Entertainment Computing 52 (2025) 100810

Contents lists available at ScienceDirect

Entertainment Computing

journal homepage: www.elsevier.com/locate/entcom

An AI-powered approach to the semiotic reconstruction of narratives

Edirlei Soares de Lima [a,*], Margot M.E. Neggers [a], Bruno Feijó [b], Marco A. Casanova [b], Antonio L. Furtado [b]

[a] Breda University of Applied Sciences, Academy for AI, Games and Media, Breda, the Netherlands
[b] Pontifical Catholic University of Rio de Janeiro, Department of Informatics, Rio de Janeiro, Brazil

ARTICLE INFO

Keywords:
Computational narratology
Interactive story composition
Deconstruction
Semiotic relations
Semiotic reconstruction
Book narratives
Movie narratives
Storyboards
Artificial intelligence
Intelligent agents

ABSTRACT

This article presents a novel and highly interactive process to generate natural language narratives based on our ongoing work on semiotic relations, providing four criteria for composing new narratives from existing stories. The wide applicability of this semiotic reconstruction process is suggested by a reputed literary scholar's deconstructive claim that new narratives can often be shown to be a tissue of previous narratives. Along, respectively, three semiotic axes – syntagmatic, paradigmatic, and meronymic – existing stories can yield new stories by the combination, imitation, or expansion of an iconic scene; lastly, a new story may emerge through reversal via an antithetic consideration, i.e., through the adoption of opposite values. Targeting casual users, we present a fully operational prototype with a simple and user-friendly interface that incorporates an AI agent, namely ChatGPT. The prototype, in a coauthor capacity, generates context-compatible sequences of events in storyboard format using backward-chaining abductive reasoning (employing Stable Diffusion to draw scene illustrations), conforming as much as possible to the user's authorial instructions. The extensive repertoire of book and movie summaries available to the AI agent obviates the need to manually supply laborious and error-prone context specifications. A user study was conducted to evaluate user experience and satisfaction with the generated narratives. The preliminary findings suggest that our approach has the potential to enhance story quality while offering a positive user experience.

1. Introduction

In this article, we describe a novel and highly interactive process to generate natural language narratives, in storyboard format, on the basis of our ongoing work on *semiotic relations* [1–4]. The starting point of our present effort to offer a friendly interface to casual users (without any previous knowledge requirement) is an assertion of the French literary scholar Roland Barthes [5], significantly referenced in a study of *deconstruction* [6]:

> "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."

This seemingly radical assertion should not be understood as a claim that there is no true originality and that new texts are essentially limited to repeating previous texts. If this were so, the backward search for sources would incur infinite recursion. The truly far-reaching claim is that, alongside their varying degrees of originality, human authors share a rich universal heritage that deserves attention as a fundamental factor in story composition. This heritage is sometimes subtle and not easily identified through deconstructive analysis.

Our proposed semiotic relations, derived from the so-called rhetorical four *master tropes* [7], provide, as suitable *semiotic deconstruction* criteria, four ways of composing new narratives from existing narratives. Along what semioticians call the horizontal syntagmatic axis, one can form a 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

* Corresponding author.
E-mail address: soaresdelima.e@buas.nl (E.S. de Lima).

https://doi.org/10.1016/j.entcom.2024.100810
Received 17 March 2024; Received in revised form 22 June 2024; Accepted 29 June 2024
Available online 4 July 2024

than adding a dimension, aims at some form of *reversal*, through the adoption of opposite values.

Our semiotic approach to interactive story composition is modeled after the seminal theoretical works of Ferdinand de Saussure [8] and Charles Sanders Peirce [9]. It applies the same guidelines employed in semiotic deconstruction for building a prototype that can perform a process we propose to call *semiotic narrative reconstruction*. This process, operated with the help of AI agents – namely ChatGPT and Stable Diffusion – works in a way that is akin to *case-based reasoning* [10], with access to an extensive library of existing book and movie narratives that serve to configure the current involved context, thus dispensing the labor-intensive need to formally supply meta-data specifications for each different chosen domain. Moreover, all user-prototype communication is in English, in a precisely idiomatic style on the part of the AI agent, which in turn remains fairly tolerant of errors and vague or clumsy utterances from the user.

Thanks to the flexible employment of natural language, it became possible to compose narratives in the form of storyboards. Our early prototype tools, developed in logic programming notation, basically remained at what Mieke Bal [11] 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). Arguably, storyboards are situated at the intermediate story layer, potentially more engaging than our early bare plots consisting of sequences of events in abstract notation.

Today, the relationship between human agents (including writers, readers, and computer users) and AI-powered prototypes is a matter of deep interest. By the late nineteen-sixties, running the risk of incurring subjectivity, literary critics such as Wolfgang Iser [12] turned their attention to the reader, affirming that it was in each reader's mind that a narrative was completed. Conceivably, they imagined 'model readers', as not inferior to writers in world knowledge and reasoning ability.

The philosopher, novelist, and semiotician,[1] Umberto Eco, sensitive to this liberal orientation, talked in special of 'open works' and 'works in movement', i.e. works that deliberately left decisions on the meaning of certain passages to the reader's care. Still mindful of the authors' privileges,[2] he explained how authorship could be shared between author and reader, the latter assuming the role of performer [15].

Thus, to sum up, we can say that Eco's "work in movement" is the possibility of numerous different personal interventions, but it is not an amorphous invitation to indiscriminate participation. It just aims at offering the performer the opportunity for an oriented insertion into something that always remains in the worldly context intended by the author.

Eco's open work agenda, barely viable in written texts, suggests what began to be feasible in the present day through *computer-based interactive storytelling*, games included, wherein the author-performer partnership takes the form of a permission to choose among alternatives whenever the storyline reaches a branching point. The final resulting plots should already, to some extent, deserve to be perceived by the user – in the performer role – as a personal new creation.

Our approach goes much further in empowering the user to issue instructions about the events being generated and their visual rendition, which the AI agents try to follow within the constraints imposed by the context and by the commonsense reasoning imbued in their training.

And still, ironically, the user's dependence is subverted when the artificially knowledgeable prototype has to rely on the user's help in conflicting situations, as happens with case-based reasoning's appeal to creative adaptations when trying to apply the solution adopted in a library case to the similar, but not identical, new problem being examined. The general topic of adaptation has been intensively studied under the label of *blending* [16] (Subsection 3.1.2 contains a typical literary example involving the combination of two narratives).

All in all, we believe that originality is and will remain a prerogative of humankind, and that prototypes like ours or, arguably, like any others with similar or superior performance, will not replace human professionals, so that the nightmarish robot dominance in [17] will – hopefully – remain in the science-fiction realm. Recognizing these limitations, and the fact that our system is positioned to assist non-professional authors, its integration into professional content creation workflows (such as in book and film production, and game development) is bound to remain outside the present scope of our work.

This work is an extension of our paper "ChatGeppetto - an AI-powered Storyteller", where we proposed our original semiotic model for interactive storytelling empowered by AI agents [18]. The paper introduced a prototype working on the context of book narratives, which complemented an earlier version that explored movie narratives in a preliminary work [19].

In the present work, the original model is reformulated in the ampler terms of semiotic reconstruction, stemming from the implicitly deconstructive claim of Roland Barthes, and the structural dependence on semiotic relations is better elucidated and is interrelated with an abductive interpretation of event enactment. In turn, the functionality of the prototype, covering both movie and book narratives, has been considerably expanded. A variety of interaction commands is now offered to users, who can now request the reformulation of each generated event, as well as of the images produced to illustrate the scenes, in both cases with the permission to issue instructions and prompts to guide the narrative generation process according to user preferences. In addition to the use of existing book and movie plots, the user-produced stories can now serve as the basis for the composition of new stories, which allows, in particular, to combine more than two stories, as also to initiate serial franchise projects. Each composed story is kept in two formats: simply as a conventional storyboard that can be printed or saved as a PDF file, and as a video displaying and reading aloud the text of the events and projecting the corresponding scene images. Another very useful narrative is kept by the system, namely a complete record of the composition steps, indicating the user's positive or negative reactions towards the events and scenes and whatever guiding or corrective prompts have been manually inserted.

As another important extension of our previous work, we conducted a user study to evaluate how well the prototype met our fundamental goal of assisting non-professional authors to reuse elements from well-known stories to create their own compositions. As part of the experiment, applying rigorous widely-used metrics, such as the Perceived Story Quality Index proposed by Baron and Bluck [91] and the User Experience Questionnaire proposed by Laugwitz et al. [90], we evaluated the overall user experience, the user satisfaction with the generated narratives, and the impact of semiotic relations in the user experience and satisfaction by comparing the full version of our system with a simplified version that generates stories based on user-provided story premises instead of semiotic operations over existing stories.

The rest of the article is organized as follows. Section 2 refers to related works. Section 3 explains in an informal style, with the help of examples, how semiotic relations can support the process of deconstructing narratives, thus providing structuring criteria for the reconstruction process, and discusses our abductive view of event enactment. Section 4 describes the design and implementation of the narrative reconstruction prototype. Section 5 gives the result of user experiments, concentrating on the users' evaluation of the interface's adequacy and the level of their satisfaction with the narratives they generated. Section 6 contains concluding remarks.

[1] Like we endeavored to do here, Umberto Eco applied semiotic techniques to the analysis of movies [13].

[2] However, in spite of this firm defense of authorship, he is in full agreement with Barthes about the intertextual liberties in the authors' community [14]: "books are made only from other books and around other books.".

2. Related work

Our approach to support narrative generation is akin to case-based reasoning [10], accessing some "virtual library" of fairly known book or movie narratives to search for items whose structure would meet certain characteristics, and could then be adapted in order to compose new narrative structures. Previous attempts to employ case-based reasoning in narrative generation include, among others, the work reported by Pablo Gervás [20,21]. 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 [16,22], calling for human interaction, ideally combined with the development of domain-oriented heuristics.

The use of Large Language Models (LLMs) for narrative generation is a recent research topic. While some works focus on the use of fine-tuning and training techniques to improve narrative quality [23–25], others explore new strategies to handle the narrative generation process [26–30]. There are also works that focus on the evaluation of the writing capabilities of LLMs, such as Yuan et al. [31], which analyzes the potentials of the GPT-3 model in the collaborative co-writing process. Another example is Yang et al. [32], which presents the results of a study comparing several aspects of narratives generated by ChatGPT with narratives authored by humans. A similar comparison is also presented in the work by Chu and Liu [33].

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. [34], which demonstrates the effectiveness of the GPT-4 model in reproducing H.P. Lovecraft's writing style. The use of LLMs as evaluators for story quality was also recently explored in the work by Gursesli et al. [35], where the authors evaluated the use of ChatGPT to analyze narrative quality considering coherence, inspiration, readability, word complexity, and narrative fluency.

In the context of fine-tuning LLMs for narrative generation, Värtinen et al. [23] presented a fine-tuned GPT-2 model to generate quest descriptions for roleplaying games. However, a user study revealed that only one in five quest descriptions was considered acceptable by human critics. A different approach is explored by Fan et al. [24], which uses a large dataset of 300 K 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. [25], 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 have been proposed in recent years. In this context, Wang et al. [27] 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. [28], 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 approach is explored in Yang et al.[29], 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. A similar solution is explored by Xie and Riedl [36], who employ LLMs to generate suspenseful stories using an iterative-prompting-based planning approach. They compared their method with a baseline consisting of directly asking ChatGPT for a suspenseful story with straightforward prompts. The results indicate that their approach is more capable of producing suspenseful stories than the baseline method.

The use of external knowledge to enhance the controllability of narrative generation is explored by Xu et al. [26], 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 by Xiang et al. [30], 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.

Many recent works also explored the use of LLMs for narrative generation in the context of games. Kumaran et al. [37] presented a narrative scene generation framework that automates the generation of non-player character (NPC) interactions for games. Their system employs a large language model to generate game scenes and branching conversation paths aligned with authorial intents, which are specified in terms of scene objectives, NPC traits, and location. The use of ChatGPT to generate branching narratives is also discussed by Schetinger et al. [38]. A similar approach is explored by Ashby et al. [39], who use a large language model to create fluent and unique accompanying dialogues for NPCs. The use of large language models for the generation of quests for games was also recently explored in the work by Al-Nassar et al. [40].

The generation of content to support narratives in games has also been explored in recent publications, such as the work by Johnson-Bey et al. [41], which uses ChatGPT to generate theme-relevant content for a character-driven settlement simulation game. Given a textual description of the narrative setting, their system uses ChatGPT to generate Python code for the simulation of characters' businesses, occupations, and traits. Another approach is explored by Wang and Gordon [42], who developed and evaluated a web game that uses GPT-3.5 to create stories based on a set of initial story elements. Their experiments show that GPT-3.5 performs very well in generating narratives based on random sets of initial story elements, even incorporating suspense and flashbacks into its stories. Also, in the context of game narratives, Taveekitworachai et al. [43] investigated the biases present in large language models when utilized for the generation of stories. The results of their study reveal a notable bias towards positive-ending stories in GPT-3.5, GPT-4, and Llama 2.

Dealing specifically with films, related works with stimulating results include a growing number of publications pertaining to the active research topic of film semiotics, a particularly influential contribution being Umberto Eco's characterization of what he calls 'cinematic code' [44]. Another essential reference to the semiology of the cinema is the work by Christian Metz [45]. The relevance of iconic scenes that function as leading motifs is stressed by Folgert Karsdorp [46], who sees them as "the primary building blocks of stories".

To provide a clearer understanding of the various approaches to narrative generation using LLM, we categorized the existing works based on their underlying strategies, as summarized in Table 1. Our approach, in contrast to previous works, utilizes semiotic relations to generate new narratives through the combination, imitation, expansion, and reversal of existing stories.

3. Narrative semiotics

When dealing with narrative semiotics, one must take into consideration Jonathan Culler's remark that a seemingly contradictory *double logic* [47] is involved, one concerning "its plot as a sequence of events which is prior to and independent of the given perspective", and the other, to the contrary, requiring that these events be justified by "their appropriateness to a thematic structure." He further elaborates: "We adopt the second perspective when we discuss the appropriateness or inappropriateness of an ending (when we debate whether these actions are appropriate expressions of the thematic structure which ought to determine them).".

To conciliate the two logics, we start from the second perspective, preliminarily defining the *thematic structure,* and then, in order to cope with the first perspective, proceed, by some sort of backward reasoning,

Table 1
Categorization of the common strategies for narrative generation using large language models.

Strategy	Study	Key Aspect
Fine-tuning and Training	Värtinen et al. [23]	Fine-tuning an LLM for generating roleplaying game quests.
	Fan et al. [24]	Training convolutional seq2seq model with a large dataset for story generation.
	Alabdulkarim et al. [25]	Fine-tuning LLMs using reinforcement learning for story continuation.
Structured Planning and Prompting	Wang et al. [27]	Interpolation techniques with LLMs to generate stories with target endings.
	Castricato et al. [28]	Question-answering approach for backward story generation using LLMs.
	Yang et al. [29]	Recursive reprompting LLMs with a structured story plan.
	Xie and Riedl [36]	Using LLMs and iterative-prompting-based planning for suspenseful story generation.
External Knowledge Integration	Xu et al. [26]	Controllable story generation with external knowledge and LLMs.
	Xiang et al. [30]	Combining LLM with a symbolic story generation system.
NPC and Scene Generation for Games	Kumaran et al. [37]	Using LLMs for narrative generation for games with branching conversation.
	Ashby et al. [39]	Generation of fluent and unique non-player character dialogues with LLMs.
Quest and Content Generation for Games	Al-Nassar et al. [40]	Combining procedural content generation with quest generation using LLMs.
	Johnson-Bey et al. [41]	Theme-relevant content generation for character-driven simulation games using LLMs.
	Wang and Gordon [42]	Story generation for games using a LLM and a set of initial story elements.
Evaluation and Bias Analysis	Garrido-Merchán et al. [34]	Simulating H.P. Lovecraft's writing style with an LLM.
	Gursesli et al. [35]	Using LLMs to evaluate story quality.
	Taveekitworachai et al. [43]	Investigating LLMs' biases in story generation.

to *enact* one or more alternative plots as sequences of events that are capable of leading to endings compatible with the context established by the structure. Intermediate target states to be reached by each appropriate event are also processed in this reasoning style, which opens the possibility of user participation in choosing alternatives and issuing prompts.

