ChatGeppetto - an AI-powered Storyteller

Edirlei Soares de Lima
Academy for Artificial Intelligence, Games and Media
Breda University of Applied Sciences
Breda, The Netherlands
soaresdelima.e@buas.nl

Bruno Feijó
Department of Informatics
PUC-Rio
Rio de Janeiro, Brazil
bfeijo@inf.puc-rio.br

Marco A. Casanova
Department of Informatics
PUC-Rio
Rio de Janeiro, Brazil
casanova@inf.puc-rio.br

Antonio L. Furtado
Department of Informatics
PUC-Rio
Rio de Janeiro, Brazil
furtado@inf.puc-rio.br

ABSTRACT

In this paper we introduce a novel highly interactive process to generate natural language narratives on the basis of our ongoing work on semiotic relations. To the two basic components of interactive systems, namely, a software tool and a user interface, we add a third component – AI agents, understood as an upgraded rendition of software agents. Our semiotic relations approach considers four ways of composing new narratives from existing narratives. Along what semioticians call the horizontal syntagmatic axis, one can form the new narrative by combining two or more previous narratives. Along the vertical paradigmatic axis, the new narrative may emerge as a similar version, which imitates the previous one, possibly in a different context. Along the depth meronymic axis, the hierarchic narrative levels, such as plot, event, and scene, are explored, allowing either expansion or summarization. Lastly, the antithetic consideration, rather than adding a dimension, aims at some form of reversal, through the adoption of opposite values. A fully operational prototype is described. Its name, ChatGeppetto, conflates the skilled Geppetto, who fashioned Pinocchio, an early case of artisanship-produced human level intelligence, with ChatGPT, which operates as the main AI agent component. To run the experiments, we concentrated on book narratives.

CCS CONCEPTS

• Computing Methodologies ~ Artificial Intelligence ~ Distributed Artificial Intelligence ~ Intelligent Agents

SBGames 2023, November, 2023, Rio Grande, Brazil

https://doi.org/XXX

KEYWORDS

Interactive Story Composition, Semiotic Relations, Artificial Intelligence, Book Narratives, Storyboards, Chatbots, ChatGPT

ACM Reference format:

Edirlei Soares de Lima, Bruno Feijó, Marco A. Casanova, and Antonio L. Furtado. 2023. ChatGeppetto - an AI-powered Storyteller. In *Proceedings of the Brazilian Symposium on Computer Games and Digital Entertainment (SBGames 2023). ACM, New York, NY, USA, 10 pages.* https://doi.org/XXX

1 Introduction

Generative AI can change how we generate game content, write stories, and even produce code for video games within a co-creation process. The video game industry already started to take advantage of the latest advancements on Artificial Intelligence. A recent example is Ubisoft, who revealed that they are using a generative AI tool called Ghostwriter to create "barks" (phrases and sounds made by non-player characters that occur mainly in the background during a triggered event). Blizzard is also developing its own AI tool, Blizzard Diffusion, in an apparent reference to the well-known Stable Diffusion model. Another example is Niantic, which used ChatGPT to help the development of content for the augmented reality game Peridot. The abovementioned examples and other recent attempts of the game industry to incorporate generative AI in game development are discussed in an article published by The New York Times [1].

In the present paper, we contribute to a more in-depth process to generate narratives using modules of generative AI as agents, where *semiotic relations* – previously introduced in the course of our Logtell[1] project [2][3][4] – are applied to support and guide

[1] http://www.icad.puc-rio.br/~logtell/

the narrative generation process. In this work, we concentrated our experiments on book narratives as they offer sufficiently complex narratives and a comprehensive database of stories.

Our early work in the Logtell project dealt mainly with interactive plan-based plot generation [5][6][7][8] and dramatization [9][10][11] wherein plots consisted of partially-ordered sequences of events. Our prototype tools, developed in logic programming notation, basically remained at what Mieke Bal [12] calls the *fabula* layer, where a simple abstract account of *what* happens is presented, to be followed by the *story layer* indicating *how* what happens is told – i.e. how the author structures the narrative – and, lastly, by the *text layer* which materializes the narrative in some chosen medium (which could be animation, video, film, etc., not necessarily written text).

On the one hand, the tools we introduced in our previous works were generally well-received by non-professional users [13][14]. They were happy to see the tools considering their decisions at branching points to produce plots conforming to their taste [15][16]. Also, they noticed the curious, unexpected solutions the plan generator found to seemingly impossible goals [17]. On the other hand, by not trying to contemplate Mieke Bal's story layer, the plots invariably looked relatively shallow, lacking the creative touch of talented professional authors.

We accordingly decided to extend our approach to the story layer, by providing access to sequences of events extracted from narratives told by professional writers. In an early attempt in this direction, we applied our previous work on semiotic relations to the analysis of folktale narratives, investigating how the sometimes-innumerous *variants* of a folktale type might have emerged, and then proceeding to show how to automatically create new variants by the interactive application of the semiotic rules. One of our papers [3][4], based on the *Index* compiled by Aarne and Thompson [18], dealt in particular with the folktale type AT 333, to which pertains the *Little Red Riding Hood* story, reported to have appeared in no less than 56 variants in different regional and temporal contexts [19].

In the present paper, we propose to advance two fundamental steps further in the application of the semiotic relations to the process of interactive narrative analysis and generation. Firstly, we decided to cope with natural language representation, under the form of *synopses*. As sources for the new narratives, we have concentrated on book narratives, restricting our choice to books that came to be popular enough to have their synopses recorded in publicly available electronic documents.

The second step, even more ambitious, was to raise the degree of user interactivity beyond the traditional permission to choose from a limited set of options at the branching points of some plot network structure. Now, as the new narrative generation proceeds along an event-by-event order, the user should be allowed to issue *instructions* – always in the form of natural language utterances – to influence the composition of the next event.

What makes our proposal look viable is the fact that current efforts to cope with the difficulties of natural language understanding and generation are obtaining increasing success, thanks to technical novelties, such as *transformers* [20], which allow the development of Large Language Models (LLMs). The GPT (Generative Pre-trained Transformer) model [21] has guided the production of the **ChatGPT** tool[2], now being tried for a variety of purposes – and to which we resort, as an *AI agent*, to attempt the transition from plots in abstract notation towards fluent and hopefully attractive literary storytelling.

To test our approach, we designed and implemented a prototype, named **ChatGeppetto**, which is now fully operational. The name is a homage to the Italian novelist Carlo Collodi, in whose famous book [22] the skilled Geppetto (Figure 1) fashioned Pinocchio from a piece of wood – an early case of human intelligence created by artisanship.

The present paper is organized as follows. Section 2 refers to related works. Section 3 explains how semiotic relations can support the process of generating new narratives. Section 4 describes the design and implementation of ChatGeppetto. Section 5 contains concluding remarks.

