

Striving towards automated writing – Views on authorship in story generation research

Tuuli Hongisto <tuuli_dot_hongisto_at_helsinki_dot_fi>, University of Helsinki  <https://orcid.org/0000-0003-0228-7255>

Abstract

This article examines how authorship is approached in story generation research publications. Story generation research forms its own, distinctive context in which computer-generated literary texts are being produced. Through text analysis and comparisons to other forms of computer-generated literature, the article examines what kind of rhetoric is used when the authorship of computer-generated texts is described in this context, and how the roles of the programmers and the program are characterised. The findings suggest that the approach to authorship in story generation research is mainly technical, referring strictly to the production process of the text, and leaving out the meaning of authorship as responsibility and accountability of the work as an aesthetic whole. This technical view affects how human-computer relations are discussed in the research, as “human-generated” and “computer-generated” texts are contrasted with each other. Furthermore, this dichotomy of the human and the machine affects how the produced stories are evaluated.

Computer-generated literature has risen prominently into public discussion as large language models (LLMs) have become available for mainstream use. These models are only the most recent development in a long tradition of developing text-generating programs for digital computers that dates to the 1950s and 1960s [Ryan 2017] [Rettberg 2018]. Part of this tradition is story generation research, a sub-field of computer science where programs that write creative texts are developed. In this article, I examine how the authorship of computer-generated texts is approached in story generation research.

Story generation research deals with core issues of literature in many ways. As researchers Nick Montfort and Rafael Pérez y Pérez argue, the developers of story generation programs can both discuss and operationalise key concepts of literary theory, such as plot and character [Montfort and Pérez y Pérez 2023, 99]. Correspondingly, text-generation programs are considered to be a part of the wide spectrum of electronic literature: works that explore the capabilities of computers and networks, as defined by the Electronic Literature Organization (ELO).^[1] Many story-generating algorithms from past decades have been examined in literary research (see e.g. [Ryan 1991] [Wardrip-Fruin 2009] [Rettberg 2018]). At its best, the intercommunication between the disciplines of computer science and literary research can both aid the development of story generation programs and enrich the understanding of electronic literature. However, the potential of this dialogue has not yet been fully achieved. Commenting on recent research, Chhun et al. point out that humanities research could have much more to contribute to the development of story-generating algorithms and the study of computational creativity [Chhun et al. 2022, 5082]. I aim to contribute to this exchange of ideas by examining newer programs in story generation that have thus far received little attention in literary studies. I focus on the concept of authorship, which is both thoroughly theorised in literary studies and central to the objectives of story generation research.

Actual author and electronic literature

In literary research, a division is often made between the “actual author,” who is the intellectual creator of the work, and

1

2

3

the “implied author,” which is an image of the author that the reader can infer from the text (see e.g. [Booth 1983] [Chatman 1994] [Schmidt 2010]). In this article, my focus is on the question of how authorship is viewed in story generation research, a question that is primarily related to the actual author.

This aspect of authorship is generally associated with the assumption of an individual who can profit from their literary work and is responsible for and deserving of credit for it. This perception of authorship has evolved alongside print technology and the publishing industry [Woodmansee 1994, 35–36] [Finkelstein and McCleery 2005, 69–71] [Febvre and Martin 1984, 165–166]. Before the invention of printing in Europe, authorship was not defined in the same way as it is today, and the authors of texts were not considered their owners [Finkelstein and McCleery 2005, 69]. Writing was seen more as a form of craftsmanship or a channelling of inspiration from a higher power [Woodmansee 1994, 35]. The perception of authorship began to shift as distribution of hundreds of books became possible with print technology, and as copyright legislation began to be drafted in the 18th century [Woodmansee 1994, 36–37]. In addition, inspiration and genius began to be used to characterise authors around the same time [Jefferson 2014, 4, 61, 67–71]. These developments influenced the formation of the current understanding of authorship, as they led to the forming of the profession of writers who could support themselves with the proceeds of their writings [Woodmansee 1994, 36–37]. The author’s role as the party that is responsible for the aesthetic whole of the work is also often associated with the term [Lamarque 2009, 109–111].

In many respects, electronic literature often goes against the grain of this perception of authorship. For instance, in the 1990s, many researchers viewed the changes brought about by computers to writing and reading as a counterforce to the old concept of authorship that emphasised the individual author (e.g. [Gaggi 1997, 106–107] [Landow 1992, 70–77] [Bolter 1991, 4]). Many phenomena that challenge the established perception of authorship, such as collective writing, adaptations, plagiarism, and crowdsourcing ideas, are also made considerably easier with digital technologies [Goldsmith 2011, 2–3]. With computer-generated texts, ascribing authorship is complicated even further, as many factors, such as the design of the program and the data that the program has been given to work on, shape the produced text [Henrickson 2021]. Furthermore, with computer-generated texts, the program’s role as a co-author also comes into consideration [Lebrun 2017], and it has been even argued that computer-generated works can undermine the image of the author as a unique genius [Bök 2002, 10].