In consonance with this double logic conciliatory understanding, our approach separates narrative semiotics into semiotic *structuring* and semiotic *enactment*, following, respectively, the guidance of Ferdinand de Saussure and of Charles Sanders Peirce, the two highly reputed initiators of Semiotics, whose models are briefly exposed and compared in an introductory survey [48]. In both subsections we strive to keep the presentation as informal as possible, concentrating on the analytical deconstruction aspect, in preparation for the more formal description of the prototype whereby our AI-powered reconstruction methods are put to work.

3.1. Semiotic structuring

The notion of *semiotic relations*, as initially introduced by the Swiss scholar Ferdinand de Saussure [8], leads 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 structuring of both factual and narrative information, based on four semiotic relations [1–4].

Our purpose, now, is to argue – and expose by way of simple examples – that our proposed semiotic relations provide suitable criteria to *deconstruct* existing narratives. And, as a practical result of what is learned in this process of analysis, we shall later show (in Section 4) that these same criteria serve to create new narratives by the inverse process of *reconstruction*, similar to Chomsky grammars, which are known as capable of both parsing and generating syntactically correct sentences.

As will be explained when describing the prototype, both processes are taught to the AI agent under the guise of the same set of rules, which will then be shown as they have been specified in natural language. For instance, the rule that verifies whether a given book or film narrative N can be recognized as a *combination* of two narratives N1 and N2, can be used to create a new narrative from a pair of given narratives.

3.1.1. Deriving semiotic relations and operations from rhetorical tropes

Our four basic semiotic relations treated here – *syntagmatic*, *paradigmatic*, *meronymic*, *antithetic* – have been drawn from the so-called four *master tropes*, a topic of prime interest in the area of semiotic research [48]. 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. The so-called master tropes are *metonymy*, *metaphor*, *synecdoche,* and *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) [48]. In modern times they were revived in Kenneth Burke's seminal study [7].

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" [49]. Applications to several topics have been reported, for instance to worldviews and ideologies by the historian Hayden White [50] and, in our own work, to mathematical proof methods [51], to database conceptual modeling [1], and – more relevant to the present discussion – to the digital interactive composition of story-plots [2–4]. In the course of these previous works, we introduced the notion of *semiotic completeness*, founded on the sufficiency of the four master tropes as convincingly demonstrated by Culler and his predecessors, from which our semiotic relations were derived.

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 [8], who originally called the second axis "associative". His horizontal syntagmatic axis modeled sentences as aligned sequences of words connected in obedience to the language syntax, whilst the vertical associative dimension would offer alternative choices to be depicted below some of the sentence components.

The now universally adopted renaming of "associative" to "paradigmatic" was promoted by Jakobson [52], 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 [53], 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." [7]. 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 [48], either simply binary or graduated. No additional dimension is implied. According to Burke [7], the perspective induced by the associated irony trope refers to dialectic, which includes antithesis as a critical phase, expressing negation.

Corresponding to each of the four semiotic relations, four *semiotic operations* can be derived with the double intent of deconstructing and reconstructing narratives, working on broadly known books and films. Letting N, N1, and N2 refer to book or film narratives and S to a scene from N, the operations are:

- Syntagmatic relation: *combine* N1 and N2.
- Paradigmatic relation: *imitate* N.
- Meronymic relation: *expand* S that is part of N.
- Antithetic relation: *reverse* N.

Differently from the semiotic relations, whose abstract formality might cause perplexity in casual users, the names of these four operations favor intuition, and for this motive were employed, in lieu of the relation names, in the prototype's user interface described in Section 4.5.

In what follows we endeavor, by way of examples, to give an idea of how deconstruction works in each of the four cases.

3.1.2. Syntagmatic relation – Lancelot-Grail

The existence of a *syntagmatic relation* between two narratives may justify their joining, often requiring laborious adaptations, a tricky complementary task that has been examined in depth under the denomination of *blending* [16]. Two medieval *Lancelot* and *Grail* stories, both first narrated by Chrétien de Troyes [54] and closely related by their common attachment to Arthur's Round Table chivalric community, were ***combined*** in the so-called *Lancelot-Grail* cycle [55]. But adaptations were required – in special, Lancelot's adulterous behavior with Arthur's queen disqualified him as protagonist in such a 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.

And yet a difficulty still persisted: it was necessary to make Galahad's tricky conception acceptable from a religious standpoint. 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 reissues the Biblical narration of how Laban induced Jacob to heavy drinking during the marriage ceremony, so that the young man would not perceive that he was lying with Laban's eldest daughter, Lia, rather than with his beloved Rachel (*Genesis* 29:25–27).

Actually, 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) [55].

3.1.3. Paradigmatic relation – The Magnificent Seven

The opening scene of the film *The Magnificent Seven* (United Artists, 1960)[3] overtly announces that it ***imitates*** Akira Kurosawa's film *Seven Samurai* (Toho, 1954).[4] First of all, the similarity between the American and the original Japanese narratives stems from their casting of solitary heroes – gunmen in the former, ronin in the latter – in both cases brought together for a minimally rewarded quest: to protect a poor defenseless farmer community against the repeated assaults of an army of bandits.

A ronin[5] was a type of samurai[6] who had no lord or master and, in some cases, had also severed all links with his family or clan. Being unemployed, they were customarily poor, and thus not unwilling to accept the farmers' appeal simply in exchange for regular daily meals. On the other hand, as all types of samurai, they were hereditary members of noble families. Gunmen, too, moved independently most of the time, admired as folkloric figures by some, but execrated for their often-unlawful conduct by those who made efforts to raise the primitive farming settlements in the old American west to the status of official political units. As opposed to the samurai, they were common people, with no pretense to nobility. The farmers begging for the help of the samurai are equally Japanese, whereas, in the American retelling, they come from a neighboring Mexican locality.

The intended narrative similarity is confirmed by a close correspondence, in terms of a number of *mappings*, between the characters and between the plot events. Each group of heroes has a leader, Kambei in the Japanese film, and Chris in the American film, who accepts the calling and recruits his companions, one of which proves his superior fighting ability in a Japanese sword duel, mapped into a typical Old West gun showdown. A radical twist in the American version is the conflation of Katshushiro, the young apprentice of the leader, with the turbulent falsely pretending but brave and resourceful samurai, Kikuchiyo. The conflated character, the hot-blooded Chico, like Kikuchiyo, has a lot to do before overcoming the initial rejection of the group and, meanwhile, like Katshushiro, has a paternally opposed – but most touching – love case with a farmer girl.

3.1.4. Meronymic relation – Barabbas

All four *Gospels*, to be quoted here in the prestigious King James Version,[7] narrate the scene wherein Pontius Pilate, governor of the Roman province of Judaea, prepares to release a prisoner, as was a custom at the feast of Passover. He asks the people to make a choice between Jesus – in whom he claims to see no fault at all – and a certain Barabbas.

The latter is unmistakably a criminal in the four texts:

- John 18:40 – "Now Barabbas was a robber"
- Matthew 27:16 – "a notable prisoner, called Barabbas"
- Mark 15:7 – "had made insurrection with him, who had committed murder in the insurrection"
- Luke 23:19 – "for a certain sedition made in the city, and for murder, was cast into prison"

thus rendering the scene iconic as a supreme injustice, motivated by the envy of the chief priests who, according to Mark 15:11, "moved the people, that he should rather release Barabbas unto them." Also noteworthy is the murderous Roman tyranny, erupting again in the assiduous persecution of the Christians, culminating with their massacre as alleged causes of the great fire of Rome in 64 AD, which could more plausibly be attributed to emperor Nero's folly.

But especially moving is a question left open by the Gospels: how would seeing an innocent victim being sacrificed in his place affect the life of Barabbas from then on?

In an inspired attempt to find a response, the Swedish writer Pär Fabian Lagerkvist ***expanded*** the scene into a Nobel-winning[8] biographic novel entitled *Barabbas* (1950) [56], where the originally non-believing bandit tried to absorb from his contact with Christians a faith he found difficult to share. Punished by renewed imprisonment and lastly by

[3] https://en.wikipedia.org/wiki/The_Magnificent_Seven.
[4] https://en.wikipedia.org/wiki/Seven_Samurai.

[5] https://www.britannica.com/topic/ronin.
[6] https://www.britannica.com/topic/samurai.
[7] https://www.kingjamesbibleonline.org/.
[8] The much revealing writer's Nobel prize speech, which takes the form of an unpublished narrative, is available at: https://www.nobelprize.org/prizes/literature/1951/lagerkvist/speech/.

enslavement in Rome, he absurdly accepts the assertion that Christians are causing the great fire and is caught when trying to contribute to the conflagration. To the writer, the great fire would very likely call to mind the Reichstag fire incident that would foster Hitler's tyrannical domination in Germany, which in turn would soon spread to the neighboring Scandinavian countries [57]. Barabbas, as inexorably fated, is crucified together with Christian martyrs and dies, in a moment of darkness, uttering an enigmatic sentence [ibidem]: "To thee I deliver up my soul.".

Lagerkvist's novel gave origin to a film,[9] equally entitled *Barabbas* (1961), directed by Richard Fleischer and exhibiting some differences from the book.

3.1.5. Antithetic relation – The Shape of Water

The narrative of the film *The Shape of Water* (Fox Searchlight Pictures, 2017)[10] has been classified as a romantic fantasy, deliberately based but in radical opposition to, therefore *reversing*, the horror movie *Creature from the Black Lagoon* (Universal-International, 1954).[11] The Mexican director of the former, Guillermo del Toro, who was also the producer and main author of the screenplay, provided, in an interview,[12] an eloquent account of how the project developed, breaking some preconceived biases of the scarcely tolerant nineteen fifties (a signal virtue of antithetic relations). Watching, when he was 6 years old, *Creature from the Black Lagoon*, he thought "I hope they end up together.".

The villain in the horror movie was definitely the creature, an amphibious humanoid found in the Amazon jungle, predestined to be killed by the hero, the scientist Dr. David Reed, to save his girlfriend Kay Lawrence from the monster's clutches. In *The Shape of Water* romantic fantasy, the creature becomes the victim, and the villain is the man who captured, kept for military purposes, and continuously tortured him: a colonel, called Richard Strickland. Most remarkably, the early victim, promoted to become the hero in the transgressive new movie (situated in the United States environment of 1962), was Elisa Esposito, a custodian at the secret government laboratory where the creature was kept. Found abandoned by the side of a river with wounds on her neck, she was mute, limited to communicating through signs. Her dream was to save the life and restore the freedom of the creature.

Her frantic quest seems to terminate in failure as she dies, but the creature's magic healing power not only revives her but also, when he touches the scars on her neck, they open to reveal gills like his own. So, while the creature is introduced as a survivor from the Devonian period, Elisa turns out to be a mutation with compatible traits. At first sight, such narratives would seem affiliated with the *Beauty and the Beast* folktale [58], but there is a snag: the monster does not change into a prince, which would make him acceptable as a consort. In fact, a more closely related folktale might be evoked, *The Ugly Duckling* [59], since Elisa's birthmarks – the scars on her neck and her muteness – were a sign of her future felicitous metamorphosis, rather than marks of inferiority.

3.2. Semiotic enactment

Having treated *thematic* structure in Subsection 3.1, as the preliminary perspective to be considered within Jonathan Culler's double logic theory, we must now turn to the other perspective, involving *event enactment*. One or more alternative sequences of events must be determined – by some sort of backward reasoning, as previously stressed – able to reach a target situation compatible with the defined thematic structure, starting from an equally compatible initial situation.

Clearly, this backward reasoning event enactment process corresponds to the semiotic notion of *abduction*, as elaborated by Charles Sanders Peirce [9]. Having compared our strategy to case-based approaches [10,20], we should remark that the possibility of achieving case-based reasoning via abduction has been demonstrated [60], in particular via abductive logic programming [61], which incidentally establishes a connection with our early projects.

In addition, as will be indicated when the prototype interface is described in Section 4, the user can optionally issue instructions on the composition of each individual event, which the prototype will try to follow as closely as possible while preserving the overall consistency of the narrative – again resorting to an application of abduction.

The abduction process consists not only of finding a suitable event or event sequence but also of choosing, among the often-plural alternatives, which one is the best according to some criterion. Umberto Eco and Thomas Sebeok [62] have already proclaimed Peirce's abduction as the method used by fictional detectives, taking advantage of their *training* on a number of subjects that may provide information pertinent to their trade; Sherlock, for instance, took instruction on biology, chemistry, anatomy, English law, and sensational criminal cases (once again, other previous narratives). Working on this unsystematic knowledge, they gained fame for their skill to discover one or more viable stories of the crime and determine which alternative actually corresponded to the observed facts.

But let us hear how Sherlock Holmes himself explains to Watson his abductive style of investigation, quite appropriately described as reasoning in inverse direction [63]:

> "Most people, if you describe a train of events to them, will tell you what the result would be. They can put those events together in their minds, and argue from them that something will come to pass. There are few people, however, who, if you told them a result, would be able to evolve from their own inner consciousness what the steps were which led up to that result. This power is what I mean when I talk of reasoning backwards, or analytically."

Clearly "reasoning backwards" brings to mind Peirce's abduction. But, even more importantly, one characteristic of detective stories, pointed out by Tzvetan Todorov, can arguably be shown to be directly applicable to our project. According to Todorov [64], any detective story actually "contains not one but two stories: the story of the crime and the story of the investigation … the first – the story of the crime – tells 'what really happened,' whereas the second – the story of the investigation – explains 'how the reader (or the *narrator*) has come to know about it.".

A conventionally composed story that lacks this so-called 'whodunit' ('who [has] done it?') interrogative conundrum of detective stories usually comprises no more than a single narrative and its author. In turn, it just so happens that the event-by-event composition process, led by the prototype in the same abductive ('reasoning backward') style of a 'detective investigation', involves the continuous performance of a narrator, the AI-Agent, under the delegation of the user-as-author, and two narratives emerge, namely the generated story itself and the record of the interactive story composition process.[13]

Typically, friendly supporting characters, such as Watson, Lestrade, Hastings, and Japp, among others, who were aware of the same facts, failed to perceive them as clues. The privileged position of the great detectives – and hence of the AI agent, playing the detective role while trying to satisfy the structural determinant and the user's prompts – commonly attributed to "intuition", is explainable thanks to an additional feature of Peirce's theory, namely his triadic model of the sign, consisting of object, representamen, and *interpretant*. Roughly speaking, representamen and object correspond to Saussure's signifier and signified [8], but the interpretant, which introduces a deeper level of meaning, is what may mark the superior analytic power of a Dupin or

[9] https://en.wikipedia.org/wiki/Barabbas_(1961_film).

[10] https://en.wikipedia.org/wiki/The_Shape_of_Water.

[11] https://en.wikipedia.org/wiki/Creature_from_the_Black_Lagoon

[12] https://www.latimes.com/entertainment/movies/la-et-mn-oscars-best-picture-20180304-story.html.

[13] As announced in the Introduction, we recognize the relevance of this record, and our prototype is programmed to store both narratives.

Sherlock. It is like a translation of the signifier that takes form in the human memory, and its structure may significantly differ in complexity from person to person – or, more precisely, in Peirce's words [9]:

> "A sign, or representamen, is something which stands to somebody for something in some respect or capacity. It addresses somebody, that is, creates in the mind of that person an equivalent sign, or perhaps *a more developed sign*. That sign which it creates I call the interpretant of the first sign."

Having alluded to "intuition", we find it useful to recall that much psychological research effort has seriously concentrated on the study of *Intuitive Decision Making* (IDM), based on the argument that in real-life situations, the more traditional Analytical Decision Making methods can hardly be practiced when one encounters [65][14] "time-constrained, complex, dynamic conditions in response to incomplete or uncertain information." Intuitive decisions can occasionally be sub-optimal or even fail altogether, but prior to the recent boom of analytical decision-making, thanks to advances in Data Science, Data and Process Mining, and a flurry of processes to handle Big Data, successful management has often relied on IDM.

The same authors associate IDM with a special type of learning, namely *Implicit Learning* (IL), while referring to special *neurocognitive structures* as their common location:

> "… these studies suggest that IDM processes share some of the same underlying neurocognitive structures and processes as a type of learning known as Implicit Learning (IL) [which] appears to accelerate a novice's ability to learn deep structures and rules embedded within information if delivered in an appropriate manner."