Figure 1: Geppetto in his workshop, according to ChatGPT and Stable Diffusion

2 Related Work

Our approach to support narrative generation is akin to *case-based reasoning* [23], accessing some "virtual library" of popular books to *search* for books whose structure has certain characteristics and could then be *adapted* in order to compose new narrative structures. Previous attempts to employ case-based reasoning to narrative generation include, among others, the work reported by Pablo Gervás [24]. *Searching* through a library is a pattern-matching activity that should not be too hard to implement if the library items are conveniently indexed. The subsequent *adaptation phase*, on the contrary, poses widely different sorts of unpredictable *blending* problems [25][26] calling for human interaction, ideally combined with the development of domain-oriented heuristics to provide assistance to the users.

The use of LLMs for narrative generation is a recent research topic. While some works focus on the use of fine-tuning and

[2] https://openai.com/blog/chatgpt

training techniques to improve narrative quality [27][28][29], others explore new strategies to handle the narrative generation process [30][31][32][33][34]. There are also works that focus on the evaluation of the writing capabilities of LLMs, such as Yuan et al. [35], which analyzes the potentials of the GPT-3 model in the collaborative co-writing process. Another example is Yang et al. [36], which present the results of a study comparing several aspects of narratives generated by ChatGPT with narratives authored by humans. Analyzing and reproducing the writing style of human authors is also a current research topic, as evidenced by the recent work of Garrido-Merchán et al. [37], which demonstrates the effectiveness of the GPT-4 model in reproducing H.P. Lovecraft's writing style.

In the context of fine-tuning, Värtinen et al. [27] presented a fine-tuned GPT-2 model to generate quest descriptions for role-playing games. However, a user study revealed that only one in five quest descriptions were considered acceptable by human critics. A different approach is explored by Fan et al. [28], which uses a large dataset of 300K human-written stories paired with corresponding writing prompts to train a convolutional seq2seq model for generating fluent text passages for storytelling. The use of reinforcement learning for fine-tuning LLMs is explored by Alabdulkarim et al. [29], who proposed a fine-tuned transformer-based language model using reinforcement learning to generate story continuations while seeking predefined goals extracted from a knowledge graph, which is used to guide the selection of candidate continuations.

Many strategies to handle narrative generation using LLMs were proposed in the recent years. In this context, Wang et al. [31] introduced a method that utilizes interpolation techniques to guide the GPT-2 model in the process of producing coherent narratives with user-specified target endings. A similar strategy is explored by Castricato et al. [32], who propose to treat story generation as a question-answering process that starts with the description of the final event of the story. By iteratively generating and answering "why" questions about character motivations, the system generates story events in a backward manner. Another similar approach is explored in Yang et al. [33], which proposes a framework that employs an LLM to generate a structured plan for the story, and then uses recursive reprompting to incorporate relevant information from the plan and story context into the generated narrative. The use of external knowledge to enhance the controllability of narrative generation is also explored by Xu et al. [30], who use a keyword predictor, a knowledge retriever, a knowledge ranker, and a conditional text generator to control the text generation. A different strategy is explored in Xiang et al. [34], which combines an LLM (GPT-3) with a symbolic story generation system (TALE-SPIN). By interleaving the GPT-3 output with the output of TALE-SPIN, their method ensures the logical coherence of the narrative, but limits story variation.

Although narrative generation using LLMs has now become a very active research topic, to the best of our knowledge, no previous work has explored the generation of new narratives by reusing and adapting existing stories through semiotic relations as proposed in this paper.

3 Narrative Generation Based on Semiotic Relations

Semiotic relations, as originally proposed by the Swiss scholar Ferdinand de Saussure [38], correspond to different forms of mental activity which, taken together, provide a model to structure linguistic discourse. One major guideline of our Logtell project is the *semiotic treatment* of both factual and narrative information, on the basis of four *semiotic relations* [2][3][4]. In this section we shall describe how they provide four different criteria to generate new narratives from existing narratives.

3.1 Semiotic Relations

The four basic semiotic relations treated here – *syntagmatic*, *paradigmatic*, *meronymic*, *antithetic* – have been drawn from the so-called four master tropes, a topic of high interest in the area of *semiotic research* [39]. The word "trope" comes from the Greek "τροποσ" from "τρεπειν", "to turn", with the intended meaning that such rhetorical figures "turn" i.e. alter the meaning of a word. Our four semiotic relations, together with their intuitive meaning, associated logical connectives, and corresponding tropes are displayed in Table 1.

Table 1. Semiotic relations

relation	**meaning**	**operator**	**trope**
syntagmatic	connectivity	and	metonymy
paradigmatic	similarity	or	metaphor
meronymic	hierarchy	part-of	synecdoche
antithetic	negation	not	irony

Present among the numerous rhetorical tropes compiled in Greco-Roman antiquity by Quintilian (c. 35-c. 100), these four tropes were later characterized as fundamental, first by Ramus (1515-72) and again by Vico (1668-1744) [39]. In modern times they were revived in Kenneth Burke's seminal study [40]. Their universality and completeness have been repeatedly emphasized, with the indication that they do constitute, according to Jonathan Culler, "a system, indeed *the* system, by which the mind comes to grasp the world conceptually in language" [41]. Applications to several topics have been reported, for instance to worldviews and ideologies by the historian Hayden White [42] and, in our own work, to mathematical proof methods [43], to database conceptual modelling [2] and, more relevant to the present discussion, to digital interactive composition of story-plots [3][4].

With respect to the names, we assigned to the proposed semiotic relations, the terms "syntagmatic" and "paradigmatic" correspond to the two linguistic axes postulated by Saussure [Saussure], who originally called the second axis "associative". His *horizontal* syntagmatic axis modelled sentences as aligned sequences of words connected in obedience to language syntax, whilst the *vertical* associative dimension would offer alternative choices to be placed below some of the sentence components.

The now universally adopted renaming of "associative" to "paradigmatic" was promoted by Jakobson [44], who argued convincingly, while discussing aphasia disorders, that "the

development of a discourse may take place along two different semantic lines: one topic may lead to another either through their *similarity* or through their *contiguity*. The metaphoric way would be the most appropriate term for the first case and the metonymic way for the second", in clear reference, respectively, to the paradigmatic and syntagmatic terminology.

In [45], wherein six types of *part-of* links are distinguished, one reads: "We will refer to relationships that can be expressed with the term 'part' in the above frames as 'meronymic' relations after the Greek 'meros' for part". About its associated trope, Burke affirms that "for synecdoche, we could substitute *representation*." [40]. In fact, breaking some signifying term into detail is an effective way to achieve its representation, whereas its identification should be possible by summarizing a detailed view. This zooming in/out variation of granularity level suggests that meronymic relations introduce *depth* as a third dimension.

Lastly, the term "antithetic" reflects the attribution of a *value scale* [39], either simply binary or graduated. No additional dimension is implied. According to Burke [40], the perspective induced by the associated irony trope refers to *dialectic*, which includes *antithesis* as a critical phase, expressing *negation*.

