive Systems*. Arlington, Virginia.

Gervás, P., E. Concepción, and G. Méndez. 2022. "Evolutionary Construction of Stories that Combine Several Plot Lines." In: *Computational Intelligence in Music, Sound, Art and Design – 11th International Conference, EvoMUSART 2022*. Madrid: Springer. doi: https://doi.org/10.1007/978-3-031-03789-4_5.

Gervás, P., and C. León. 2016. "Integrating Purpose and Revision into a Computational Model of Literary Generation." In: *Creativity and Universality in Language*, edited by Mirko Degli Esposti, Eduardo G. Altmann, and François Pachet, 105–121. Cham: Springer.

Gervás, P, B. Lönneker, J. C. Meister, and F. Peinado. 2006. "Narrative Models: Narratology Meets Artificial Intelligence." In: *International Conference on Language Resources and Evaluation. Satellite Workshop: Toward Computational Models of Literary Analysis*. 44–51. Genova.

Herman, D, M. Jahn, and M. Ryan. 2010. *Routledge Encyclopedia of Narrative Theory*. London and New York: Taylor & Francis.

Kybartas, B., and R. Bidarra. 2016. "A Survey on Story Generation Techniques for Authoring Computational Narratives." *IEEE Transactions on Computational Intelligence and AI in Games* 9 (3): 239–253.

Lang, R. R. 1999. "A Declarative Model for Simple Narratives." In: *Proceedings of the AAAI Fall Symposium on Narrative Intelligence*, 134–141. AAAI Press.

Meehan, J. R. 1977. "Tale-spin, an Interactive Program that Writes Stories." In: *Proc. IJCAI:* 91–98.

Nau, D., M. Ghallab, and P. Traverso. 2004. *Automated Planning: Theory & Practice*. San Francisco: Morgan Kaufmann Publishers.

Niederhoff, B. 2009. "Perspective/point of View." In: *Handbook of Narratology*. edited by P. Hühn, J. Pier, W. Schmid, and J. Schönert, 384–397. Berlin, New York: De Gruyter. doi: https://doi.org/10.1515/9783110217445.384.

Niehaus, J. M., and R. M. Young. 2014. "Cognitive Models of Discourse Comprehension for Narrative Generation." *Lit. Linguistic Comput.* 29: 561–582.

O'Neill, B. 2013. *" Computational Model of Suspense for the Augmentation of Intelligent Story Generation*. PhD Thesis, Atlanta: Georgia institute of Technology.

Palmer, A. 2004. *Fictional Minds*. Lincoln: University of Nebraska Press.

Pemberton, L. 1989. "A Modular Approach to Story Generation." In: *Proceedings of the Fourth Conference on European Chapter of the Association for Computational Linguistics*, EACL '89. Stroudsburg, PA: Association for Computational Linguistics, 217–224. https://doi.org/10.3115/976815.976845.

Pérez y Pérez, R. 1999. "MEXICA: A Computer Model of Creativity in Writing." PhD Thesis, Sussex: The University of Sussex.

Porteous, J, M. Cavazza, and F. Charles. 2010. ''Narrative Generation Through Characters' Point of View." In: *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010),* Vol. 1-3, Toronto. 1297–1304.

Porteous, J, F. Charles, and M. Cavazza. 2016. "Plan-based Narrative Generation with Coordinated Subplots." In: *European Conference on Artificial Intelligence (ECAI 2016)*, Vol. 285. IOS Press, 846–854.

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

Ricoeur, P. 1976. *Interpretation Theory: Discourse and the Surplus of Meaning*. Fort Worth: TCU press.

Riedl, M. 2004. "Narrative Planning: Balancing Plot and Character." PhD Thesis, Department of Computer Science, Raleigh: North Carolina State University.

Riedl, M. O., and V. Bulitko. 2013. "Interactive Narrative: An Intelligent Systems Approach." *AI Magazine* 34 (1): 67

Rumelhart, D. E. 1975. "Notes on a Schema for Stories." *Representation and Understanding: Studies in Cognitive Science:* 211–236.

Ryan, M. L. 1991. *Possible Worlds, Artificial Intelligence, and Narrative Theory*. Bloomington: Indiana University Press.

Ryan, J. 2018. *Curating Simulated Storyworlds*. Santa Cruz: University of California.

Ryan, M. L. 2019. “From Possible Worlds to Story Worlds.” In *Possible Worlds Theory and Contemporary Narratology*, 62–87. University of Nebraska.

Shakespeare, W. 2000. *The Merchant of Venice* (1598), edited by Cedric Watts. Ware: Wordsworth Editions.

Shoham, Y., and K. Leyton-Brown. 2008. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. New York: Cambridge University Press.

Stevenson, R. L. 1883. *Treasure Island*. 1st ed. London: Cassell & Company.

Terras, M, J. Nyhan, and E. Vanhoutte. 2016. *Defining Digital Humanities: A Reader*. London: Routledge.

Theune, M., S. Faas, A. Nijholt, and D. Heylen. 2003. “The Virtual Storyteller: Story Creation by Intelligent Agents.” In: *Proceedings of the Technologies for Interactive Digital Storytelling and Entertainment (TIDSE) Conference*, vol. 204215: 116.

Todorov, T. 1980. “The Categories of Literary Narrative.” *Papers on Language and Literature* 16 (1): 3–36.

Turner, S. R. 1993. “Minstrel: A Computer Model of Creativity and Storytelling.” PhD Thesis, Los Angeles: University of California.

Viktorovič, T. B. 1965. “Teorija Literatury. poėtika.” In: *Russian formalist criticism: Four essays*, edited by L. T. Lemon and M.J. Reis, vol. 405, 61–95. Lincoln: University of Nebraska Press.

Ware, S. G. 2014. “A Plan-based Model of Conflict for Narrative Reasoning and Generation.” PhD Thesis, Raleigh: North Carolina State University.

Winston, P. H. 2014. "The Genesis Story Understanding and Storytelling System: A 21st-century Step toward Artificial Intelligence." Technical report, *Center for Brains, Minds and Machines* (CBMM). https://dspace.mit.edu/handle/1721.1/100183.