Such special neurocognitive structures seem to us the material residence appropriate to host Peirce's interpretant. Another famous detective, Hercule Poirot, draws his friend's attention to the processing unit necessary to perform intelligent processes [66]: "These little grey cells. It is 'up to them'—as you say over here." In an obvious transition from fiction to modern science, Poirot's reliance on his methodical little grey cells brings into consideration the intricate work of neurons and synapses [67].[15]

More specifically, as reported in [68],[16] neuroscientific investigations have revealed the key brain regions involved in learning: medial temporal lobe activity is associated with explicit memory for prior examples; posterior caudate activity correlates with key brain systems associated with implicit learning; and dorsolateral prefrontal cortex activity is associated with resolving competition between implicit and explicit processing. Also, repeatedly stressing the practical advantages of IDM and IL, the previously mentioned project of Squire and colleagues [65] investigated how both could be enhanced by special *training* methods. It has been argued that training by way of implicit learning, besides expanding the person's knowledge of a domain of interest, serves to develop *commonsense reasoning* [69] in substitution to formally specified algorithms.

In contrast, in our previous work, we applied an entirely analytical paradigm based on plan generation and plan recognition, inspired by Propp's functions [70], which we adapted to diverse fictional and serious domains and formalized in terms of pre-conditions and post-conditions, as part of a three-schemata model (specifying classes of facts, events, and agents), and expressed in temporal logic [71]. In common with the present discussion, we may observe in retrospect that semiotic abduction was already present in our work, since any one of possibly several alternative plan-generated plot P_i, leading to a final desired situation S, could be viewed as the antecedent of a rule $P_i \rightarrow S$. Indeed, methods of *abductive planning* have already been proposed, based on event calculus [72,73], among several other AI approaches to abduction [74].

Since the very beginning, our early prototypes [75–78] that followed this paradigm, similarly to the prototype now under discussion, allowed user interaction at intermediate points reached by the planner in response to short-term goals. One serious inconvenience inherent in this paradigm was the difficulty of specifying the logic schemas to ensure the satisfaction of integrity constraints and conform to the conventions of the chosen genre. We had to acknowledge that the formal specification of each domain was clearly beyond the reach of final casual users, and consequently required the intermediate participation of a domain designer, with a minimum of logic expertise, to supplement the general-purpose prototypes created by us. This difficulty affecting a number of symbolic story generation systems that work in a backward (i.e. abductive) fashion is clearly expressed in [28]:

> "Some of these systems start with an end state—the state the fictional world should be in at the end of the story—and work backward, determining what must have happened to transform an initial world state into the goal. These systems often generate coherent stories guaranteed to end in a given state. Their drawback is that they require significant hand-authored domain knowledge."

The current prototype does not require domain designers; indeed, it does not require domain specifications, since these are supplied by the user's choice of book or film narratives, whose synopses would have become familiar to ChatGPT as part of its training. As said at the beginning of this section, event enactment is an abductive task. On the other hand, the valuable psychological research investment in training for implicit learning has reached an unexpected climax with the advent of Chat Generative Transformers capabilities [79].

With the continuing progress of successive generations of such large language models, one begins to count on hidden artificial structures to emulate the special human neurocognitive structures, thereby allowing to accommodate *the more developed signs* – the subtle interpretants – postulated by Peirce. Through a two-way translation process, all reasoning activities (including abductive reasoning) are internally processed over the vector representation, which serves at the *interpretant* level, while, to the users, the natural language keeps the status of a single explicit *representamen*. Hopefully, it has become possible to effectively cope with Implicit Learning, as claimed in [80]:

> "There is consensus on the fact that an associative model of cognition can explain psychological phenomena such as learning, priming, *implicit learning* and associations … and not many are dubious on the fact that LLMs, being complex associators, can simulate such mental processes."

– and perhaps, hopefully, being apt to contribute narrative "intuitive" turns, despite the risk, also affecting humankind, of incurring in errors and hallucinations.

In an attempt to give a more concrete idea of how our version of event enactment is supposed to work, we found it convenient to briefly describe one reconstruction experiment with the prototype, thereby anticipating the detailed technical presentation of the next section. The experiment started as an application of the *reverse* semiotic operation to William Shirer's book *The Rise and Fall of the Third Reich* [57] and in the "Protagonist" field we specified: "Chamberlain and Daladier in the surrender at Munich".

To begin with, we just let the prototype compose the narrative without providing any event orientation. Under the title *The Bittersweet Triumph: Munich's Surrender*, a counterfactual account was produced, wherein the two allied leaders stayed until the end of the war. The

[14] Work performed by a team of six specialists, partially under contract with the Office of Naval Research (ONR).

[15] The first author, C. R. Gallistel, is an Emeritus Professor of Psychology at Rutgers University, specialized in the cognitive processes of learning and memory.

[16] The author, Paul J. Reber, professor of Psychology & Communication Sciences and Disorders at Northwestern University, was a distinguished participant of the ONR-sponsored project.

reverse structural determinant was obeyed through an ultimately successful change in their political conduct, which transformed the disastrous surrender into vigorous resistance leading to the ultimate defeat of the aggressor:

1. In their endless pursuit of peace, Chamberlain and Daladier yield to Hitler's demands at Munich, signifying an act of surrender unheard of in their history.
2. Back in Britain and France, Chamberlain and Daladier defend their actions with the phrase "Peace for our time", a hollow phrase echoing through somber, disillusioned crowds.
3. As war breaks out, despite their efforts for peace, Chamberlain and Daladier must steer their nations with an immense burden of guilt and responsibility.
4. As the direness of their situation sets in, Chamberlain and Daladier rally their nations, inspiring resilience and courage in the face of adversity, the seeds of their redemption beginning to germinate.
5. Under their leadership, Britain and France start to turn the tides of war, their earlier surrender at Munich becoming a catalyst for their unyielding spirit to fight back.
6. Chamberlain and Daladier, once seen as architects of surrender, are transformed into symbols of resilience, their initial failure yielding an unexpected victory.
7. In the aftermath of the war, Chamberlain and Daladier face their nations as they come to grips with their roles in history, their surrender at Munich is now seen as a bitter catalyst for eventual victory.

In a second try, we proceeded to interfere with the events in order to satisfy our preferences. The prototype's reactions, when processing this more controlled form of composition, served to illustrate its concern to accommodate the user's (in this case, our own) interactive prompts while preserving consistency with the context involved.

The first event generated was "Chamberlain and Daladier come to an agreement with Hitler in Munich, consenting 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.".

The resulting story, entitled *The Honor of a Nation* by the prototype, thus comprised the following events:

1. Chamberlain and Daladier come to an agreement with Hitler in Munich, agreeing to cede the Sudetenland to Nazi Germany in exchange for "peace."
2. The annexation of Czechoslovakia by Nazi Germany prompts Chamberlain to realize his grave mistake.
3. 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.
4. A group of German generals, led by Erwin Rommel, successfully deposes Hitler and takes control of the government in a peaceful coup.
5. The generals begin the process of negotiating a peace treaty with the Allied powers to end the war and restore democracy in Germany.

This example raises two important points. Firstly, it is an illustration that our model and its prototype implementation are not restricted to fiction. Secondly, the real historical occurrence treated here involving the conquering ambition of an aggressive government is not simply a memory of the past. Indeed, it occurs again in the present time in the far more grievous context of nuclear armament, reopening the dilemma of whether the danger of war should be faced by friendly negotiation or by the menace of prompt military retaliation.

In contrast, contemporary treatments of historical facts are not necessarily more informative than the older ones. Contrary to book authors' usual practice, screenwriters of films said to be "based" on historical facts commonly feel free to distort events to produce "dramatic effect" – whereas we all tend to assume that whatever is shown on the screen is authentic. A recent example is the film *Oppenheimer* (Christopher Nolan, 2023), with respect to which we learn that the specific conversations between the protagonist and Einstein "which the film revolves around never happened."[17].

4. The ChatGeppetto prototype

The exploration of semiotic relations within existing narratives allowed us to identify crucial strategies for repurposing narrative elements to produce novel stories. By taking advantage of the recent advancements in LLMs, we leveraged ChatGPT's remarkable comprehension of the proposed semiotic operations and its vast narrative knowledge base to construct ChatGeppetto, an AI-powered system for automated narrative generation.

4.1. The architecture of ChatGeppetto

The architecture of ChatGeppetto, as shown in Fig. 1, follows a multi-agent approach for narrative generation and visualization. Two distinct AI agents are employed: a Storywriter AI Agent and an Illustrator AI Agent. The Storywriter AI Agent, specialized in natural language generation, generates story events and scene descriptions. Concurrently, the Illustrator AI Agent utilizes these textual descriptions to produce corresponding visual representations. Coordinating these agents is the Plot Manager, who acts as the central orchestrator, directing the overall narrative generation process by requesting story events and illustrations from each agent respectively. User interaction is facilitated through a user-friendly interface, enabling the composition and visualization of novel narratives in video and storyboard formats. All generated narratives are stored within a dedicated Stories Database.

ChatGeppetto uses the embedded knowledge within its AI agents to create novel narratives by applying the proposed semiotic operations. For seamless integration and future upgrades, a plugin architecture was adopted for the system, facilitating the replacement of AI agents with newer and more powerful models as they become available. In the current iteration, the Storywriter AI Agent uses the text generation capabilities of the ChatGPT model (GPT-4) accessed via the OpenAI API.[18] The Illustrator AI Agent, on the other hand, is built upon a fine-tuned version of the Stable Diffusion 1.5 model hosted on a private server and accessible through a REST API.

The system adopts a web application architecture, enabling access through any web browser, encompassing desktop and mobile devices.

[17] https://en.wikipedia.org/wiki/Oppenheimer_(film)#Accuracy_and_omissions.

[18] https://platform.openai.com/docs/api-reference.

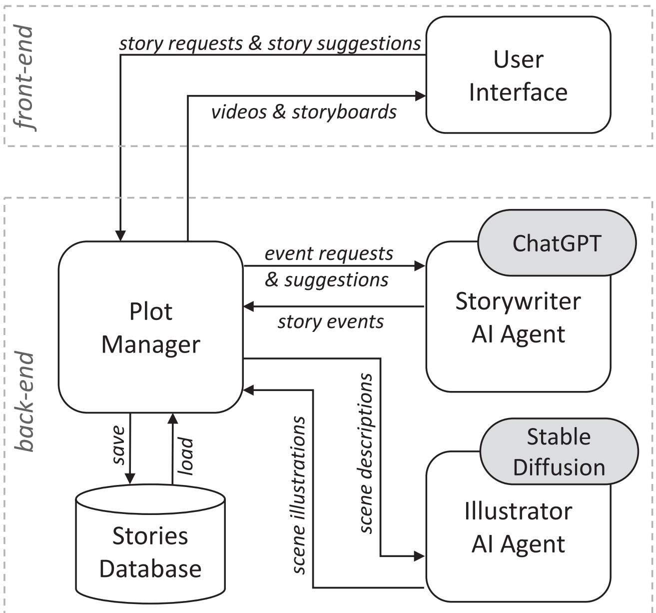

Fig. 1. The architecture of ChatGeppetto.

This architecture follows a three-tiered model for web applications, separating its components and promoting maintainability. The back end, implemented using the server-side scripting languages PHP and Python, hosts the Plot Manager, which orchestrates the narrative generation process, and the communication interfaces for interacting with the AI agents. The front-end, built with HTML and JavaScript, provides the user interface, a visually intuitive platform for user interaction and story visualization. Finally, a MySQL database serves as the persistent storage layer, securely storing generated narratives for future access and retrieval.

4.2. Narrative generation

Drawing upon ChatGPT's extensive literary and movie knowledge, ChatGeppetto employs a set of semiotic operations to generate novel narratives through strategically reusing and adapting existing plots. As detailed in Section 3.1.1, the *combine* operation merges plot elements from two *media sources* ($M1$ and $M2$) to create a new story. The *expand* operation delves deeper into a specific scene within a single source ($M1$) using a provided scene description (S) to craft a new narrative. Finally, the *imitate* and *reverse* operations require a single source ($M1$). While *imitate* generates a new story resembling $M1$, *reverse* constructs a narrative antithetical to the plot of $M1$. Movies or books can be used as the *media type* (MT) for the sources. The media sources ($M1$ and $M2$) can be any existing movie or book (depending on the selected MT) or any story previously created within the system. In both cases, the sources are identified by their titles.

In addition to the selected semiotic operation, media type, existing media sources, and scene descriptions, the *protagonist description* (PD) plays an important role in the narrative generation process. The protagonist(s) can originate from the chosen stories, be assigned arbitrary names, or even be drawn from entirely unrelated contexts, as exemplified by the classic novel "*A Connecticut Yankee in King Arthur's Court*" [81] where a contemporary individual finds himself thrust into a historical setting. Beyond the protagonist's name, the PD parameter allows users to specify preferences related to their characteristics (physical attributes, emotions, desires) or their current situation (location, actions, past events).

ChatGeppetto uses the aforementioned parameters to generate a set of prompts for the Storywriter AI Agent (currently based on the GPT-4 model). These prompts, created by initial experiments assessing ChatGPT's ability to identify semiotic relations within existing narratives, were designed in a standardized three-part structure:

1. Semiotic Operation ($P_{semiotic}$): This section defines the pertinent semiotic relation to be used in the narrative generation process.
2. Task Description (P_{task}): This section details the specific task the AI agent needs to perform, indicating how the semiotic operation must be applied.
3. Output Format (P_{output}): This section specifies the expected format of the generated narrative, allowing the system to interpret the content generated by the AI agent.

Table 2 provides an overview of the definitions and task descriptions associated with each available semiotic operation.

The output format specification was established to ensure that it can be easily interpreted by ChatGeppetto. This format dictates that the Storywriter AI Agent generate narratives adhering to a specific structure containing the following elements:

1. Event Description: This initial line, denoted by "EVENT:", provides a concise description of the key event occurring within the narrative.
2. Scene Description: The subsequent line, beginning with "IMAGE:", offers a brief description of an image visually depicting the aforementioned event.
3. Story Title: Finally, a creative title is generated for the newly produced story, preceding the narrative content, and marked by "TITLE:".

The exact instructions provided to the Storywriter AI Agent 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*

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 stories. The parameters are: <MT> – media type (movie or book); <M1> – related media source 1; <M2> – related media source 2; <S> – scene description; and <PD> – protagonist's description.

Semiotic Operation		Prompt
combine	Definition	"*The fictional protagonist in the plot of <MT> A is X. The fictional protagonist in the plot of <MT> C is Y. <MT> A has a syntagmatic relation with <MT> C if X does not appear in C, Y does not appear in A, and there exists a <MT> B that features both X and Y.*"
	Task	"*Considering this definition, write the first event for a new <MT> B considering that <MT> A is '<M1>' and <MT> C is '<M2>'. The protagonist of <MT> B must be <PD>.*"
imitate	Definition	"*<MT> A has a paradigmatic relation with <MT> 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 <MT> B considering that <MT> A is '<M1>' (<MT> B must have a paradigmatic relation with <MT> A). The protagonist of <MT> B must be <PD>.*"
expand	Definition	"*There is a meronymic relation between a <MT> A and a <MT> 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 <MT> '<M1>', S be a scene taken from the plot of A where <S>, and letting B be an imaginary new <MT>, write the first event for the plot of B, so that there is a meronymic relation between A and B. The protagonist of <MT> B must be <PD>.*"
reverse	Definition	"*<MT> A has an antithetic relation with <MT> B if there is some similarity between their plots but the objective of the protagonist of <MT> A is the complete opposite of the objective of the protagonist of <MT> B.*"
	Task	"*Considering this definition, write the first event for a new <MT> B considering that <MT> A is '<M1>' (<MT> B must have an antithetic relation with <MT> A). The protagonist of <MT> B must be <PD>.*"

image that illustrates the narrative event. Also, generate a new creative title for <MT> B and add it at the beginning of the response starting with 'TITLE:'.". These standardized instructions, provided to the Storywriter AI Agent, guarantee the generation of narratives in a format readily interpretable by ChatGeppetto's processing mechanisms.

When the user decides to use a story previously generated within the system as a media source, an additional instruction describing the events of the selected story is added to the beginning of the initial prompt used by ChatGeppetto to instruct the AI agent on how to generate stories. Considering that $M1$ is a story generated by ChatGeppetto, which comprises a sequence of events ($M1_{events}$), the additional instruction will be defined by: "*Consider the following storyline of a <MT> called '<M1>': <M1events>.*". When both $M1$ and $M2$ are stories generated by ChatGeppetto, the events and titles of both stories are described. This additional prompt is called $P_{stories}$.

A key feature of the initial prompt lies in its directive for the Storywriter AI Agent to only generate the first narrative event. This allows users to actively engage in a step-by-step narrative generation process, fostering a dynamic and collaborative storytelling experience. At every step, four operations related to the narrative generation process are available for users: (1) *continue*, which allows users to request the Storywriter AI Agent to write the next event for the story (optionally suggesting what should happen next); (2) *finish*, which is used to finalize and save the current story; (3) *regenerate*, which allows users to request the Storywriter AI Agent to write an alternative version of the last story event; and (4) *delete*, which is used to remove the last event from the story. This interactive story generation process allows users to propose new events, refine existing ones, determine the narrative's pace and progression, and ultimately signal when the story should reach its conclusion.