3.1.1 The Syntagmatic Relation

The existence of a *syntagmatic relation* between two narratives may justify their joining, often requiring laborious adaptations. The *Lancelot* and the *Grail* stories, both first narrated by Chrétien de Troyes [46] and closely related by their common attachment to Arthur's Round Table chivalric community, were *combined* in the so-called *Lancelot-Grail* cycle [47]. But adaptations were required – in special Lancelot's adulterous behavior with Arthur's queen disqualified him as protagonist in such religious mission. The solution was to give him a son, Galahad, who would inherit his heroic qualities but retain the purity of heart needed to achieve the Grail quest. Also, the two narratives are not joined by simply telling one and then the other. Instead, episodes involving Lancelot are interspersed with Grail quest incidents, in what French scholars have termed "entrelacement" (interlacing) [47].

3.1.2 The Paradigmatic Relation

A *paradigmatic relation* was applied to turn Galahad's tricky conception acceptable from a religious viewpoint. Since Lancelot idolized Arthur's queen Guinevere and could never be induced to make love with any other woman, a damsel, with the consent of her father, was placed in bed with Lancelot, after he had drunk a potion that made him believe that he had the queen on his side. What made the trick tolerable was that it *imitates* the Biblical narration of how Laban induced Jacob to heavy drinking during the marriage ceremony, so that the young man would believe that he was lying with Laban's eldest daughter, Lia, rather than with his beloved Rachel (*Genesis* 29:25-27).

3.1.3 The Meronymic Relation

As a case of employment of *meronymic relations*, a new narrative may be a retelling of a previous one, either by going into a deeper level of detail or through summarization. The *Lancelot* of Chrétien de Troyes [46] introduces the hero by way of a single adventure, when he rescued queen Guinevere from her abduction by a villain, whereas the *Lancelot-Grail* [47] largely *expands* this early text, so as to cover his life since from his birth, even explaining that his full name "Lancelot of the Lake" came from his passing his childhood below the waters of a lake – actually a fairy-produced visual illusion, since a knight breathing through gills would have no place in Arthurian stories. On the opposite direction, the fairy tale of the *Beauty and the Beast*, first recorded (1740) by Gabrielle-Suzanne Barbot de Villeneuve, was later radically *summarized* (1756), to become its most popular version by Jeanne-Marie Leprince de Beaumont [48].

Still about meronymic relations, one must realize that hierarchy in fictional genres may work at more than two levels. Besides types, Aarne-Thomson's *Index* also catalogues *motifs* [18][49], which are short sequences of events recurring in a vast number of types. Of special interest to us is the case where a pivotal scene is located in one story, and then extracted and *expanded* to full story dimension. As an example, we may refer to the *Volsunga Saga* (*Saga of the Völsungs*) [50], a late 13th-century poetic rendition in Old Norse of the origin and decline of the Völsung clan, in which the dwarf Andvari curses a ring called Andvaranaut ("Andvari's gift"), warning that it would bring death to anyone who owned it. This cursed ring motif is reputed to be the primary source of Tolkien's richly developed *The Lord of the Rings* [51].

3.1.4 The Antithetic Relation

Two narratives may be in *antithetic relation* if they somehow *reverse* each other. They may involve mutually incompatible event sequences or may assume contradictory factual situations or opposite objectives. They also serve to express a different viewpoint when judging the story's moral or political implications. To find an example, we may refer to the very popular folktale broadly known as *Little Red Riding Hood*. In Aarne-Thompson's *Index*, the story is classified under type AT 333, characteristically named "The Glutton". In the canonical versions, the character impersonating the "glutton" is a wolf or other monster which devours human beings until all of them are rescued alive from his belly. A striking transgressive reversal is the *Uncle Wolf* story, collected by Italo Calvino [52], in which the girl is the first to reveal herself as "glutton". She does not resist the temptation to eat and drink all that her mother was sending to Uncle Wolf in return for the loan of a skillet, offering him instead an ugly mess composed of donkey manure, dirty water, and lime. He is not deceived and threatens her: "Tonight I'm coming to eat you!" He sneaks into the house, repeatedly announcing where he is at each moment until reaching the girl's room and eating her, in a frightening ghost-like sequence.

3.2 Semiotic Operations for Narrative Generation

Corresponding to each of the four semiotic relations, four *semiotic operations* were derived with the intent of generating new stories from the plot of existing books. Letting B1 and B2 be books and S a scene from book B1, the operations are:

- Syntagmatic relation: **combine** B1 and B2

- Paradigmatic relation: **imitate** B1
- Meronymic relation: **expand** S that is part of B1
- Antithetic relation: **reverse** B1

As we prepare to describe our proposed approach, as adopted in the prototype that is the object in the next section, whereby existing book narratives serve as inspiring sources of new narratives, we should have in mind Barthes's assertion that reusing other authors' work, in honest and imaginative fashion, is a culturally consecrated artistic practice, and that *intertextuality* can be detected everywhere [53]:

> "Any text is a new tissue of past citations. Bits of code, formulae, rhythmic models, fragments of social languages, etc., pass into the text and are redistributed within it, for there is always language before and around the text. Intertextuality, the condition of any text whatsoever, cannot, of course, be reduced to a problem of sources or influences; the intertext is a general field of anonymous formulae whose origin can scarcely ever be located; of unconscious or automatic quotations, given without quotation marks."

4 The ChatGeppetto Prototype

The previous section analyzed the manifestation of semiotic relations in existing narratives, which yielded valuable insights regarding how story writers can effectively reuse narrative ideas to create new stories. The remarkable capabilities of ChatGPT to understand the proposed semiotic operations, together with its extensive knowledge of narrative books, led us to the development of ChatGeppetto, an AI-powered narrative generation system, which is presented in the next subsections.

4.1 The Architecture of ChatGeppetto

The architecture of ChatGeppetto is illustrated in Figure 2. The system comprises two AI agents: (1) the *Storywriter AI Agent*, which is responsible for writing story events and scene descriptions in natural language; and (2) the *Illustrator AI Agent*, which handles the process of visually depicting the narrative events through images (using the natural language descriptions of the scenes generated by the *Storywriter AI Agent*). The AI agents are controlled by the *Plot Manager*, which conducts the plot generation process by requesting story events and illustrations from the *Storywriter AI Agent* and from the *Illustrator AI Agent*, respectively. Users interact with ChatGeppetto through the *User Interface* component, which provides a visual interface that allows users to compose and visualize new narratives in video and storyboard formats. All narratives generated by the system are stored in the *Stories Database*.

ChatGeppetto relies on the knowledge of the AI agents to generate new narratives using the proposed semiotic operations. A plugin approach is used to integrate the AI agents into the system, so as to simplify their substitution when new and more powerful models are released. In the current implementation of ChatGeppetto, the *Storywriter AI Agent* is based on the ChatGPT model (GPT-3.5-turbo), which is accessed through the OpenAI API;[3] and the *Illustrator AI Agent* is based on the Stable Diffusion 2.1 model,[4] which runs on a private server and is accessed through a REST API.