The *continue* operation empowers users to actively engage with the Storywriter AI Agent by providing suggestions in natural language that can influence the direction of the narrative. These suggestions are optional, which allows users to either actively shape the story through creative suggestions or passively observe the AI Agent's autonomous generation of the next events. Internally, a specific prompt ($P_{continue}$) is used 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 (user suggestion), the continue prompt is complemented by the suggestions prompt ($P_{suggestions}$): "*Consider the following suggestions for the story:* " + user suggestions. This approach fosters a highly effective user-driven interactive storytelling experience, offering both autonomous and collaborative narrative generation options.

In order to generate stories, ChatGeppetto explores ChatGPT's ability (GPT-4) to engage in conversational interactions, enabling the AI agent to comprehend the story context and instructions provided within prior messages. This collaborative narrative generation process follows a predefined conversation structure:

1. ChatGeppetto: $P_{stories} + P_{semiotic} + P_{task} + P_{output}$
2. ChatGPT: TITLE: *story title*. EVENT: *event description*. IMAGE: *scene description*.
3. ChatGeppetto: $P_{continue} + P_{suggestions}$
4. ChatGPT: EVENT: *event description*. IMAGE: *scene description*.

Messages 3 and 4 of the conversation structure are repeated until the user decides to finish the story (i.e., when the user feels that the story is complete), which is indicated in the user interface of ChatGeppetto by pressing a "Save Story" button.

In addition to the *continue* operation, users can refine the narrative by requesting the regeneration of the last story event. The *regenerate* operation removes the most recent event from the story plot and prompts the Storywriter AI Agent to generate a new, alternative event. By exploring the stochastic nature of ChatGPT, this operation introduces different variations of story events, allowing users to explore various narrative possibilities until they find one that aligns with their preferences.

4.3. Image generation

The proposed narrative generation process explores the recent advancements in text-to-image machine learning models for the creation of visual representations to accompany the generated story events. These models, including DALL-E, Midjourney, and Stable Diffusion, possess the remarkable capability of producing detailed and intricate images based on natural language descriptions. For our implementation, we specifically employ the Juggernaut model,[19] which is a fine-tuned version of the Stable Diffusion 1.5 model.

The scene descriptions generated by the Storywriter AI Agent play a crucial role in driving the image generation process. While directly feeding these textual descriptions into the Stable Diffusion model would suffice for producing illustrative representations of the narrative events, we prioritize user control. Therefore, ChatGeppetto incorporates an additional parameter, the illustration style (*IS*), which allows users to customize the visual aesthetic of the generated images. This parameter is seamlessly combined with the Storywriter AI Agent's scene description to form the final prompt provided to the Illustrator AI Agent: $SCENEDESCRIPTION + IS$. ChatGeppetto's default illustration style is "*digital painting, hyperrealistic, highly detailed, sharp focus, stunningly beautiful, cinematic lighting*", which generates aesthetically pleasing and cinematic-looking images. However, users are allowed to modify this style by providing any desired descriptions, including famous artists (e.g., "*painted by Pablo Picasso*") or specific artistic movements (e.g., "*cubism painting*"). This user-driven approach fosters increased creative control over the visual representation of the narrative.

The proposed image generation process also complements the basic prompt with a general negative prompt, which aims to guide the Stable Diffusion model toward generating more visually accurate and aesthetically pleasing images by explicitly identifying undesirable visual elements. The negative prompt was meticulously crafted based on weaknesses observed in the model's outputs 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 convey emphasis, with more parentheses representing a greater emphasis on the undesired elements.

The use of parenthesis to increase the emphasis of certain elements in the scene illustrations is also applied to the scene descriptions generated by the Storywriter AI Agent. By emphasizing nouns and adjectives, the Illustrator AI Agent is capable of producing more clear and concise scene illustrations. The process of automatically adding parenthesis to scene descriptions is implemented using the natural language processing library spaCy.[20] The process begins with parsing the scene description using spaCy's *en_core_web_lg* model, which provides linguistic annotations, enabling the identification of noun phrases and their modifying adjectives. For each noun chunk identified in the scene description, the algorithm distinguishes nouns and proper nouns (NOUN, PROPN) from their associated adjectives (ADJ) and coordinating conjunctions (CCONJ). This differentiation is important for understanding the semantic structure of the description and for identifying the main elements within the scene description. Then, the method constructs noun-adjective pairs, mapping each noun into its preceding adjectives, if any. This mapping is crucial for determining the emphasis strategy:

[19] https://civitai.com/models/46422/juggernaut.

[20] https://spacy.io/.

nouns without adjectives receive a double set of parentheses, signaling a lesser degree of emphasis, whereas noun-adjective pairs are enclosed in triple parentheses for a higher degree of emphasis (improving the representation of attributes of nouns in the scene illustrations).

For example, by giving the following scene description: "*A dimly lit, dusty room overflowing with books, in the midst of which a surprised Sherlock Holmes stands before a shimmering, circular gateway emitting a mystical glow",* our method will produce the following description: "*A dimly lit, (((dusty room))) overflowing with ((books)), in the ((midst)) of which a (((surprised Sherlock Holmes))) stands before a shimmering, (((circular gateway))) emitting a (((mystical glow)))*".

When generating illustrations for the story, we also control the *seed* utilized by the Stable Diffusion model. This parameter acts as an initialization point for the model, enabling the production of similar images (the same seed and prompt consistently yield identical output images). Additionally, minor variations in the prompt while maintaining the same seed often result in visually similar images. To ensure visual coherence throughout the narrative, the Illustrator AI Agent strategically employs the same seed for all images generated within a single story. This approach helps to maintain a consistent visual aesthetic across the various events depicted in the story, reinforcing the narrative's overall cohesion.

4.4. Video generation

The proposed system extends beyond static storyboards by incorporating a video generation functionality. This feature seamlessly blends image, text, and audio narration, producing an audiovisual rendition of the generated stories. Once a user finalizes the narrative composition, the video generation process is performed using the generated story plot. This data, encompassing both event descriptions and image illustrations, serves as the input for the video generation process.

Given as input a story plot $S = \{e_1, e_2, \cdots, e_n\}$, where each event is a pair $e_i = (text_i, image_i)$ containing the textual description of the event ($text_i$) and a corresponding image illustrating the event ($image_i$), the video generation process 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.

The video generation process relies on a combination of well-established libraries, APIs, and tools. The Python Imaging Library (Pillow),[21] facilitates the seamless blending of text and image elements, forming the foundation for the video's visual content. The OpenAI TTS API[22] allows the system to generate an audio narration that complements the visual storytelling. Finally, the FFmpeg tool[23] plays a crucial role in both synchronizing and seamlessly integrating the generated audio narration with the corresponding visual elements. Additionally, it enables the merging of individual video clips into the final video output.

4.5. The user interface of ChatGeppetto

ChatGeppetto is implemented as a web application that is accessible through any web browser. Users can access ChatGeppetto's interface through our public website: https://narrativelab.org/chatgeppetto/.

For the initial user interface design of ChatGeppetto we have used most of the design principles as outlined by Blair-Early and Zender [82] to optimize user engagement and ease of use. As illustrated in Fig. 2, the user interface of the system is designed to facilitate intuitive story composition. At the top of the page, an explanation of the tool is given, which is an obvious starting point that allows users to understand how to use the tool [82]. Users begin by selecting a desired semiotic operation from the provided options (action 1 in Fig. 2 (a)). They then input the necessary parameters (action 2 in Fig. 2 (a)) according to the chosen operation. Initiating the story generation process is achieved by pressing the "Generate Story" button (action 3 in Fig. 2 (a)). This button is in a different color than the rest of the interface, which increases its salience and the probability of a user being able to locate it [89]. Furthermore, all clickable links and buttons are in a color that is easy to notice to our users. After pressing the "Generate Story" button, the system transitions to the story composition screen (Fig. 2 (b)), where users can view the generated story's title, the first event with its corresponding scene description, and a visual illustration.

In the example showcased in Fig. 2, the user first selected "movie" as the media type, the "combine" operation (action 1 in Fig. 2 (a)), entered the movie titles "Sherlock Holmes (2009)" and "The Hobbit: An Unexpected Journey", and described the protagonist as "Sherlock Holmes" (action 2 in Fig. 2 (a)). Then, clicking on the "Generate Story" button (action 3 in Fig. 2 (a)), prompted ChatGeppetto to generate a new story titled "Holmes & Hobbit: The Intersection of Worlds". After viewing the first event, the user has various options: either regenerate the entire event by pressing the reload button on the scene's bottom-right corner, for which a familiar icon is used (Fig. 2 (b)) [82], or regenerate only the visual illustration, or continue the story by pressing the button "Continue Story" (action 4 in Fig. 2 (b)) to generate the next event as a storyboard (Fig. 2(c)) or provide suggestions in the "Optional Suggestions" field. The system will consider these suggestions during the next event generation or regeneration.

Upon concluding the story composition, the user presses the "Save Story" button (action 5 in Fig. 2 (c)) to signal the end of the narrative development. ChatGeppetto then automatically generates a video rendition of the completed story. Once the video generation process is complete, the user is presented with the final narrative (Fig. 2 (d)). Additionally, clicking the "View Story in Storyboard Format" link (action 6 in Fig. 2 (d)) allows users to revisit the entire story in its original storyboard format (Fig. 2 (e)).

Saved narratives are automatically added to the user's personal library (Fig. 2 (f)), which includes a few illustrative examples we preloaded. This library can be accessed from various points within the system: the initial screen (Fig. 2 (a)), the story visualization screen in video format (Fig. 2 (d)), and the storyboard visualization screen (Fig. 2 (e)). To access the library from these screens, users simply click the "View User Library" link (action 8 in Fig. 2 (e)). User privacy is maintained by associating each saved story with a unique user key. This key is automatically generated upon a user's first interaction with the system and is stored locally in the user's web browser using a web cookie.

Besides saving the generated narratives in both video and storyboard format, ChatGeppetto also keeps track of all composition steps performed by the user to produce the final narrative, including suggested events and image regenerations. Users can visualize the composition steps of a story by clicking on the "View Composition Steps" link in the story visualization screens (video and storyboard) (action 7 in Fig. 2 (d)). The visualization of the composition steps is most useful to users, since it documents their work, keeping track of their positive or negative reactions towards the events and scenes generated by the system, and allowing them to inspect their manually inserted corrective prompts. Hopefully, reasoning about such documentation will contribute to improving not only their future experiments with the prototype but even their own composition skills.

Examples of stories generated by ChatGeppetto are available at: https://narrativelab.org/chatgeppetto/list.php. In particular, a video narrating the story "Holmes & Hobbit: The Intersection of Worlds",

[21] https://pypi.org/project/Pillow/.
[22] https://platform.openai.com/docs/guides/text-to-speech.
[23] https://ffmpeg.org/.

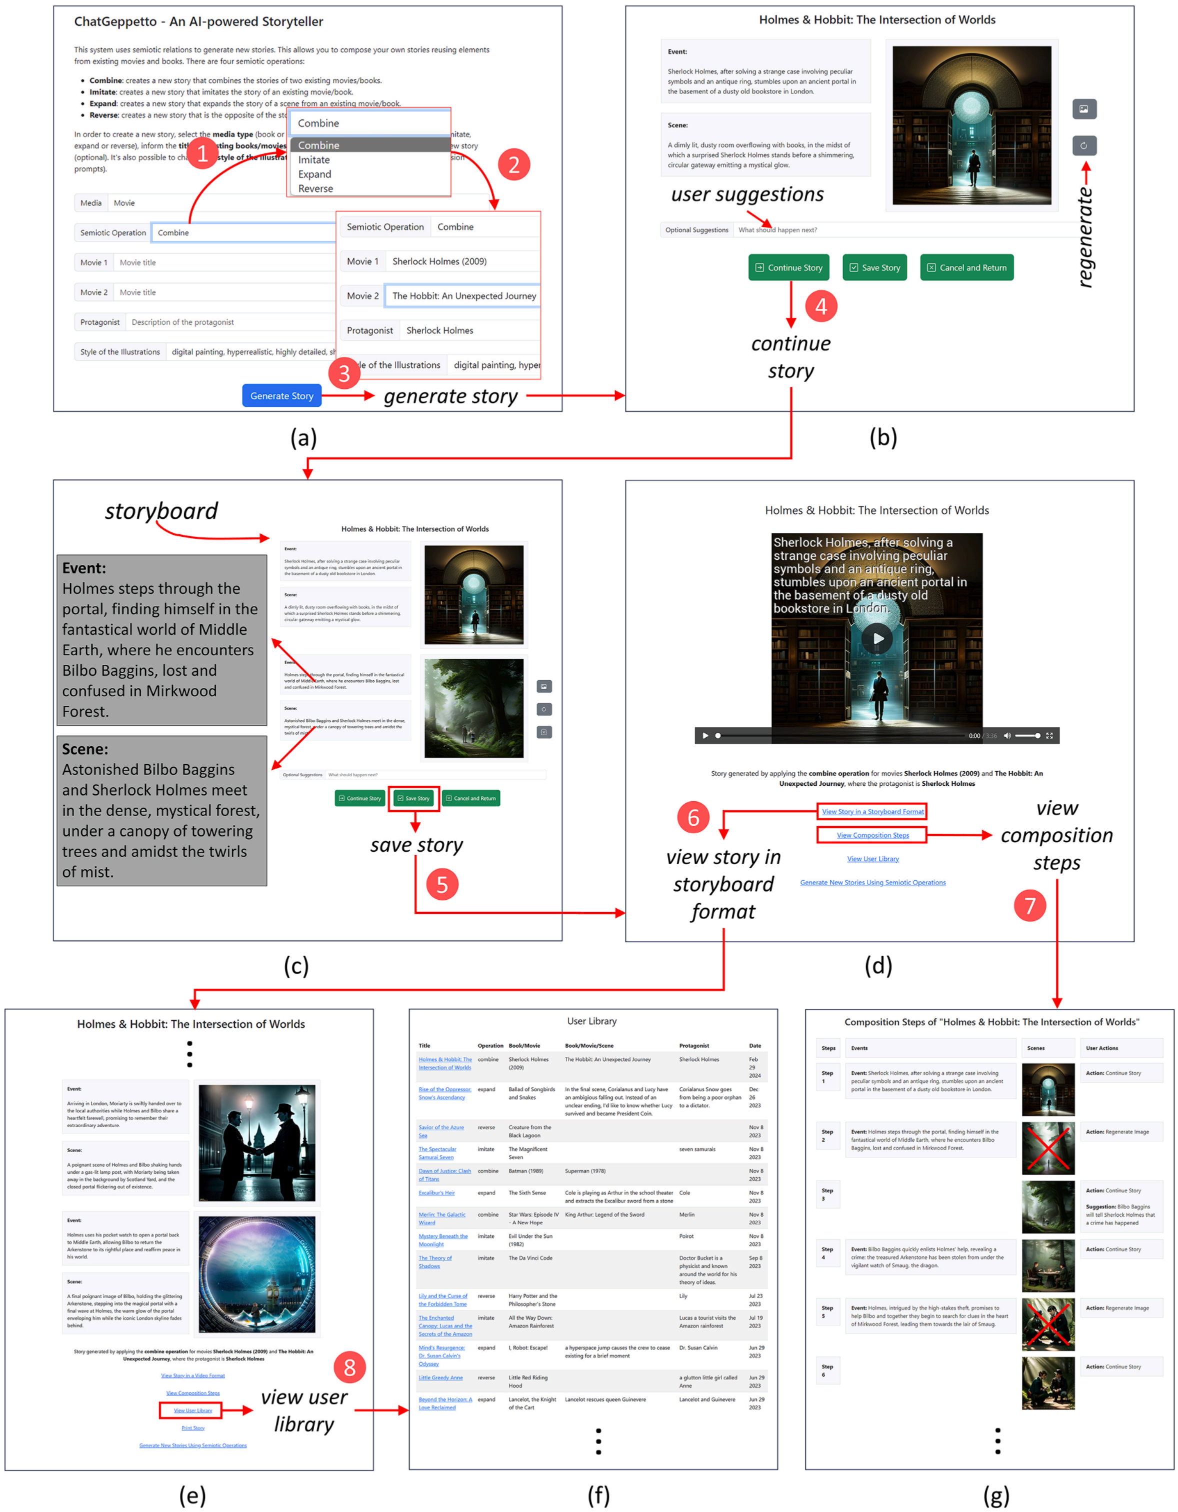

Fig. 2. The user interface of ChatGeppetto.

whose opening screen is shown in Fig. 2, can be accessed at: https://narrativelab.org/chatgeppetto/video.php?id=548d. Appendix A presents all the composition steps performed to create this story.

5. Evaluation and results

ChatGeppetto and the proposed method for the reconstruction of narratives using semiotic relations were designed to assist non-professional authors in the process of creating stories. In order to evaluate whether our tool and method have achieved their goals, we conducted a usability test aiming at:

1. Assessing the overall user experience in our tool;
2. Evaluating user satisfaction with the narratives created using our method; and
3. Analyzing the impact of semiotic relations on the user experience and satisfaction.

5.1. Participants and design

A group of 42 participants (28 males, 13 females, and 1 who preferred not to indicate the gender; M_{age}= 23.44, SD_{age} = 6.00, *Range* = [18 - 48]) was recruited for this study. 21 participants had Dutch nationality, 5 had Polish nationality, 3 had Bulgarian nationality, 2 had Portuguese nationality, and every other nationality only occurred once in the dataset. 31 participants reported no experience with story writing, and the other 11 had varying levels of experience, ranging from writing Dungeons & Dragons scripts to professional copywriting experience. 30 participants use generative AI at least once per week.

Participants were either bachelor's students or teachers of the Applied Data Science and Artificial Intelligence program at Breda University of Applied Sciences, or they were invited by means of convenience sampling. Informed consent of each participant was obtained before collecting any data.

The subjects participated in a between-subjects design with two conditions: they either interacted with Full ChatGeppetto, or with a modified version called Simplified ChatGeppetto. The main dependent variables are user experience with the tool and user satisfaction with the final story.

5.2. Materials and measurements

To evaluate the impact of the proposed method for reconstructing narratives using semiotic relations, we implemented a simplified version of ChatGeppetto that does not use semiotic relations in the narrative generation process. Instead of having the user select a semiotic operation and existing books/movies as input, in the Simplified ChatGeppetto the user informs of a short description of the desired story (a premise) and describes the protagonist. The initial screens of the two versions of the system are presented in Fig. 3. Except for the input, all other features and interaction processes are the same in both versions of the system. In this way, we can test the effect of introducing the proposed semiotic operations in the story generation process. It is also important to notice that the initial screen of Full ChatGeppetto includes an explanation about the semiotic operations, which allows users to understand the concepts before applying them. To avoid the influence of the narratives created by us in the user experience, we removed all examples of narratives from the user libraries of both versions of the system during the period of the experiment.

The story generation in Simplified ChatGeppetto is based on the process described in Section 4.2, but $P_{stories}$ and $P_{semiotic}$ are omitted from the initial prompt provided to the Storywriter AI Agent, and the description of P_{task} is replaced by "*Write the first event of a story about <SP>.*", where *<SP>* is the text of the story premise provided by the user.

To measure the user experience of our participants, we employed the User Experience Questionnaire (UEQ) as originally proposed by Laugwitz et al. [89]. The UEQ comprises 26 items focusing on six main concepts: (1) attractiveness, which is related to the users' overall impression of the system; (2) perspicuity, which measures how easy it is to get familiar with and learn how to use the system; (3) efficiency, which evaluates the effort required to use the system; (4) dependability, which measures whether users feel in control of the interaction; (5) stimulation, which assesses users' excitement and motivation while using the system; and (6) novelty, which is related to users' perception of how creative and innovative the system is. The items are measured on a 7-point semantic differential scale with two items presented with an opposite meaning.

To evaluate user satisfaction with the final story, we complemented the UEQ with 7 additional items derived from the Perceived Story Quality Index proposed by Baron and Bluck [90]. These items, measured on a 7-point semantic differential scale, comprise the following story adjectives: entertaining/boring, coherent/uncoherent, emotional/unemotional, memorable/unmemorable, original/unoriginal, creative/unimaginative, and engaging/unappealing.

We also added four open questions to assess what features of the system were the most useful and whether any features were difficult to use. Furthermore, we assessed the challenges and allowed participants to add any comments or feedback.

Finally, we collected usage data in the system to evaluate how users interact with the system. During the experiment, the system automatically registered all input data used to create stories (media type, semiotic operation, books/movies, premise, protagonist, and style of illustrations), all generated story data (titles, events, and images), and all user interaction data (time spent on each story, time spent on each event, number of deleted/regenerated events per story, number images regenerated per story, and number of suggestions per story).

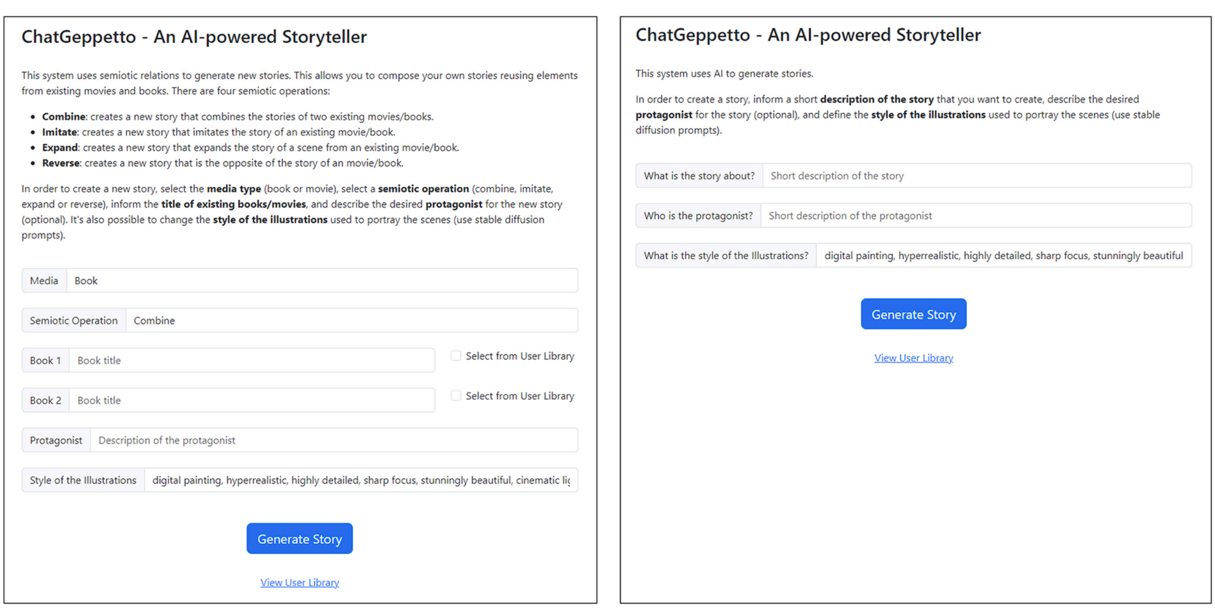

(a) (b)

Fig. 3. The initial screen of Full ChatGeppetto (a) and Simplified ChatGeppetto (b).

5.3. Procedure

Participants were invited to the study via the student portal of Breda University of Applied Sciences. They were informed to fill in the survey on a laptop or desktop PC. They first read the informed consent form, which provided them with information about the procedure and the data collection and storage policy. After this, they provided their informed consent for voluntary participation.

In the next screen of the survey, participants were asked to interact with one of the two versions of the system: Full ChatGeppetto or Simplified ChatGeppetto. They were instructed to create at least one story in collaboration with the system on any topic that interests them. The only requirement suggested is that the story should contain at least 8 events. They were allowed to interact longer with the system and create more stories if they wanted to explore the tool. The instructions were similar for both versions of the system.

After interacting with the system, participants were directed to a page where they had to fill in the questions concerning user experience, satisfaction with the final narrative, and open questions related to their opinions on the system's features.

Finally, some demographic questions were asked. First, they filled in their gender, age, and nationality. We also included additional questions to assess their story writing experience, and experience with Generative AI tools, such as ChatGPT.

At the end of the survey, they were thanked for participating in the study.

5.4. Results

We start the analysis of the results by assessing the usage data collected by the system. On average, participants who tested Full ChatGeppetto spent 12.24 min interacting with the system (SD = 6.65). Similarly, participants who tested Simplified ChatGeppetto spent 12.63 min interacting with the system (SD = 7.24). The number of stories created by each participant is also similar in both versions of the system. In Full ChatGeppetto, 72.2 % of participants created a single story during the experiment, 22.7 % of participants created two stories, and 4.5 % of participants created more than two stories. In Simplified ChatGeppetto, 70.0 % of participants created a single story during the experiment, 20.0 % of participants created two stories, and 10.0 % of participants created more than two stories. A total of 31 stories were generated in Full ChatGeppetto and a total of 29 stories were generated in Simplified ChatGeppetto.

The time participants spent per story is similar in both versions of the system: an average of 8.68 min in Full ChatGeppetto (SD = 4.51) and an average of 8.58 min in Simplified ChatGeppetto (SD = 5.14). The same can be observed in the number of events per story: while users of the Full ChatGeppetto generated stories with an average of 9.16 events per story (SD = 5.44), users of Simplified ChatGeppetto composed stories with an average of 7.20 events per story (SD = 3.70).

Regarding the data provided as input to the story generation process, users of Full ChatGeppetto explored different media types: 54.5 % of participants selected books as media source, 40.9 % of participants selected movies, and 4.5 % of participants tested both media types. Regarding the semiotic operations, most participants used only the *combine* operation (50.0 %), followed by the *imitate* and *expand* operations (each used by 18.2 % of participants). Only one participant (4.5 %) used the *reverse* operation. Two participants (9.1 %) explored more than one semiotic operation: one tested both *reverse* and *imitate* operations, and the other tested both *imitate* and *combine* operations.

The movies and books used as input by participants who tested Full ChatGeppetto during the experiment are:

- Movies: *Fear and Loathing in Las Vegas*, *Shrek*, *Titanic*, *Despicable Me*, *Avatar*, *Interstellar*, *John Wick: Chapter 4*, *Requiem for a Dream*, *Avengers: Endgame*, *How to Train Your Dragon*, *Die Hard*, *12 Angry Men*, *Back to the Future*, *The Hunger Games*, *Saving Private Ryan*, *Pirates of the Caribbean*, *Grease*, *Snow White*, *The Matrix*, *Barbie*, *The Lord of the Rings*, and *The Muppets*.
- Books: *In A New York Minute*, *The Course of Love*, *Boy Swallows Universe*, *Hell Divers*, *Harry Potter and the Goblet of Fire*, *Parable of the Sower*, *Don Quixote*, *The Eye of the Queen*, *Jane Eyre*, *The Last Summer*, *Kafka on the Shore*, *And Then There Were None*, and *The Lord of the Rings*.

The premises provided by users of Simplified ChatGeppetto vary from very simple concepts to well-elaborated story descriptions. While simple premises are very open, such as "*The world's worst nightmare*" or "*A man discovers a hidden room in his house*", the more elaborated descriptions provide detailed aspects of the story, such as "*The story needs to be about the struggles of a 25 year old, Gen-Z woman who struggles with modern society. She wants to keep up with all of the trends in the city and live a lavish life style, but in this modern society in combination with the high peer pressure and social media, it is hard to decide what her real desires are and what she thinks she wants because of peer pressure*". Most of the participants that tested Simplified ChatGeppetto (75.0 %), provided simple premises (with less than 100 characters).

We also analyzed some statistics related the user interaction to understand how users interacted with the system during the story composition process. Most participants of both versions of the system did not delete or regenerate any of the events generated by the system. In Full ChatGeppetto, 74.2 % of the stories had no event deletion/regeneration, 19.4 % had only one event deletion/regeneration, and 6.5 % had more than one event deletion/regeneration. In Simplified ChatGeppetto, and 82.8 % of the stories had no event deletion/regeneration, 6.9 % had only one event deletion/regeneration, and 10.3 % had more than one event deletion/regeneration. A similar behavior can be observed regarding image regenerations. In Full ChatGeppetto, 77.4 % of the stories had no image regenerations, 19.4 % had only one image regeneration, and 3.2 % had more than one image regeneration. In Simplified ChatGeppetto, 58.6 % of the stories had no image regenerations, 27.6 % had only one image regeneration, and 13.8 % had more than one image regeneration. Out of the image regenerations in Full ChatGeppetto, 45.5 % of them included user prompts. In Simplified ChatGeppetto, 43.8 % of the image regenerations included user prompts.

A more active user participation can be observed in the user suggestions for story events. In Full ChatGeppetto, 32.3 % of the generated stories contain one or two user suggestions, and 32.3 % contain three or more user suggestions. In Simplified ChatGeppetto, 27.6 % of the stories were created using one or two user suggestions, and 44.8 % contained three or more user suggestions.

After analyzing the statistical data registered by the system, we concentrated our attention on the data obtained with the UEQ. To generate a variable describing the concepts of the UEQ scale, we first measured the reliability of the scales. The attractiveness subscale was measured using 6 items ($\alpha = 0.86$), perspicuity was measured using 4 items ($\alpha = 0.70$), efficiency was measured using 4 items ($\alpha = 0.73$), dependability was measured using 4 items ($\alpha = 0.56$), stimulation was measured using 4 items ($\alpha = 0.86$), and finally novelty was also measured using 4 items ($\alpha = 0.74$). For each subscale, we created an average value of all items that were used to measure it, ranging from -3 to 3. The final results are shown in Fig. 4.

We tested the required assumptions (e.g. normality and homogeneity of variance) before using a *t*-test statistic to check for any significant differences between Full ChatGeppetto and Simplified ChatGeppetto. No assumptions were violated, so we were able to conduct the required tests. None of the t-tests executed showed a significant difference. All the relevant statistics are shown in Table 3.

User satisfaction of the final story was measured with seven items ($\alpha = 0.92$). We averaged all items to form a new variable, ranging from -3 to 3. All separate items and the newly created overall user satisfaction of

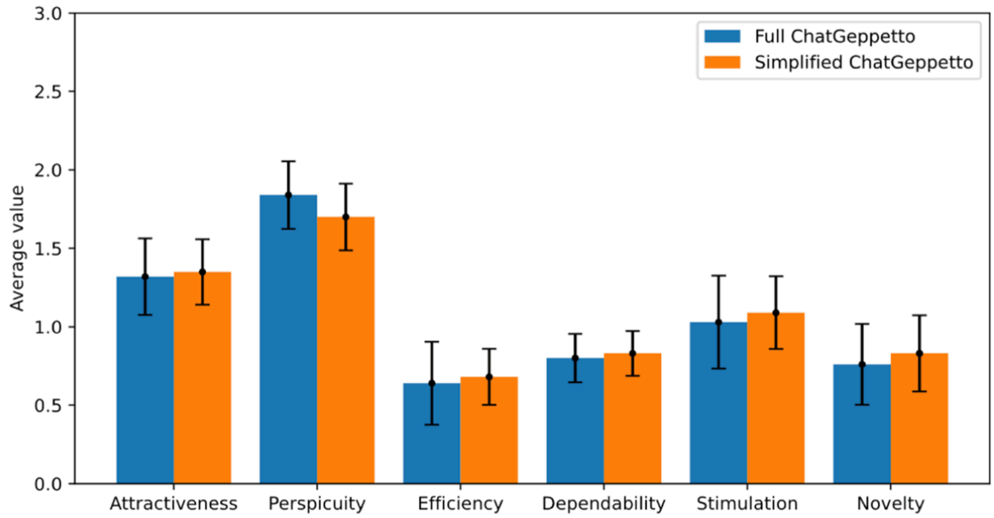

Fig. 4. Average values of all items of the UEQ for both versions of the system. The error bars represent Standard Errors.

Table 3
Summary Statistics of the UEQ Questionnaire.

UEQ variable	Full ChatGeppetto Mean (SD)	Simplified ChatGeppetto Mean (SD)	t-value	p-value
Attractiveness	1.32 (1.13)	1.35 (0.93)	t(40) = −0.12	0.90
Perspicuity	1.84 (1.01)	1.70 (0.95)	t(40) = 0.46	0.64
Efficiency	0.64 (1.24)	0.68 (0.80)	t(40) = −0.11	0.91
Dependability	0.80 (0.82)	0.83 (0.64)	t(40) = −0.12	0.90
Stimulation	1.03 (1.39)	1.09 (1.04)	t(40) = −0.14	0.89
Novelty	0.76 (1.21)	0.83 (1.09)	t(40) = −0.18	0.86

the final story are visualized in Fig. 5.