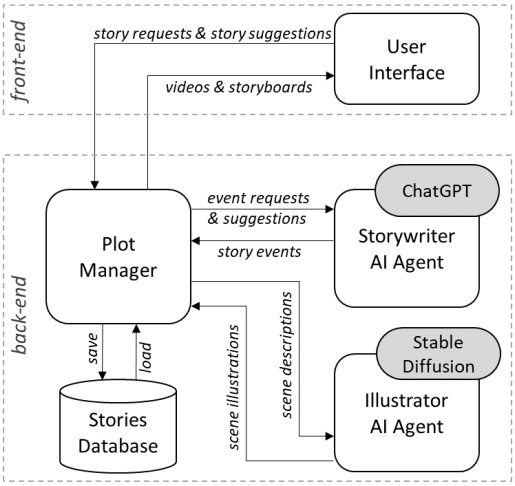

Figure 2: Architecture of ChatGeppetto

The system is implemented as a Web application that can be accessed through any Web browser (desktop or mobile). The *Plot Manager* and the communication interfaces of the *AI Agents* are implemented in PHP and Python (running on the back-end), and the *User Interface* is implemented in HTML and JavaScript (operating as the front-end of the system). The *Stories Database* is implemented as a MySQL database.

4.2 Story Generation

The proposed process to generate new narratives using semiotic operations relies on the extensive book knowledge of ChatGPT to create new stories reusing and adapting the plot of existing books. As introduced in section 3.2, new stories are generated according to a set of parameters that vary for each type of semiotic operation: the *combine* operation requires two *book titles*, B1 and B2 (creating a new story that draws from the plot of both books); the *expand* operation functions over a single *related book* B1 and a *scene description* S (creating a new story that details (zoom-in) the scene S of book B1), and both *imitate* and *reverse* operations require only one *related book* B1 (while the imitate operation creates a new story that is similar to B1, the antithetic operation generates a new story that is the opposite of B1).

Besides the preferred semiotic operation and the selected books and scenes, another important parameter for the story generation process is the *description of the protagonist* (PD). The protagonist (or protagonists) can be characters taken from related books, can be arbitrarily chosen names, or can even be recruited from totally extraneous contexts – noting as a classic example of such deliberate contrasts the book *A Connecticut Yankee in King*

[3] https://platform.openai.com/docs/api-reference
[4] https://github.com/Stability-AI/stablediffusion

Arthur's Court [54], an 1889 novel by American writer Mark Twain. In addition to the protagonists' names, this parameter can also be used to indicate preferences related to the protagonists' characteristics (e.g., physical attributes, emotions, desires) or their current situation (e.g., where they are, what they are doing, what happened before).

The aforementioned parameters are used by ChatGeppetto to compose a set of prompts to instruct the AI agent responsible for writing the story (i.e., the Storywriter AI Agent, which is currently based one the ChatGPT model). These prompts were designed based on initial experiments conducted with ChatGPT to test its capacity to identify semiotic relations in the plot of existing narratives. Each prompt has three parts: (1) the definition of a semiotic relation; (2) a task description; and (3) a description of the output format. Table 2 presents the definitions and task descriptions created for each semiotic operation.

The description of the output format was designed to instruct ChatGPT to write the story in a format that can be interpreted by ChatGeppetto, which contains a description of the story event (*event description*), a description of an image that illustrates the scene (*scene description*), and the title of the new story (*story title*). The exact instructions provided to the Storywriter AI Agent to produce stories in the desired output format are: "*Always write just one line for the story, starting with 'EVENT:', followed by a second line starting with 'IMAGE:' containing a short description of an image that illustrates the narrative event. Also generate a new creative title for book B and add it at the beginning of the response starting with 'TITLE:'.*".

A key feature of the initial prompt is the fact that it instructs ChatGPT to write only the first event of the story, which leaves ChatGeppetto free to produce narratives in a stepwise manner. This approach allows users to compose stories in an interactive manner, allowing them to suggest new events for the story, decide when to regenerate certain events, when to continue the narrative, and when to finish the story.

The *continue* operation confers to users the ability to creatively interact with the Storywriter AI Agent by entering suggestions for the story in natural language. These suggestions are optional, which allows users to decide whether they choose to influence in the event composition process, or to simply let the AI Agent proceed without interference. ChatGeppetto uses a specific prompt to instruct the AI Agent to generate the next event for the story: "*Continue the story by generating another pair of lines (only 'EVENT' and 'IMAGE:').*". Whenever a suggestion is provided by the user, the continue prompt is complemented by the following instruction: "*Consider the following suggestions for the story:* " + user suggestions.

An essential feature of ChatGPT, extensively explored by our system, is its capacity to operate in a conversational way, allowing the AI agent to recognize the context and instructions provided in previous messages. The general structure of the conversation between ChatGeppetto and ChatGPT is:

1. **ChatGeppetto:** SEMIOTIC RELATION DEFINITION + SEMIOTIC RELATION TASK + OUTPUT FORMAT
2. **ChatGPT:** TITLE: *story title*. EVENT: *event description*. IMAGE: *scene description*.
3. **ChatGeppetto:** CONTINUE PROMPT
4. **ChatGPT:** EVENT: *event description*. IMAGE: *scene description*.

where messages 3 and 4 are repeated until the user decides to end the story (i.e., when the user feels that the story is complete), which is indicated by pressing a "Finish and Save" button.

Besides continuing the story, users can also request the regeneration of the last story event, which allows them to find alternatives that best suit their preferences. The *regenerate* operation consists of removing the last event from story plot and requesting a new event to the Storywriter AI Agent. Thanks to the stochastic nature of the ChatGPT model, an alternative – and sometimes very different – version of the story event will be generated and added to the plot.

4.3 Image and Video Generation

The process of generating a visual representation for the narrative events relies on the recent advancements on text-to-image machine learning models, such as DALL-E, Midjourney, and Stable Diffusion, which are capable of producing detailed images based on natural language descriptions. In our implementation, we use the Stable Diffusion 2.1 model [55].

The scene descriptions generated by the Storywriter AI Agent play a fundamental role in the image generation process. Providing these textual descriptions directly as input to the Stable Diffusion model would be enough to produce interesting illustrations for the narrative events. However, to provide users with more control over the generated images, we included an additional parameter in ChatGeppetto that allows users to describe the *style of the illustrations* (IS). This parameter is directly combined with the scene description generated by the Storywriter AI Agent to establish the prompt that is provided to the Illustrator AI Agent: SCENE DESCRIPTION + IS. The default illustration style used by ChatGeppetto is "*digital painting, hyperrealistic, highly detailed, sharp focus, stunningly beautiful, cinematic lighting*", which we found to be good for producing cinematic-looking images. However, users are free to input any style description, including styles of famous artists (e.g.: "*painted by Pablo Picasso*") and art styles (e.g.: "*cubism painting*").