All necessary assumptions before executing a *t*-test were checked. The overall user satisfaction of the stories generated in collaboration with Full ChatGeppetto ($M = 0.82$, $SD = 1.49$) and of the stories generated in collaboration with Simplified ChatGeppetto ($M = 0.46$, $SD = 1.36$) did not differ significantly ($t(40) = 0.83$, $p = 0.41$).

Furthermore, all individual items were also compared with t-tests. We had to adjust the significance label to correct for multiple tests, using a Bonferroni correction. None of the comparisons were significantly different, but some main trends could be identified. All statistics are shown in Table 4.

After analyzing the responses to the UEQ and the user satisfaction with the final story, we moved to the analysis of the open questions, which gave us more insight into the best features of ChatGeppetto and the challenges participants faced while using it.

For both Full ChatGeppetto and Simplified ChatGeppetto a feature that was appreciated was the option to add suggestions for the continuation of the story. According to PP29: *"I like it that we can change the story ourselves […] just in case it goes on a wrong path"*. Specifically for Full ChatGeppetto, the semiotic operations were mentioned as a positive aspect. For example, PP26 mentioned: *"The combining of several elements together was cool to create a wild new story."*. For Simplified ChatGeppetto, the ease-of-use of the system was mentioned. According to PP24: *"Easy to generate and easy to write a simple prompt"*.

Additionally, the participants mentioned some challenges for both systems. The time required by the system to generate the events of the story was often mentioned as an issue. For example, P21 mentioned: *"The time it took to load got in the way of enjoying the story. Especially*

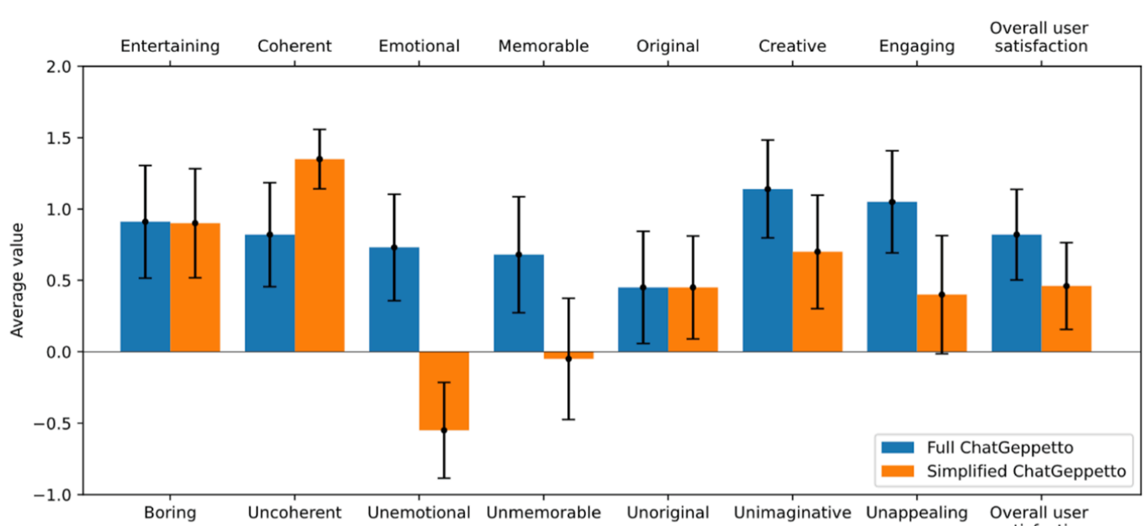

Fig. 5. Average values of all items to measure user satisfaction with the generated story for both versions of the system. The error bars represent Standard Errors.

Table 4
Summary Statistics of the individual items for user satisfaction with the final generated story.

UEQ variable	Full ChatGeppetto Mean (SD)	Simplified ChatGeppetto Mean (SD)	t-value	p-value
Boring-Entertaining	0.91 (1.85)	0.90 (1.71)	t(40) = −0.02	0.99
Coherent-Uncoherent	0.82 (1.71)	1.35 (0.93)	t(40) = 1.23	0.23
Emotional-Unemotional	0.73 (1.75)	−0.55 (1.50)	t(40) = −2.52	0.02
Memorable-Unmemorable	0.68 (1.91)	−0.05 (1.90)	t(40) = −1.24	0.22
Original-Unoriginal	0.45 (1.84)	0.45 (1.61)	t(40) = −0.01	0.99
Creative-Unimaginative	1.14 (1.61)	0.70 (1.78)	t(40) = −0.83	0.40
Engaging-Unappealing	1.05 (1.68)	0.40 (1.85)	t(40) = −1.19	0.24

because it creates very small pieces of the story every round.". For Full ChatGeppetto, one participant mentioned being unsure about the influence of the provided books in the generated narrative. PP10: *"I found it hard to determine how the reference books were represented in the final narrative. How would I see/notice that these books are of influence on the writing or the story."*. For Simplified ChatGeppetto, multiple participants found it challenging that the appearance of the characters changed in the images the system generated. For example, PP14 mentioned: *"The characters look different in the images"*. This might have been a more significant issue in the simplified version, because no universally known characters were used there. Furthermore, for Simplified ChatGeppetto, multiple participants mentioned that they thought the final stories were not creative. For example, PP15 mentioned: *"There should be more twists in the story to make it more interesting."*. Perhaps it was difficult for people to come up with their own storylines, while borrowing the initial structure from a known book or movie helps a person to create a more interesting storyline.

In both systems, some participants were unhappy with the images and were suggesting being able to remove that feature from the system. According to PP20: *"I don't really care about the images. If you just want to make a story without images, there should be an easy way to turn that off"*. On the other hand, some participants mentioned they liked the combination of text with pictures. For example, PP27 stated: *"I really like that it combines pictures with scenes"*.

5.5. Discussion

An important objective of this study was to assess user experience with ChatGeppetto through the UEQ. While statistically significant differences between the two versions of the system were not detected in the UEQ scores, it is important to consider that our sample size might have been a factor limiting the power of the statistical tests to detect subtle variations in user experience. It is also essential to notice that the UEQ scores for both systems leaned toward the positive end of the scale. This indicates that participants generally found interacting with both versions of ChatGeppetto to be attractive, perspicuous, efficient, dependable, stimulating, and novel. This overall positive user experience aligns well with the qualitative data obtained through open-ended questions. For instance, participants appreciated features like the ability to suggest story events in both systems, which allows them to exert control over the narrative flow and personalize the story to their preferences. In Full ChatGeppetto, the semiotic operations were highlighted as a positive aspect that facilitated storytelling, allowing users to manipulate existing stories in creative ways. For Simplified ChatGeppetto, ease-of-use was a key strength. Despite the absence of statistically significant differences, these findings suggest that ChatGeppetto offers a promising user experience for story creation.

Similar to the user experience findings, user satisfaction scores did not reveal a statistically significant difference between stories generated with the two versions of the system. However, by examining individual user satisfaction items and qualitative data, it is possible to highlight some trends that deserve discussion. The data suggests that Full ChatGeppetto might be able to create stories that are perceived as more emotional and memorable. This could be attributed to the incorporation of semiotic relations, which might guide the narrative towards more emotional and impactful moments. Additionally, qualitative data from the open-ended questions revealed that participants found stories created with Simplified ChatGeppetto to be less interesting. This aligns with the observed trends of lower user ratings for memorability and emotional impact. The lack of emotional engagement and memorable elements in Simplified ChatGeppetto stories could be a contributing factor to users perceiving them as less interesting.

A relevant contrast emerged regarding story coherence. While user satisfaction scores indicated a slight preference for coherence in Simplified ChatGeppetto stories, open-ended responses and the input data registered by the system suggest a possible explanation. In Simplified ChatGeppetto, participants were able to provide open-ended story premises, potentially allowing the system more flexibility to generate internally coherent narratives. Conversely, Full ChatGeppetto relies on referencing existing stories, which might introduce a challenge in user perception of coherence. Even though the system utilizes semiotic operations to reuse the structure of existing stories, the resulting narrative structure might not perfectly align with how users envisioned the reinterpretation of the original stories.

The study revealed mixed user opinions regarding the image generation aspect of ChatGeppetto. While some participants expressed a desire to remove images entirely, others valued the combination of text and images for storytelling. This positive perception of the text-image pairing aligns with the finding that some users appreciated the system's ability to generate illustrations alongside the narrative. Additionally, some participants pointed out inconsistencies in character appearance across the generated images, which could be particularly problematic when dealing with stories that lack established characters.

6. Concluding remarks

The contributions of this article are twofold. As a theoretical result, in the area of computational narratology, we mention our semiotic reconstruction approach to interactive story composition with the participation of AI agents. Aiming at a practical result, we implemented a prototype based on the proposed theory, with the main purpose of helping users not necessarily trained as professional writers to compose engaging stories as well-structured variants of popular book or movie narratives. Our claim concerning the *semiotic completeness* of our semiotic relations model (cf. Section 3.1.1), extensively explored in our previous works, motivated our confidence on the prototype's potential to cover an ample variety of narrative genres.

To test whether this overarching practical purpose was met, we conducted a user study with participants interacting with two versions of our prototype: one that incorporates semiotic relations based on the proposed theoretical framework, and another that does not use semiotic relations. The results of the study provided valuable insights into the overall user experience and satisfaction with different aspects of the system, including perceptions of story quality. While statistically significant differences in user experience and overall satisfaction with the generated stories were not observed, positive trends emerged. Participants appreciated several functionalities in both systems, such as the ability to suggest story continuations. The version that incorporates the proposed theoretical framework was specifically praised for its use of semiotic operations as a valuable feature for creative storytelling. The data also suggests that stories created using semiotic operations might be perceived as more emotional and memorable. The evaluation of elements related to story quality offered a preliminary glimpse into how users perceived different dimensions of narrative creativity and coherence. These findings provide initial support for the theoretical foundation of the semiotic reconstruction approach and its potential to enhance user experience and story quality in interactive narrative generation. With the continuation of our project, we expect to benefit from an even ampler and more widely diversified audience, by encouraging eventual readers of the present article to play with the prototype, which can be freely accessed, via the link provided in Section 4.5.

The integration of AI into creative processes, as exemplified by our prototype, raises several ethical considerations that require careful attention. As AI-generated narratives become more sophisticated, it becomes increasingly essential to disclose the role of AI in the creative process to maintain transparency and avoid misleading audiences [83]. This concern is not merely a matter of ethics but also a potential legal issue, as the question of who holds the rights to AI-generated creations remains a subject of ongoing debate [84,85]. Furthermore, while our approach seeks to empower users and enhance creativity, it also intersects with complex issues regarding authorship and originality. The extent to which we consider AI-generated narratives original and how deep we understand their implications in traditional notions of authorship are open questions requiring further exploration in law and ethics. Fisher [86] investigates an interesting human-centric approach to these questions. Moreover, as we designed our system to augment, not replace, human creativity, it is crucial to consider the broader societal implications of AI's growing role in creative fields. The increasing automation of creative tasks through AI raises concerns about the potential displacement of human creators and the devaluing of human creative labor, as discussed by Caporusso [87] and Zarifhonarvar [88]. It is our responsibility to address these ethical considerations, not just for the responsible development and deployment of AI in creative fields but also to foster public trust and mitigate potential negative impacts on individual creators and society.

Besides stressing that creating a narrative justifies itself as a pleasurable entertainment experience, we should add that a prototype like ours can serve in a variety of other applications, always in a helper capacity, avoiding any unethical and truly absurd claims to supersede human agents. A straightforward application is as a teaching resource to train language proficiency and to develop literary skills. Another instigating educational application, suggested by our example with William Shirer's book (Section 3.2), is to exercise counterfactual "what if" speculation over serious historical narratives.

While the current implementation of ChatGeppetto utilizes the generative capabilities of ChatGPT and Stable Diffusion, it is important to emphasize that the architecture of our system is not intrinsically tied to these specific models. The modular plugin design explicitly addresses the dynamic nature of artificial intelligence advancements. As newer text generation models or image synthesis techniques emerge and demonstrate superior performance, they can be seamlessly integrated into the system by replacing the existing AI agents. This ensures that our approach remains adaptable and can continually benefit from the rapid progress in the field of artificial intelligence.

Our current implementation demonstrates the functionality and effectiveness of the proposed approach within a controlled environment. However, we recognize the importance of addressing scalability and performance considerations for production environments. Currently, narrative generation time is inherently dependent on the responsiveness of the OpenAI API, which is designed to scale with user demand [92–94]. However, further investigation is needed to assess potential bottlenecks and optimization strategies within our system's integration with the API. For image generation, our system utilizes a privately hosted Stable Diffusion model. Recent research has demonstrated the potential for scaling such models on cloud infrastructure [95,96], suggesting viable solutions for ensuring efficient image generation even with increased user traffic. However, as pointed out by Xu et al. [95], the scalability of foundation models for large-scale production deployment remains an open research problem.

A key consideration in our approach is the reliance on existing narratives as a source of inspiration for new stories. While this allows for creative reinterpretations and remixes, it also inherently carries the risk of perpetuating biases present in those source materials. Biases in character representation, plot structures, and thematic elements may inadvertently be reinforced in the generation of new stories. However, this limitation can be mitigated with the use of the antithetic relation (reverse operation), which allows for the creation of narratives that intentionally challenge or invert the motifs, structure, or stereotypes found in the source material. The possibility of selecting a story generated by other semiotic operations – conventionally stored in the user library – as the input for the reverse operation also extends the possibility of generating counter-narratives using the other semiotic operations. Nonetheless, the complex issue of bias in AI-generated content remains an ongoing challenge in the field of AI that requires further investigation.

Another important consideration in our approach is that any LLM model, including GPT-4, is inherently limited by its training data. It is worth noting that GPT-4 has demonstrated a remarkable ability to understand and generate stories within a wide variety of genres. In our recent work [97], we explored the use of genre patterns to guide story generation using GPT-4, revealing the model's capacity to adapt to diverse narrative structures, thematic conventions, and genres. While the extensiveness of GPT-4's training data allows for considerable diversity in generated narratives, limitations may arise when dealing with niche titles or very recent works that fall outside the model's knowledge cutoff date (currently, April 2023).[24] Therefore, further research is needed to explore strategies for incorporating more up-to-date and diverse data to further enhance the storytelling capabilities of the proposed approach.

Future work may aim at further expanding the proposed semiotic reconstruction approach, and, simultaneously, to pursue more comprehensive user studies. As a research direction, one may investigate how to extend the proposed approach to face the theoretical and practical problems of adapting book narratives to movie screenplays, taking into consideration the different characteristics and requirements of the two media. As an equally seductive objective, one may concentrate on the creation of stimulating and deeply immersive stories to be integrated into existing game development workflows. To improve the generalizability of the user evaluation, future studies may consider a larger and more heterogeneous user base, perhaps considering the target users of extended versions of the proposed system. This can encompass professional screenplay writers or videogame story writers, whose feedback may provide more focused perspectives on the creative potential of the proposed system, particularly regarding their opinion on the use of semiotic operations for narrative generation in their area.

Improving specific components of the proposed system also represents interesting directions for future research. Our user study revealed both user appreciation for the text-image combination and concerns regarding image consistency across different scenes. In this context, future efforts may target optimizing the image generation process. A promising solution involves the development of a library with detailed descriptions of characters and their physical attributes. By allowing users to select the protagonists of their stories from this set of predefined characters, the system would be able to provide more detailed descriptions of the scenes to the Stable Diffusion model, which can possibly enhance character consistency across images within the same story.

Another captivating direction for future research involves the

[24] https://openai.com/index/new-models-and-developer-products-announced-at-devday/.

exploration of the recent advancements in text-to-video models, such as OpenAI's Sora model.[25] By integrating a text-to-video component, the system could not only generate textual narratives with static illustrations but also produce corresponding animated video sequences for each story event. This would transform ChatGeppetto into an actual movie-generation tool.

CRediT authorship contribution statement

Edirlei Soares de Lima: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Margot M.E. Neggers:** Writing – review & editing, Writing – original draft, Validation, Methodology, Formal analysis, Data curation. **Bruno Feijó:** Writing – review & editing, Conceptualization. **Marco A. Casanova:** Writing – review & editing, Conceptualization. **Antonio L. Furtado:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

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.

Appendix A

In order to demonstrate the creative capabilities of ChatGeppetto, this appendix describes all the steps performed to create the story "Holmes & Hobbit: The Intersection of Worlds". This story was generated based on the following parameters:

- **Media:** Movie
- **Semiotic Operation:** Combine
- **Movie 1:** Sherlock Holmes (2009)
- **Movie 2:** The Hobbit: An Unexpected Journey
- **Protagonist:** Sherlock Holmes
- **Style of the Illustrations:** digital painting, hyperrealistic, highly detailed, sharp focus, stunningly beautiful, cinematic lighting

All the steps performed by the user and the resulting events and illustrations for this story are presented in Table 5. The same composition steps – as registered in our system – are available at: https://narrativelab.org/chatgeppetto/log.php?id=548.

Table 5
Composition steps of the story "Holmes & Hobbit: The Intersection of Worlds".

Steps	Events	Scene Illustrations	User Actions
Step 1	**Event:** *Sherlock Holmes, after solving a strange case involving peculiar symbols and an antique ring, stumbles upon an ancient portal in the basement of a dusty old bookstore in London.* **Scene:** *A dimly lit, dusty room overflowing with books, in the midst of which a surprised Sherlock Holmes stands before a shimmering, circular gateway emitting a mystical glow.*		**Action:** Continue Story
Step 2	**Event:** *Holmes steps through the portal, finding himself in the fantastical world of Middle Earth, where he encounters Bilbo Baggins, lost and confused in Mirkwood Forest.* **Scene:** *Astonished Bilbo Baggins and Sherlock Holmes meet in the dense, mystical forest, under a canopy of towering trees and amidst the twirls of mist.*	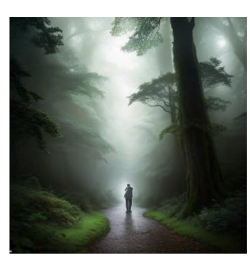	**Action:** Regenerate Image

(continued on next page)

[25] https://openai.com/sora.

Table 5 (*continued*)

Steps	Events	Scene Illustrations	User Actions
Step 3			**Action:** Continue Story **Suggestion:** Bilbo Baggins will tell Sherlock Holmes that a crime has happened
Step 4	**Event:** *Bilbo Baggins quickly enlists Holmes' help, revealing a crime: the treasured Arkenstone has been stolen from under the vigilant watch of Smaug, the dragon.* **Scene:** *A worried Bilbo Baggins talking to a bemused Sherlock Holmes, with a backdrop of thick forest and a crude map sketching the whereabouts of the dragon's lair, spread out on a makeshift table made from a tree stump.*		**Action:** Continue Story
Step 5	**Event:** *Holmes, intrigued by the high-stakes theft, promises to help Bilbo and together they begin to search for clues in the heart of Mirkwood Forest, leading them towards the lair of Smaug.* **Scene:** *Sherlock Holmes, on his knees, examining the ground while Bilbo looks on nervously, their journey illuminated by the dappled sunlight piercing the dense foliage overhead.*		**Action:** Regenerate Image
Step 6		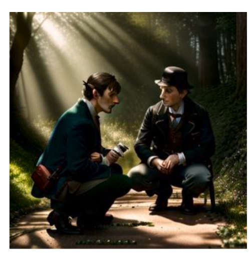	**Action:** Continue Story
Step 7	**Event:** *As they approach the dragon's lair, they discover a trail of peculiar footprints different from any Middle Earth creature, leading them to suspect an outsider might be involved.* **Scene:** *A close-up of Holmes' hand, holding a magnifying glass over an exotic footprint imprinted in the smoldering ground near Smaug's lair, with a watchful Bilbo nearby.*		**Action:** Regenerate Image **Extra Prompt:** ((((Sherlock Holmes checking a footprint)))), (((((wide shot camera angle)))))
Step 8		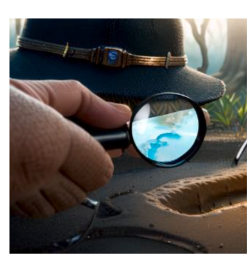	**Action:** Regenerate Image

(*continued on next page*)

Table 5 (*continued*)

Steps	Events	Scene Illustrations	User Actions
Step 9			**Action:** Continue Story
Step 10	**Event:** *They stealthily infiltrate the lair of Smaug, using a secret passage discovered with the aid of Sherlock's deductive reasoning and Bilbo's knowledge of the terrain.* **Scene:** *A suspenseful scene showing Holmes and Bilbo cautiously navigating a narrow, winding and dark pathway towards a massive, fiery cavern in the heart of the mountain.*		**Action:** Continue Story
Step 11	**Event:** *In the lair, they find a dropped handkerchief monogrammed with a "M", a clue suggesting that this crime extends beyond the realm of Middle Earth.* **Scene:** *A dramatic shot of Sherlock Holmes meticulously examining a monogrammed handkerchief, its white fabric stark against piles of golden treasure, with a narrow beam of light illuminating the scene.*		**Action:** Regenerate Image
Step 12			**Action:** Continue Story
Step 13	**Event:** *Realizing the gravity of the situation, Holmes and Bilbo decide to confront Smaug about the theft, approaching the mighty dragon with a mixture of caution and determination.* **Scene:** *A wide shot of Smaug's massive, awe-inspiring form being approached by the seemingly small figures of Sherlock Holmes and Bilbo, his fiery eyes gleaming in the dark cavern.*	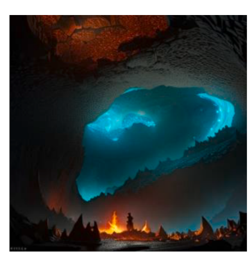	**Action:** Regenerate Image **Extra Prompt:** (((((dragon)))))
Step 14			**Action:** Continue Story

(*continued on next page*)

Table 5 (*continued*)

Steps	Events	Scene Illustrations	User Actions
Step 15	**Event:** *Holmes, with his characteristic aplomb, manages to talk Smaug into revealing that he had seen a strange man, clad in Victorian attire, at the lair the day the Arkenstone went missing.* **Scene:** *A riveting tableau of Sherlock Holmes engaged in conversation with Smaug, the massive dragon, amidst a myriad of scattered treasure, with a visibly tense Bilbo Baggins looking on.*		**Action:** Regenerate Image
Step 16			**Action:** Continue Story
Step 17	**Event:** *Piecing together the clues, Holmes deducts that the intruder is none other than his arch-nemesis, Moriarty, who has somehow found his way to Middle Earth and stolen the Arkenstone.* **Scene:** *Sherlock Holmes, in deep thought, stands against the backdrop of the majestic dragon, linking the dots marked on a parchment, with a profile sketch of Moriarty casting a sinister shadow.*		**Action:** Regenerate Image
Step 18			**Action:** Continue Story
Step 19	**Event:** **Scene:**.	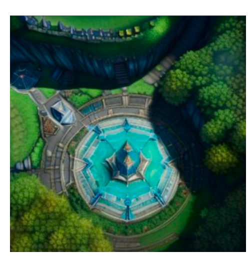	**Action:** Regenerate Event
Step 20	**Event:** *With the revelation, Holmes and Bilbo, now with a clear adversary, set off on a quest to find Moriarty and retrieve the Arkenstone before it's used for nefarious purposes.* **Scene:** *Holmes and Bilbo, silhouetted against the setting sun, leave the dragon's lair, embarking on their dangerous mission, with a foreboding sense of adventure enveloping the scene.*	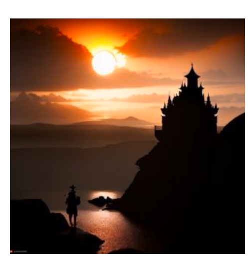	**Action:** Regenerate Image **Extra Prompt:** ((((Sherlock Holmes and Bilbo Baggins))))

(*continued on next page*)

Table 5 (*continued*)

Steps	Events	Scene Illustrations	User Actions
Step 21			**Action:** Continue Story
Step 22	**Event:** *As they journey through Middle Earth, Holmes trains Bilbo in the art of deduction and survival, preparing him for the inevitable confrontation with Moriarty.* **Scene:** *A montage of shots showing Holmes and Bilbo in various scenarios – studying maps, examining tracks, practicing combat – against the sprawling landscapes and diverse terrains of Middle Earth.*		**Action:** Continue Story
Step 23	**Event:** *The duo's journey leads them to the Elves of Rivendell, who reveal that Moriarty has last been seen heading towards the haunted mountains of Angmar.* **Scene:** *A tranquil scene of Holmes and Bilbo in conversation with regal Elves in the enchanting city of Rivendell, lit under a moonlit sky with the ominous silhouette of the mountains looming in the distance.*		**Action:** Continue Story
Step 24	**Event:** *With the new lead, Holmes and Bilbo venture towards Angmar, braving the treacherous terrains and battling supernatural creatures which guard the haunted mountains.* **Scene:** *The moonlight casting long, monstrous shadows as Holmes and Bilbo fight off grotesque, shadowy figures on the rocky slopes of Angmar, their faces a mix of focus and determination.*	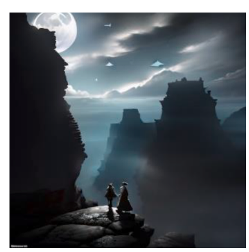	**Action:** Continue Story **Suggestion:** Holmes and Bilbo will confront Moriarty
Step 25	**Event:** *Reaching the peak of Angmar, Holmes and Bilbo finally confront Moriarty, who stands smirking with the shining Arkenstone in his hand, ready for an epic showdown.* **Scene:** *The silhouette of Moriarty standing at the edge of a cliff, the Arkenstone radiating an eerie glow in his hand, with a resolute Holmes and Bilbo confronting him against a tumultuous sky.*		**Action:** Continue Story
Step 26	**Event:** *An intense battle ensues between Holmes, Bilbo, and Moriarty, their moves a dance of strategy and wit on the precipice of the haunted mountain.* **Scene:** *A dynamic shot of the trio locked in combat, their figures dramatically outlined against the raw fury of an approaching storm, with lightning splitting the sky above the mountains.*	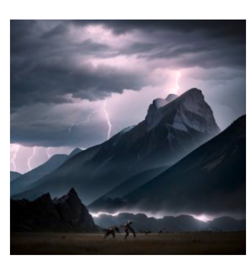	**Action:** Continue Story

(continued on next page)

Table 5 (*continued*)

Steps	Events	Scene Illustrations	User Actions
Step 27	**Event:** *Using a coordinated plan, Holmes distracts Moriarty with a feint while Bilbo, applying his newly learned skills, manages to snatch the Arkenstone from Moriarty's grasp.* **Scene:** *A high-tension image of Bilbo deftly swiping the glowing Arkenstone as Holmes and Moriarty engage in a fierce duel, sparks flying from their fast-moving bodies against a dark stormy sky.*		**Action:** Continue Story
Step 28	**Event:** *With the Arkenstone secured and Moriarty disarmed, Holmes pulls out a pocket watch, a concealed device that opens a portal back to their world, ready to take Moriarty into custody.* **Scene:** *An arresting image of a triumphant Holmes pressing the watch, activating a portal that swirls open amidst the chaos, casting an ethereal glow on the battle-worn faces of our heroes.*		**Action:** Regenerate Image
Step 29			**Action:** Continue Story
Step 30	**Event:** *Moriarty, Holmes, and Bilbo make a hasty retreat back to London via the portal, leaving Middle Earth behind with the promise of the Arkenstone's safe return.* **Scene:** *A stunning portrayal of the three figures leaping into the swirling portal, their forms blurring as they shift from the rugged mountains of Angmar to the familiar cobblestone streets of Victorian London.*	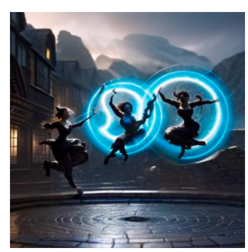	**Action:** Regenerate Image **Extra Prompt:** (((Moriarty))) (((Sherlock Holmes))) (((Bilbo Baggins)))
Step 31			**Action:** Continue Story
Step 32	**Event:** *Arriving in London, Moriarty is swiftly handed over to the local authorities while Holmes and Bilbo share a heartfelt farewell, promising to remember their extraordinary adventure.* **Scene:** *A poignant scene of Holmes and Bilbo shaking hands under a gas-lit lamp post, with Moriarty being taken away in the background by Scotland Yard, and the closed portal flickering out of existence.*		**Action:** Continue Story

(*continued on next page*)

Table 5 (*continued*)

Steps	Events	Scene Illustrations	User Actions
Step 33	**Event:** *Holmes uses his pocket watch to open a portal back to Middle Earth, allowing Bilbo to return the Arkenstone to its rightful place and reaffirm peace in his world.* **Scene:** *A final poignant image of Bilbo, holding the glittering Arkenstone, stepping into the magical portal with a final wave at Holmes, the warm glow of the portal enveloping him while the iconic London skyline fades behind.*		**Action:** Regenerate Image
Step 34			**Action:** Regenerate Image **Extra Prompt:** ((((a small magic portal over London streets)))) ((((Bilbo Baggins))))
Step 35			**Action:** Finish Story

References

[1] A.L. Furtado, M.A. Casanova, S.D.J. Barbosa, A semiotic approach to conceptual modelling, in: E. Yu, G. Dobbie, M. Jarke, S. Purao, (Eds.), Conceptual Modeling, ER 2014, Lecture Notes in Computer Science 8824 (2014) 1-12. Doi: 10.1007/978-3-319-12206-9_1.

[2] E.S. de Lima, B. Feijó, M.A. Casanova, A.L. Furtado, Storytelling variants based on semiotic relations, Entertain. Comput. 17 (2016) 31–44, https://doi.org/10.1016/j.entcom.2016.08.003.

[3] E. S. de Lima, A. L. Furtado, B. Feijó, Storytelling Variants: The Case of Little Red Riding Hood, in: K. Chorianopoulos, M. Divitini, J, Baalsrud Hauge, L. Jaccheri, R. Malaka (eds.), Entertainment Computing - ICEC 2015, ICEC 2015, Lecture Notes in Computer Science 9353 (2015) 286-300. Doi: 10.1007/978-3-319-24589-8_22.

[4] E.S. de Lima, A.L. Furtado, B. Feijó, Types, Motifs and the Emergence of Variants, in: Proceedings of the 14th Brazilian Symposium on Games and Digital Entertainment, 2015, pp. 295–303.

[5] R. Barthes, Theory of the Text, in: R. Young (Ed.), Untying the Text: A Post-Structuralist Reader, Routledge & Kegan Paul, Oxfordshire, 1981.

[6] J. Culler, On Deconstruction: Theory and Criticism After Structuralism, Cornell University Press, New York, 2007.

[7] K. Burke, A Grammar of Motives, University of California Press, Oakland, 1969.

[8] F. Saussure, Course in General Linguistics, Open Court Publishing, Chicago, 1998.

[9] C.S. Peirce, Principles of Philosophy, Harvard University Press, Cambridge, 1931.

[10] J. Kolodner, Case-Based Reasoning, Morgan Kaufmann Publishers, Cambridge, 2014.

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

[12] W. Iser, The Implied Reader: Patterns of Communication in Prose Fiction from Bunyan to Beckett, Johns Hopkins University Press, Baltimore, 1974.

[13] U. Eco, A Theory of Semiotics, Indiana University Press, Bloomington, 1976.

[14] U. Eco, The Limits of Interpretation, Indiana University Press, Bloomington, 1994.

[15] U. Eco, The Open Work, Harvard University Press, Cambridge, 1989.

[16] G. Fauconnier, M.B. Turner, Conceptual projection and middle spaces, Technical Report 9401, University of California, San Diego (1994), https://doi.org/10.2139/ssrn.1290862.

[17] I. Asimov, I, Robot, Gnome Press, New York, 1950.

[18] E.S. de Lima, M.A. Casanova, B. Feijo, A.L. Furtado, ChatGeppetto - an AI-powered Storyteller, in: Proceedings of the 22nd Brazilian Symposium on Games and Digital Entertainment, 2023, pp. 28–37, https://doi.org/10.1145/3631085.3631302.

[19] E. S. de Lima, B. Feijo, M. A. Casanova, A. L. Furtado, Semiotic Structuring in Movie Narrative Generation, in: P. Ciancarini, A. Di Iorio, H. Hlavacs, F. Poggi, (Eds.), Entertainment Computing – ICEC 2023, ICEC 2023, Lecture Notes in Computer Science 14455 (2023) pp. 161-175. Doi: 10.1007/978-981-99-8248-6_13.

[20] P. Gervás, B. Díaz-Agudo, F. Peinado, R. Hervás, Story plot generation based on CBR, in: A. Macintosh, R. Ellis, T. Allen (Eds.), Applications and Innovations in Intelligent Systems XII, SGAI 2004, (2004), pp. 33-46. Doi: 10.1007/1-84628-103-2_3.

[21] P. Gervás, Computational approaches to storytelling and creativity, AI Mag. 30 (3) (2009) 49–62, https://doi.org/10.1609/aimag.v30i3.2250.