The basic prompt is complemented by a general negative prompt that identifies the visual elements of that the Stable Diffusion model should avoid during image generation. This negative prompt was designed considering the weaknesses we observed in the images generated for general narrative events. The negative prompt used by ChatGeppetto is: "*((poorly drawn hands)), ((poorly drawn face)), (((deformed))), blurry, ((bad anatomy)), (((bad proportions))), ((extra limbs)), cloned face, (((ugly))), (((duplicate))), ((mutilated)), extra fingers, (too many fingers), mutated hands, (((disfigured))), out of frame, extra limbs, gross proportions, (malformed limbs), ((missing arms)), ((missing legs)), (((extra arms))), (((extra legs))), (fused fingers), (((long neck)))*", where parentheses are used to increase the emphasis of the inner elements (more parentheses represent more emphasis).

Table 2. Parameterized prompts designed to instruct the Storywriter AI Agent on how to write new stories using semiotic operations and the plot of existing books. The parameters are: <B1> - related book 1; <B2> - related book 2; <S> - scene description; and <PD> - protagonist's description.

Semiotic Operation		Prompt
Combine	Definition	The fictional protagonist in the plot of book A is X. The fictional protagonist in the plot of book C is Y. Book A has a syntagmatic relation with book C if X does not appear in C, Y does not appear in A, and there exists a book B that features both X and Y.
	Task	Considering this definition, write the first event for a new book B considering that book A is "<B1>" and book C is "<B2>". The protagonist of book B must be <PD>.
Imitate	Definition	Book A has a paradigmatic relation with book B if there is some similarity between their plots and their protagonists have similar objectives.
	Task	Considering this definition, write the first event for a new book B considering that book A is "<B1>" (book B must have a paradigmatic relation with book A). The protagonist of book B must be <PD>.
Expand	Definition	There is a meronymic relation between two books A and B if the plot of B is equivalent to a detailed narrative of a scene S taken from the plot of A.
	Task	Letting A be the book "<B1>", S be a scene taken from the plot of A where <S>, and letting B be an imaginary new book, write the first event for the plot of B, so that there is a meronymic relation between A and B. The protagonist of book B must be <PD>.
Reverse	Definition	Book A has an antithetic relation with book B if there is some similarity between their plots but the objective of the protagonist of book A is the complete opposite of the objective of the protagonist of book B.
	Task	Considering this definition, write the first event for a new book B considering that book A is "<B1>" (book B must have an antithetic relation with book A). The protagonist of book B must be <PD>.

In addition to the negative prompt, we also control the seed used by the Stable Diffusion. The *seed* is a parameter used to initialize the model, which can be used to produce similar images (the same seed and the same prompt always produce the same output image). Small variations in the prompt with the same seed also tend to produce similar images. Therefore, the Illustrator AI Agent uses the same seed for all images generated for the same story, which improves the coherence between the images generated for the different events of the same narrative.

Besides image generation, which allows ChatGeppetto to create storyboards by combining text and images, the system also incorporates a video generation functionality that blends image, text, and audio narration to create an audiovisual rendition for the generated stories. Once a user completes the composition of a story plot, the video generation process is performed using the plot events as input (event descriptions and image illustrations).

Given as input a story plot $S=\{e_1, e_2, \ldots, e_n\}$, where each event is a pair $e_i = (text_i, image_i)$ containing the textual event description of the event ($text_i$) and the image that illustrates the event ($image_i$), the process of creating a video rendition for the story comprises the following steps:

1. For each event $e_i \in S$:
   a. Create a new image IMG_i by drawing $text_i \in e_i$ into $image_i \in e_i$;
   b. Create a new audio AUD_i by narrating $text_i \in e_i$ using a text-to-speech model;
   c. Create a new video VID by combining IMG_i and AUD_i, and add it to set V;
2. Merge all videos of V and save the resulting video file.

In our current implementation, the process of blending text and image is implemented using the Python Imaging Library (Pillow),[5] the process of generating audio narration is performed through the Coqui TTS Library,[6] and the process of combining image and audio, as well as the process of merging video files, is performed using the FFmpeg tool.[7]

4.4 User Interface

As previously mentioned, ChatGeppetto is a Web application that can be accessed by any Web browser through our public Website: http://www.icad.puc-rio.br/~logtell/chatgeppetto/.

As illustrated in Figure 3, the user interface of ChatGeppetto allows users to compose new stories by selecting a semiotic operation (action 1 in Figure 3 (a)) and providing the required parameters (action 2 in Figure 3 (a)). The story generation process starts when the user presses the button "Generate Story" (action 3 in Figure 3 (a)), which leads to the story composition screen (Figure 3 (b)), wherein the user can see the title and the first event of the story with a scene description and a visual illustration.

[5] https://pypi.org/project/Pillow/
[6] https://github.com/coqui-ai/TTS
[7] https://ffmpeg.org/