[22] E. S. de Lima, B. Feijó, A. L. Furtado, Computational Narrative Blending Based on Planning, in: Entertainment Computing – ICEC 2021, Lecture Notes in Computer Science 13056 (2021), 289–303. Doi: 10.1007/978-3-030-89394-1_22.

[23] S. Värtinen, P. Hämäläinen, C. Guckelsberger, Generating role-playing game quests with GPT language models, IEEE Trans. Games (2022) 1–12, https://doi.org/10.1109/TG.2022.3228480.

[24] A. Fan, M. Lewis, Y. Dauphin, Hierarchical Neural Story Generation, in, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, 2018, pp. 889–898, https://doi.org/10.18653/v1/P18-1082.

[25] A. Alabdulkarim, W. Li, L.J. Martin, M.O. Riedl, Goal-Directed Story Generation: Augmenting Generative Language Models with Reinforcement Learning, in: arXiv: 2112.08593, 2021. Doi: 10.48550/arXiv.2112.08593.

[26] P. Xu, M. Patwary, M. Shoeybi, R. Puri, P. Fung, A. Anandkumar, B. Catanzaro, 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, 2020, pp. 2831–2845, https://doi.org/10.18653/v1/2020.emnlp-main.226.

[27] S. Wang, G. Durrett, K. Erk, Narrative Interpolation for Generating and Understanding Stories, in: arXiv:2008.07466, 2020. Doi: 10.48550/arXiv.2008.07466.

[28] L. Castricato, S. Frazier, J. Balloch, M. Riedl, Tell me a story like i'm five: story generation via question answering. Proceedings of the 3rd Workshop on Narrative Understanding, 2021.

[29] D. Yang, Y. Zhou, Z. Zhang, T. J. J. Li, L. C. Ray, AI as an Active Writer: Interaction strategies with generated text in human-AI collaborative fiction writing, in: Joint Proceedings of the ACM IUI 2022 Workshops, 2022, pp. 56-65.
[30] J. Xiang, Z. Zhao, M. Zhou, M. McKenzie, A. Kilayko, I, C. Macbeth, S. Carter, K. Sieck, M. Klenk, Interleaving a Symbolic Story Generator with a Neural Network-Based Large Language Model, in: Proceedings of Tenth Annual Conference on Advances in Cognitive Systems, 2022.
[31] A. Yuan, A. Coenen, E. Reif, D. Ippolito, Wordcraft: Story Writing With Large Language Models, in: 27th International Conference on Intelligent User Interfaces, 2022, pp. 841-852. Doi: 10.1145/3490099.3511105.
[32] K. Yang, Y. Tian, N. Peng, D. Klein, Re3: Generating Longer Stories With Recursive Reprompting and Revision, in: rXiv:2210.06774, 2022. Doi: 10.48550/arXiv.2210.06774.
[33] H. Chu, S. Liu, Can AI tell good stories? Narrative Transportation and Persuasion with ChatGPT, PsyArXiv (2023), https://doi.org/10.31234/osf.io/c3549.
[34] E. C. Garrido-Merchán, J. L. Arroyo-Barrigüete, R. Gozalo-Brizuela, Simulating H. P. Lovecraft horror literature with the ChatGPT large language model, in: arXiv: 2305.03429, 2023. Doi: 10.48550/arXiv.2305.03429.
[35] M. C. Gursesli, P. Taveekitworachai, F. Abdullah, M. F. Dewantoro, A. Lanata, A. Guazzini, V. K. Lê, A. Villars, R. Thawonmas, The Chronicles of ChatGPT: Generating and Evaluating Visual Novel Narratives on Climate Change Through ChatGPT, in: L. Holloway-Attaway, J. T. Murray (Eds.), Interactive Storytelling. ICIDS 2023, Lecture Notes in Computer Science 14384 (2023) 181-194. Doi: 10.1007/978-3-031-47658-7_16.
[36] K. Xie, M. Riedl, Creating Suspenseful Stories: Iterative Planning with Large Language Models, in: arXiv:2402.17119, 2024. Doi: 10.48550/arXiv.2402.17119.
[37] V. Kumaran, J. Rowe, B. Mott, J. Lester, SCENECRAFT: automating interactive narrative scene generation in digital games with large language models, in: Proceedings of the Nineteenth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, 2023, Article 9, pp. 86-96. Doi: 10.1609/aiide.v19i1.27504.
[38] V. Schetinger, S. D. Bartolomeo, E. S. de Lima, C. Meinecke, R. Rosa, n Walks in the Fictional Woods, in: arXiv:2308.06266, 2023. Doi: 10.48550/arXiv.2308.06266.
[39] T. Ashby, B. K. Webb, G. Knapp, J. Searle, N. Fulda, Personalized Quest and Dialogue Generation in Role-Playing Games: A Knowledge Graph- and Language Model-based Approach, in: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems, 2023, Article 290, pp. 1-20. Doi: 10.1145/3544548.3581441.
[40] S. Al-Nassar, A. Schaap, M. V. D. Zwart, M. Preuss, M. A. Gómez-Maureira, QuestVille: Procedural Quest Generation Using NLP Models, in: Proceedings of the 18th International Conference on the Foundations of Digital Games, 2023, Article 50, pp. 1-4. Doi: 10.1145/3582437.3587188.
[41] S. Johnson-Bey, M. Mateas, N. Wardrip-Fruin, Toward Using ChatGPT to Generate Theme-Relevant Simulated Storyworlds, in: Proceedings of the Experimental Artificial Intelligence in Games Workshop, 19th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, 2023.
[42] T. S. Wang, A. S. Gordon, Playing Story Creation Games with Large Language Models: Experiments with GPT-3.5, in: L. Holloway-Attaway, J. T. Murray (Eds.) Interactive Storytelling, ICIDS 2023, Lecture Notes in Computer Science 14384 (2023) 297-305. Doi: 10.1007/978-3-031-47658-7_28.
[43] P. Taveekitworachai, F. Abdullah, M. C. G., M. F. Dewantoro, S. Chen, A. Lanata, A. Guazzini, R. Thawonmas, What Is Waiting for Us at the End? Inherent Biases of Game Story Endings in Large Language Models, in: L. Holloway-Attaway, J. T. Murray (Eds.), Interactive Storytelling, ICIDS 2023, Lecture Notes in Computer Science 14384 (2023) 274–284. Doi: 10.1007/978-3-031-47658-7_26.
[44] U. Eco, Articulations of the cinematic code, in: B. Nichols (Ed.), Movies and Methods, University of California Press, Oakland, 1976, pp. 590–607.
[45] C. Metz, Language and Cinema, De Gruyter Mouton, Boston, 1974.
[46] F. Karsdorp, P. van Kranenburg, T. Meder, D. Trieschnigg, A. van den Bosch, In search of an appropriate abstraction level for motif annotations, in: Proceedings of the Third Workshop on Computational Models of Narrative, 2012, pp. 22-28.
[47] J. Culler, Fabula and Sjuzhet in the analysis of narrative, Poetics Today 1 (3) (1980) 27–37, https://doi.org/10.2307/177240.
[48] D. Chandler, Semiotics: The Basics, Routledge, Oxfordshire, 2002.
[49] J. Culler, The Pursuit of Signs: Semiotics, Literature, Deconstruction, Routledge, Oxfordshire, 1981.
[50] H. White, Tropics in Discourse: Essays in Cultural Criticism, Johns Hopkins University Press, Baltimore, 1978.
[51] A.L. Furtado, Semiotic Relations and Proof Methods, Monografias em Ciência da Computação, Technical Report 18/11, PUC-Rio, 2011.
[52] R. Jakobson, Two Aspects of Language and Two Types of Aphasic Disturbances, in: R. Jakobson, M. Halle (Eds.), Fundamentals of Language, Mouton, Berlin, 1956, pp. 69–90.
[53] M.E. Winston, R. Chaffin, D. Herrmann, A taxonomy of part-whole relations, Cognit. Sci. 11 (4) (1987) 417–444, https://doi.org/10.1207/s15516709cog1104_2.
[54] Chrétien de Troyes, The Complete Romances of Chretien de Troyes, D. Staines (trans.), Indiana University Press, Bloomington, 1991.
[55] N. Lacy, The Old French Vulgate & Post-Vulgate Cycles in Translation, Routledge, Oxfordshire, 2010.
[56] P. Lagerkvist, Barabbas, A. Blair (trans.), Vintage Books, London, 1989.
[57] W.L. Shirer, The Rise and Fall of the Third Reich, Simon & Schuster, New York, 2011.
[58] J. Zipes, Beauties, Beasts and Enchantment: Classic French Fairy Tales, Crescent Moon Publishing, Kent, 2016.
[59] H.C. Andersen, Hans Christian Andersen Complete Fairy Tales, Canterbury Classics, 2014.
[60] Z. Sun, G. Finnie, K. Weber, Abductive Case-Based Reasoning, Int. J. Intell. Syst. 20 (9) (2005) 957–983, https://doi.org/10.1002/int.20101.
[61] K. Satoh, Translating Case-Based Reasoning into Abductive Logic Programming, in: 12th European Conference on Artificial Intelligence, 1996, pp. 142-146.
[62] U. Eco, T.A. Sebeok (Eds.), The Sign of the Three: Dupin, Holmes, Peirce, Indiana University Press, Bloomington, 1988.
[63] A.C. Doyle, A Study in Scarlet, Ward Lock & Co, London, 1888.
[64] T. Todorov, The Poetics of Prose, Blackwell, Hoboken, 1977.
[65] P. Squire, J. Cohn, D. Nicholson, M. Nolan, P. J. Reber, D. Oudiette, J. Niehaus, A. Geyer, L. O'Neill, Towards Enhancing Intuitive Decision Making Through Implicit Training, in: Proceedings of the Interservice/Industry Training, Simulation, and Education Conference, 2014, Paper No. 14253.
[66] A. Christie, The Mysterious Affair at Styles, John Lane, 1920.
[67] C.R. Gallistel, L.D. Matzel, The neuroscience of learning: beyond the Hebbian Synapse, Annu. Rev. Psychol. 64 (2013) 169–200, https://doi.org/10.1146/annurev-psych-113011-143807.
[68] P. J. Reber, M. Beeman, K. A. Paller, Human Memory Systems: A Framework for Understanding the Neurocognitive Foundations of Intuition, in: D. D. Schmorrow, C. M. Fidopiastis (eds.), Foundations of Augmented Cognition, AC 2013, Lecture Notes in Computer Science 8027 (2013) 474-483. Doi: 10.1007/978-3-642-39454-6_51.
[69] B. Juba, Implicit Learning of Common Sense for Reasoning, in: Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, 2013, pp. 939–946.
[70] V. Propp, Morphology of the Folktale, University of Texas Press, Austin, 1968.
[71] A. Ciarlini, M.A. Casanova, A.L. Furtado, P. Veloso, Modeling interactive storytelling genres as application domains, J. Intell. Inf. Syst. 35 (3) (2010) 347–381, https://doi.org/10.1007/s10844-009-0108-5.
[72] K. Eshghi, Abductive Programming with Event Calculus, in: Proceedings of the Fifth International Conference on Logic Programming, 1988, pp. 562-579.
[73] M. Shanahan, An Abductive Event Calculus Planner, J. Log. Program. 44 (1–3) (2000) 207–240, https://doi.org/10.1016/S0743-1066(99)00077-1.
[74] G. Paul, AI Approaches to Abduction, in: D. M. Gabbay.
[75] Abductive Reasoning and Learning. Handbook of Defeasible Reasoning and Uncertainty Management Systems 4 (2000) 35–98, https://doi.org/10.1007/978-94-017-1733-5_2.
[76] A. E. M. Ciarlini, C. T. Pozzer, A. L. Furtado, B. Feijo, 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, 2005, pp. 133–140. Doi: 10.1145/1178477.1178495.
[77] E.S. de Lima, B. Feijó, S.D.J. Barbosa, A.L. Furtado, A.E.M. Ciarlini, C.T. Pozzer, Draw your own story: paper and pencil interactive storytelling, Entertain. Comput. 5 (1) (2014) 33–41, https://doi.org/10.1016/j.entcom.2013.06.004.
[78] E.S. de Lima, B. Feijó, A.L. Furtado, Video-based interactive storytelling using real-time video compositing techniques, Multimed. Tools Appl. 77 (2) (2018) 2333–2357, https://doi.org/10.1007/s11042-017-4423-5.
[79] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, 2017, pp. 6000-6010.
[80] G. Sartori, G. Orrù, Language Models and Psychological Sciences, Front. Psychol. 14 (2023) 1279317, https://doi.org/10.3389/fpsyg.2023.1279317.
[81] M. Twain, A Connecticut Yankee in King Arthur's Court, SeaWolf, 2018.
[82] A. Blair-Early, M. Zender, User interface design principles for interaction design, Des. Issues 24 (3) (2008) 85–107, https://doi.org/10.1162/desi.2008.24.3.85.
[83] A.E. Ali, K.P. Venkatraj, S. Morosoli, L. Naudts, N. Helberger, P. Cesar, Transparent AI Disclosure Obligations: Who, What, When, Where, Why, How, in: Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems (2024) 1–11. Doi: 10.1145/3613905.3650750.
[84] S. Chesterman, Good models borrow, great models steal: intellectual property rights and generative AI, Policy Soc. (2024) puae006, https://doi.org/10.1093/polsoc/puae006.
[85] J. Smits, T. Borghuis, Generative AI and Intellectual Property Rights, in: B. Custers, E. Fosch-Villaronga (Eds), Law and Artificial Intelligence, Information Technology and Law Series 35 (2022) 323-344. Doi: 10.1007/978-94-6265-523-2_17.
[86] J. A. Fisher, Centering the Human: Digital Humanism and the Practice of Using Generative AI in the Authoring of Interactive Digital Narratives, in: L. Holloway-Attaway, J. T. Murra (Eds), Interactive Storytelling, ICIDS 2023, Lecture Notes in Computer Science 14383 (2023) 73-88. Doi: 10.1007/978-3-031-47655-6_5.
[87] N. Caporusso, Generative artificial intelligence and the emergence of creative displacement anxiety, Res. Directs Psychol. Behavior 3 (1) (2023), https://doi.org/10.53520/rdpb2023.10795.
[88] A. Zarifhonarvar, Economics of ChatGPT: a labor market view on the occupational impact of artificial intelligence, J. Electronic Business Digital Econ. (2023), https://doi.org/10.1108/JEBDE-10-2023-0021.
[89] Y. Cao, R.W. Proctor, Y. Ding, V.G. Duffy, Y. Zhang, X. Zhang, Influences of color salience and location of website links on user performance and affective experience with a mobile web directory, Int. J. Human-Computer Interact. 37 (6) (2021) 547–559, https://doi.org/10.1080/10447318.2020.1838188.
[90] B. Laugwitz, M. Schrepp, T. Held, Construction and evaluation of a user experience questionnaire, in: A. Holzinger (Ed.), HCI and Usability for Education and Work, USAB 2008, Lecture Notes in Computer Science 5298 (2008) pp. 63-76. Doi: 10.1007/978-3-540-89350-9_6.

[91] J.M. Baron, S. Bluck, That Was a Good Story! Preliminary Construction of the Perceived Story Quality Index, Discourse Processes 48 (2) (2011), 93-118. Doi: 10.1080/01638531003702109.
[92] OpenAI et al., GPT-4 Technical Report, in: arXiv:2303.08774 [cs.CL], 2023. Doi: 10.48550/arXiv.2303.08774.
[93] P.P. Ray, ChatGPT: A comprehensive review on background, applications, key challenges, bias, ethics, limitations and future scope, Internet of Things Cyber-Physical Syst. 3 (2023) 121–154, https://doi.org/10.1016/j.iotcps.2023.04.003.
[94] OpenAI Documentation, Production Best Practices. https://platform.openai.com/docs/guides/production-best-practices (accessed 12 May 2024).
[95] M. Xu et al., A Survey of Resource-efficient LLM and Multimodal Foundation Models, in: arXiv:2401.08092 [cs.LG], 2024. Doi: 10.48550/arXiv.2401.08092.
[96] N. Haefner, V. Parida, O. Gassmann, J. Wincent, Implementing and scaling artificial intelligence: A review, framework, and research agenda, Technol. Forecast. Soc. Chang. 197 (2023) 122878, https://doi.org/10.1016/j.techfore.2023.122878.
[97] E.S. de Lima, M.M.E. Neggers, A.L. Furtado, Multigenre AI-powered story composition, in: arXiv:2405.06685 [cs.CL], 2024. Doi: 10.48550/arXiv.2405.06685.