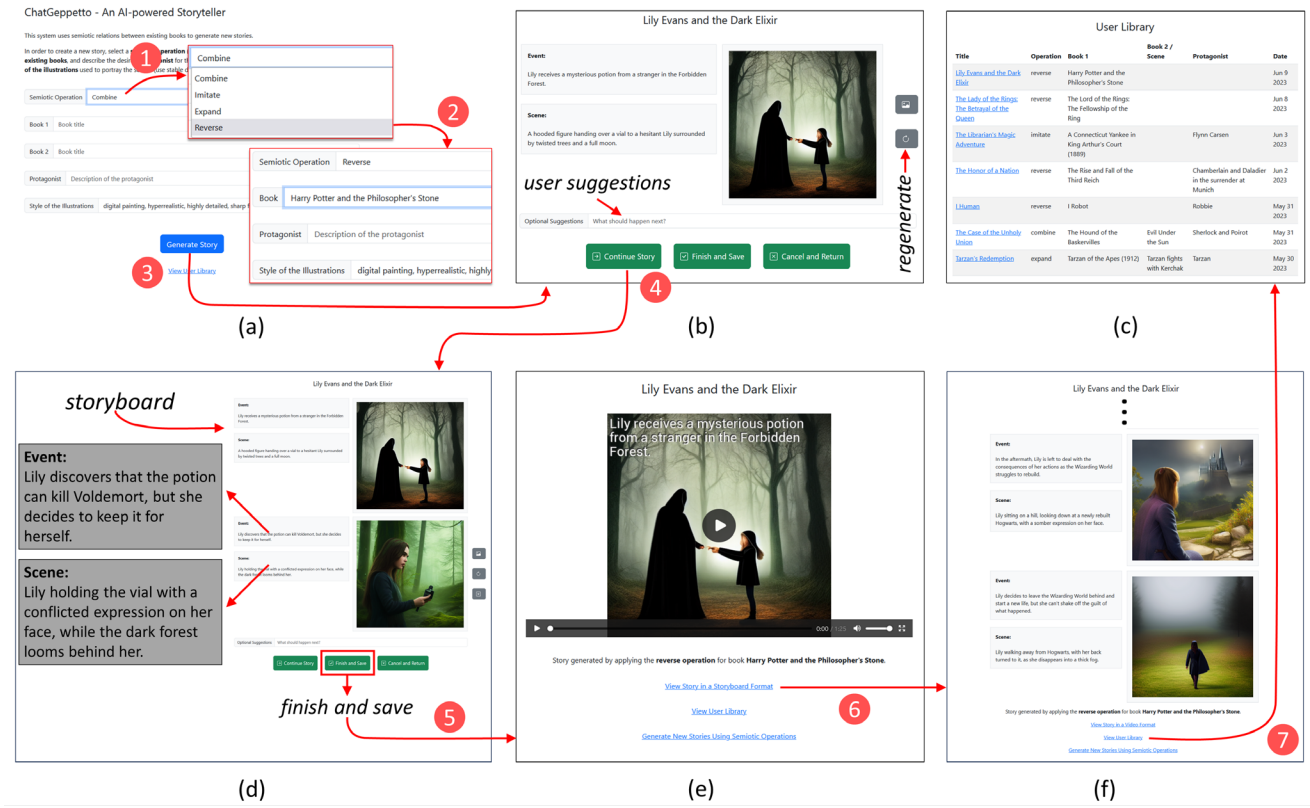

Figure 3: User interface of ChatGeppetto

In the example illustrated in Figure 3, the user selected the reverse operation, entered the title of the book "Harry Potter and the Philosopher's Stone", and ChatGeppetto generated a story called "Lily Evans and the Dark Elixir". After visualizing the first event, the user has four options: (1) regenerate the event by pressing the button on the bottom-right side of the scene illustration (button with a reload icon), which will cause the system to produce another version of the event; (2) regenerate only the visual illustration of the scene by pressing the button on the top-right side of scene illustration (button with an image icon), which will cause the system to generate an alternative visual illustration for the event (using a random seed for the Stable Diffusion model); (3) continue the story by pressing the button "Continue Story" (action 4 in Figure 3 (b)), which will lead the system to generate the next event for the story in a storyboard format (Figure 3 (d)); or (4) provide suggestions for the next story events in the field "Optional Suggestions", and then continue the story or regenerate the event, in both cases causing the system to consider the user's suggestions in the generation of the next event.

When the user decides to conclude the story, the button "Finish and Save" must be pressed (action 5 in Figure 3 (d)). With the story completed, ChatGeppetto proceeds with the generation of the video for the story. When the video is ready, the resulting narrative is presented to the user (Figure 3 (e)), who also has the option to visualize the full narrative in the storyboard format (Figure 3 (f)) by clicking in the link "View Story in Storyboard Format" (action 6 in Figure 3 (e)).

All saved narratives are added to the current user's library (Figure 3 (c)), wherein a few representative examples created by us are also kept. The user library can be accessed from the initial screen of our system (Figure 3 (a)) or from the story visualization screens (Figure 3 (e) and (f)) by clicking in the link "View User Library" (action 7 in Figure 3 (f)). User's stories are kept private in the user's library by associating them to a unique user key, which is automatically generated when a user accesses the system for the first time. The user key is stored locally in the user's web browser through a Web cookie.

Examples of stories generated by ChatGeppetto are available at: http://www.icad.puc-rio.br/~logtell/chatgeppetto/list.html. In particular, a video narrating the story "Lily Evans and the Dark Elixir", whose opening screen is shown in Figure 3, can be accessed at: http://www.icad.puc-rio.br/~logtell/chatgeppetto/view.html?id=170.

5 Concluding Remarks

What we consider to be the main contribution of the present paper is the proposal and demonstration of a novel strategy for assisting

different types of users, not necessarily professional writers, in a highly interactive story composition process. As explained in the preceding sections, based on our semiotic relations theory, the strategy comprises three components: a software tool, a user interface, and two AI-agents.

This third component, currently driven by ChatGPT, played from the start a fundamental role. The verbal definitions determining how semiotic relations instruct the prototype to react creatively to the user's prompts were first analyzed through a number of guiding experiments conducted in consultation with ChatGPT. To users, an AI agent should be regarded as a welcome collaborator, like the traditional anonymous ghostwriter, invited to enhance literary writing. No longer limited to multiple-choice decisions at branching points, they are free to issue instructions to the AI agents regarding the composition of each new event, with the power to demand alternatives for the generated events, as well as for their pictorial representation. While some users may prefer to just let the story unroll, more knowledgeable users will likely enjoy the opportunity to contribute creatively as co-authors. At first, we feared that so much freedom might compromise consistency, but then we noted how ChatGPT used its context awareness and imbued common sense to achieve plausibility.

One example should suffice to illustrate this point. We applied the reverse semiotic operation to William Shirer's book *The Rise and Fall of the Third Reich* [56], and in the "Protagonist" field we specified: "Chamberlain and Daladier in the surrender at Munich". [8] The first event generated was "Chamberlain and Daladier come to an agreement with Hitler in Munich, agreeing to cede the Sudetenland to Nazi Germany in exchange for 'peace'." Then we gave the instruction that, in the next event, Chamberlain should repent, since we expected to make him cancel the agreement. However, the ensuing events were not so trite. First, some reason for Chamberlain's misgivings should be made explicit, which was thus expressed by the AI agent: "The annexation of Czechoslovakia by Nazi Germany prompts Chamberlain to realize his grave mistake". So, the agreement had already produced a nasty effect, which could not be said to be magically undone by formal cancellation – and so cancelling the agreement should be interpreted as a change of attitude: "Chamberlain, standing in front of the House of Commons, announces the cancellation of the Munich Agreement, declaring that Great Britain will now stand up to Nazi aggression." And when we next instructed the prototype to exhort the German generals to depose Hitler, the AI agent cited, by its own initiative, the name of a famous general who, in actual world history, would later revolt and be compelled to commit suicide: "A group of German generals, led by Erwin Rommel, successfully deposes Hitler and takes control of the government in a peaceful coup." And – which is more significant – the tool, when allowed to continue, made explicit the consequences of the dictator's downfall: "The generals begin the process of negotiating a peace treaty with the Allied powers to end the war and restore democracy in Germany." As an aside, we must remark that this test suggests that not only fictional, but also serious narratives might be usefully treated, in this case reproducing a counterfactual "what if" speculation that was already present in Shirer's book.

To ourselves, while specifying our definitions in English sentences, the AI agent gave us the unique opportunity to start learning how to program in natural language, given that until recently we were exclusively used to reason in terms of logic programming, besides employing a variety of general-purpose and graphical programming languages.

The results thus far obtained are encouraging. Our previous tools, using plan generation to compose the plots [17], sometimes surprised us by finding unexpected ways to reach goals that we believed to be impossible. But the multiplicity of unpredicted variants, obtained with the autonomous thinking of the AI agents, goes much beyond what we had before. Moreover, the user interface is informal, and the generated texts are exhibited to the user in idiomatically correct natural language.

Until now, the prototype served mainly as a proof of concept and for gaining experience exploring AI agents' potential more effectively. Further research is needed to measure user satisfaction, considering both professional and non-professional users, thereby providing valuable feedback to extend the functionality and improve the prototype's interface. The set of definitions that, so to speak, parameterize the behavior of the prototype, may be extended and/or generalized to cover an increasing number of cases.

Future research may also be directed toward other uses of systems assisted by AI agents, similar to those adopted in our project. A straightforward application is as a teaching resource to train language proficiency and to develop literary skills. We chose to concentrate on book narratives – known as habitual sources of movie adaptations and originators of games. Indeed, besides helping to create underlying stories for games, an AI agent could play a part, perhaps as an adversary to be defeated, but also as a Merlin-like mentor, offering wise advice to the human players.

ACKNOWLEDGMENTS

We want to thank CNPq (National Council for Scientific and Technological Development) and FINEP (Funding Agency for Studies and Projects), which belong to the Ministry of Science, Technology, and Innovation of Brazil, for the financial support.

REFERENCES

[1] Shannon Liao. A.I. may help design your favorite video game character. The New York Times, May 22, 2023. Retrieved from www.nytimes.com/2023/05/22/arts/blizzard-diffusion-ai-video-games.html

[2] Antonio L. Furtado, Marco A. Casanova, and Simone D. J. Barbosa. 2014. A semiotic approach to conceptual modelling. In Conceptual Modeling - ER 2014, Lecture Notes in Computer Science, vol. 8824, Springer, Cham, 1–12. DOI: https://doi.org/10.1007/978-3-319-12206-9_1.

[3] Edirlei Soares de Lima, Antonio L. Furtado, and Bruno Feijó. 2015. Storytelling Variants: The Case of Little Red Riding Hood. In Entertainment Computing - ICEC 2015, Lecture Notes in Computer Science, vol. 9353, Springer, Cham, 286–300. DOI: https://doi.org/10.1007/978-3-319-24589-8_22.

[4] Edirlei Soares de Lima, Bruno Feijó, Marco A. Casanova, and Antonio L. Furtado. 2016. Storytelling Variants Based on Semiotic Relations. Entertainment Computing 17, 31-44. DOI: https://doi.org/10.1016/j.entcom.2016.08.003.

[5] Angelo E. M. Ciarlini, Marco Antonio Casanova, Antonio L. Furtado, and Paulo A. S. Veloso. 2010. Modeling interactive storytelling genres as application

[8] http://www.icad.puc-rio.br/~logtell/chatgeppetto/view.html?id=104

domains. Journal of Intelligent Information Systems 35, 3, 347–381. DOI: https://doi.org/10.1007/s10844-009-0108-5.

[6] Fabio A. Guilherme da Silva, Antonio L. Furtado, Angelo E. M. Ciarlini, Cesar Tadeu Pozzer, Bruno Feijó, and Edirlei Soares de Lima. 2012. Information-Gathering Events in Story Plots. In Entertainment Computing - ICEC 2012, Lecture Notes in Computer Science, vol. 7522, Springer, Heidelberg, 30-44. DOI: https://doi.org/10.1007/978-3-642-33542-6_3.

[7] Edirlei Soares de Lima, Bruno Feijó, and Antonio L. Furtado. 2014. Hierarchical generation of dynamic and nondeterministic quests in games. In Proceedings of the 11th Conference on Advances in Computer Entertainment Technology (ACE '14). ACM, New York, NY, USA, Article 24, 1–10. DOI: https://doi.org/10.1145/2663806.2663833.

[8] Edirlei Soares de Lima, Bruno Feijó, and Antonio L. Furtado. 2009. Procedural Generation of Quests for Games Using Genetic Algorithms and Automated Planning. In Proceedings of the 18th Brazilian Symposium on Computer Games and Digital Entertainment (SBGames), Rio de Janeiro, Brazil, 2019, 144-153. DOI: https://doi.org/10.1109/SBGames.2019.00028.

[9] Angelo E. M. Ciarlini, Cesar T. Pozzer, Antonio L. Furtado, and Bruno Feijó. 2005. A logic-based tool for interactive generation and dramatization of stories. In Proceedings of the 2005 ACM SIGCHI International Conference on Advances in computer entertainment technology (ACE '05). ACM, New York, 133–140. DOI: https://doi.org/10.1145/1178477.1178495.

[10] Edirlei Soares de Lima, Bruno Feijó, and Antonio L. Furtado. 2018. Video-based Interactive Storytelling Using Real-time Video Compositing Techniques. Multimedia Tools and Applications 77, 2, 2333–2357. DOI: https://doi.org/10.1007/s11042-017-4423-5.

[11] Edirlei Soares de Lima, Cesar T. Pozzer, Marcos C. d'Ornellas, Angelo E. M. Ciarlini, Bruno Feijó, and Antonio L. Furtado. 2009. Virtual cinematography director for interactive storytelling. In Proceedings of the International Conference on Advances in Computer Entertainment Technology (ACE '09). ACM, New York, NY, USA, 263–270. https://doi.org/10.1145/1690388.1690432.

[12] Mieke Bal. 2017. Narratology: Introduction to the Theory of Narrative. University of Toronto Press, Toronto.

[13] Edirlei Soares de Lima, Bruno Feijó, and Antonio L. Furtado. 2023. Managing the plot structure of character-based interactive narratives in games. Entertainment Computing 47, 100590. DOI: https://doi.org/10.1016/j.entcom.2023.100590.

[14] Edirlei Soares de Lima, Bruno Feijó, Simone D.J. Barbosa, Antonio L. Furtado, Angelo E.M. Ciarlini, and Cesar T. Pozzer. 2014. Draw your own story: Paper and pencil interactive storytelling. Entertainment Computing 5, 1, 33-41. DOI: https://doi.org/10.1016/j.entcom.2013.06.004.

[15] Edirlei Soares de Lima, Bruno Feijó, and Antonio L. Furtado. 2020. Adaptive storytelling based on personality and preference modeling. Entertainment Computing 34, 100342. DOI: https://doi.org/10.1016/j.entcom.2020.100342.

[16] Edirlei Soares de Lima, Bruno Feijó, and Antonio L. Furtado. 2018. Player behavior and personality modeling for interactive storytelling in games. Entertainment Computing 28, 32-48. DOI: https://doi.org/10.1016/j.entcom.2018.08.003.

[17] Angelo E. M. Ciarlini, Simone D. J. Barbosa, Marco A. Casanova, and Antonio L. Furtado. 2010. Event relations in plan-based plot composition. Computers in Entertainment 7, 4, Article 55. DOI: https://doi.org/10.1145/1658866.1658874.

[18] Antti Aarne, and Stith Thompson. 1961. The Types of the Folktale. Acad. Scientiarum Fennica.

[19] Jamshid J. Tehrani. 2013. The Phylogeny of Little Red Riding Hood. PLOS ONE 8, 11, e78871. DOI: https://doi.org/10.1371/journal.pone.0078871.

[20] James H. Martin, Daniel Jurafsky. 2023. Speech and Language Processing - An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, Third Edition Draft.

[21] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774. DOI: https://doi.org/10.48550/arXiv.2303.08774.

[22] Carlo Collodi. 1998. The Adventures of Pinocchio, Story of a Puppet. Zoland.

[23] Janet L. Kolodner. 2014. Case-Based Reasoning. Morgan Kaufmann.

[24] Pablo Gervás, Belén Díaz-Agudo, Federico Peinado & Raquel Hervás. 2004. Story Plot Generation based on CBR. In Applications and Innovations in Intelligent Systems XII, SGAI 2004. Springer, London, 33–46. DOI: https://doi.org/10.1007/1-84628-103-2_3.

[25] Gilles Fauconnier, and Mark B. Turner. 1994. Conceptual projection and middle spaces. Technical Report 9401, University of California, San Diego. DOI: https://doi.org/10.2139/ssrn.1290862.

[26] Edirlei Soares de Lima, Bruno Feijó, and Antonio L. Furtado. 2021. Computational Narrative Blending Based on Planning. In Entertainment Computing – ICEC 2021, Lecture Notes in Computer Science, vol. 13056, Springer, Cham, 289–303. DOI: https://doi.org/10.1007/978-3-030-89394-1_22.

[27] Susanna Värtinen, Perttu Hämäläinen, and Christian Guckelsberger. 2022. Generating Role-Playing Game Quests With GPT Language Models. IEEE Transactions on Games, December 2022, 1-12. DOI: https://doi.org/10.1109/TG.2022.3228480.

[28] Angela Fan, Mike Lewis, and Yann Dauphin. 2018. Hierarchical Neural Story Generation. In Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, 889–898, Melbourne, Australia. Association for Computational Linguistics. DOI: https://doi.org/10.18653/v1/P18-1082.

[29] Amal Alabdulkarim, Winston Li, Lara J. Martin, and Mark O. Riedl. 2021. Goal-Directed Story Generation: Augmenting Generative Language Models with Reinforcement Learning. In arXiv:2112.08593. DOI: https://doi.org/10.48550/arXiv.2112.08593.

[30] Peng Xu, Mostofa Patwary, Mohammad Shoeybi, Raul Puri, Pascale Fung, Anima Anandkumar, and Bryan Catanzaro. 2020. MEGATRON-CNTRL: Controllable Story Generation with External Knowledge Using Large-Scale Language Models. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2831–2845. DOI: https://doi.org/10.18653/v1/2020.emnlp-main.226.

[31] Su Wang, Greg Durrett, and Katrin Erk. 2020. Narrative Interpolation for Generating and Understanding Stories. In arXiv:2008.07466. DOI: https://doi.org/10.48550/arXiv.2008.07466.

[32] Louis Castricato, Spencer Frazier, Jonathan Balloch, and Mark Riedl. 2021. Tell Me A Story Like I'm Five: Story Generation via Question Answering. In Proceedings of the 3rd Workshop on Narrative Understanding. https://par.nsf.gov/biblio/10249509.

[33] Daijin Yang, Yanpeng Zhou, Zhiyuan Zhang, Toby Jia-Jun Li, and Ray LC. 2022. AI as an Active Writer: Interaction strategies with generated text in human-AI collaborative fiction writing. In Joint Proceedings of the ACM IUI Workshops 2022, March 2022, Helsinki, Finland, 56-65.

[34] Jingwen Xiang, Zoie Zhao, Mackie Zhou, Megan McKenzie, Alexis Kilayko, Jamie C. Macbeth, Scott Carter, Katharine Sieck, and Matthew Klenk. 2022. Interleaving a Symbolic Story Generator with a Neural Network-Based Large Language Model. In Proceedings of Tenth Annual Conference on Advances in Cognitive Systems, Arlington, November 19-22, 2022.

[35] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: Story Writing With Large Language Models. In 27th International Conference on Intelligent User Interfaces (IUI '22). ACM, New York, NY, USA, 841–852. DOI: https://doi.org/10.1145/3490099.3511105.

[36] Kevin Yang, Yuandong Tian, Nanyun Peng, and Dan Klein. 2022. Re3: Generating Longer Stories With Recursive Reprompting and Revision. In rXiv:2210.06774. DOI: https://doi.org/10.48550/arXiv.2210.06774.

[37] Eduardo C. Garrido-Merchán, José Luis Arroyo-Barrigüete, and Roberto Gozalo-Brizuela. 2023. Simulating H.P. Lovecraft horror literature with the ChatGPT large language model. In arXiv:2305.03429. DOI: https://doi.org/10.48550/arXiv.2305.03429.

[38] Ferdinand de Saussure. 1998. Course in General Linguistics. Open Court.

[39] Daniel Chandler. 2002. Semiotics: The Basics. Routledge, London.

[40] Kenneth Burke. 1969. A Grammar of Motives. Univ. of California, Oakland.

[41] Jonathan Culler. 1981. The Pursuit of Signs: Semiotics, Literature, Deconstruction. Routledge, Oxfordshire.

[42] Hayden White. 1978. Tropics in Discourse: Essays in Cultural Criticism. Johns Hopkins University Press, Baltimore.

[43] Antonio L. Furtado. 2011. Semiotic Relations and Proof Methods. Monografias em Ciência da Computação n° 18/11. PUC-Rio, Rio de Janeiro.

[44] Roman Jakobson. 1956. Two Aspects of Language and Two Types of Aphasic Disturbances. In Fundamentals of Language. Mouton, The Hague.

[45] Morton E. Winston, Roger Chaffin, and Douglas Herrmann. 1987. A taxonomy of part-whole relations. Cognitive Science 11, 4, 417–444. DOI: https://doi.org/10.1207/s15516709cog1104_2.

[46] Chretien de Troyes. 1991. The Complete Romances of Chretien de Troyes. Indiana University Press, Bloomington.

[47] Norris Lacy. 2010. The Old French Vulgate & Post-Vulgate Cycles in Translation. Routledge, Oxfordshire.

[48] Jack Zipes. 2016. Beauties, Beasts and Enchantment: Classic French Fairy Tales, 2nd edition. Crescent Moon Publishing, Maidstone.

[49] Folgert Karsdorp, Peter van Kranenburg, Theo Meder, Dolf Trieschnigg, and Antal van den Bosch. 2012. In Search of an Appropriate Abstraction Level for Motif Annotations. In: Proceedings of the Third Workshop on Computational Models of Narrative, 22-28.

[50] Jesse Byock. 2000. The Saga of the Volsungs. Penguin Classics, London.

[51] John Ronald Reuel Tolkien. 2020. The Lord of the Rings, 3 vol. Clarion Books.

[52] Italo Calvino. 2002. Italian Folktales. Penguin Books, London.

[53] Roland Barthes. 1981. Theory of the Text. In Robert Young (ed.) Untying the Text: a Post-Structuralist Reader. Routledge & Kegan Paul, Oxfordshire.

[54] Mark Twain. 2018. A Connecticut Yankee in King Arthur's Court. SeaWolf.

[55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 10674–10685. DOI: https://doi.org/10.1109/CVPR52688.2022.01042.

[56] William L. Shirer. 1960. The Rise and Fall of the Third Reich. Simon and Schuster.