These differing perspectives — the common assumption of the author as an individual who is the owner of and responsible for their work, and the varying approaches in electronic literature that challenge this expectation — provide a background for analysing the research material as well. I compare how authorship is framed in story generation research in relation to these points of view.

Story generation programs as research material

The following research material is based on a number of story generation studies published over the last decade. In selecting the publications to analyse, I rely on reviews that have tracked the state-of-the-art of story generation programs [van der Lee et al. 2019] [Herrera-González, Gelbukh, and Calvo 2020] [Alhussain and Azmi 2021] [Guan et al. 2021] [Chhun et al. 2022] [Ansag and Gonzalez 2023]. The other publications I examine [Ammanabrolu et al. 2021] [Bottoni et al. 2020] [Cheong and Young 2015] [Daza, Calvo, and Figueroa-Nazuno 2016] [Fan, Lewis, and Dauphin 2018] [Guan et al. 2020] [Ippolito et al. 2020] [Jain et al. 2017] [Li, Ding, and Liu 2018] [Liu et al. 2020] [O’Neill and Riedl 2014] [Peng et al. 2018] [Peng et al. 2022] [Porteus and Lindsay 2019] [Tambwekar et al. 2019] [Xu et al. 2018] [Yao et al. 2019] were chosen by cross-referencing these reviews and selecting the articles that are discussed in more than one of them. This is to ensure that although I cannot comprehensively cover all of the recent work on story generation in this article, my research material aligns with what multiple other researchers have found noteworthy. The reviews themselves are also a subject of analysis since they discuss the authorship of computer-generated texts as well. I examine how the term “authorship” is used across these various publications in their entirety, including the descriptions of the programs developed and the evaluations of the stories produced.

Story generation programs are part of a larger phenomenon of text generation. Text generation with digital computers has been a method of producing literary works since the mid-20th century [Rettberg 2018, 31–33]. It is a versatile

practice where multiple sources of influence are at play. One such influence is dadaist poetry, and methods similar to the cut-up technique are used in computer-generated works [Funkhouser 2007, 32–34]. Furthermore, the work of Oulipo (*Ouvroir de littérature potentielle*) has also been influential in text-generation, as they explored the intersections of literature and mathematics from the 1960s on and are known for experimenting with combinatorial procedures and methods of constrained writing, that is, creating rules and restrictions within which different works could be created [Schäfer 2007, 121] [Funkhouser 2007, 33–34]. Later, text generation has been used in a similar way to produce works based on the rules created by the author [Funkhouser 2007, 34].

One key motivation for text-generation, especially in computer science, has been the imitation of human creativity. Story generating programs that have been developed in the field of computer science have been closely tied to research in artificial intelligence [Gervás 2012]. As Sharples & Pérez y Pérez put it, developing programs that can perform various creative tasks is seen as a way of modelling the human mind, and this work aims to answer questions such as “What is intelligence?” or “What makes humans creative?” [Sharples and Pérez y Pérez 2022, xii]. Here, storytelling is viewed as a task that is particularly difficult for computers, and is thus often discussed as one of the grand challenges of artificial intelligence [Fan, Lewis, and Dauphin 2018, 889] [Xu et al. 2018, 4306] [Guan et al. 2020, 93] [Ammanabrolu et al. 2021, 5859].

The goals and practices of story generation as a field of research also affects how the text-generation process and authorship are framed. As my analysis will show, text generation of literary texts in the context of computer science research is often viewed more in terms of problem solving than in terms of creating art. This marks a difference to many computer-generated works that are created outside of research-oriented environments. Throughout the article, I contrast story generation with text generation outside of computer science research, both to explore story generation as a subtype of electronic literature and to offer an external perspective on the conceptions of authorship that might be useful in the field.

Story generation researchers have historically utilised literary theory in the development of the programs, and some (e.g. [Gervás 2012] [Montfort and Pérez y Pérez 2023, 99]) have even suggested that story generation could work as a means of “testing” different concepts and theories of literature. Story generation researchers have been especially interested in formulations of plot in narrative theory, most notably Vladimir Propp’s famous study *Morphology of a Folktale* (1968), which has been used either as a source of inspiration or even as a basis for story generation programs in many instances (see e.g. [Turner 1994] [Peinado and Gervás 2006] [Gervás 2013] [Veale 2014] [Ryan 2017]).

However, based on my review, it seems that dialogue between the fields is not necessarily a part of all story generation research. The objective of contributing to or utilising literary theory is not generally present in the articles surveyed, and literary theory concerning concepts such as story and plot is often not referenced at all or only in passing (see e.g. [Li, Ding, and Liu 2018] [Xu et al. 2018] [Fan, Lewis, and Dauphin 2018] [Peng et al. 2018] [Tambwekar et al. 2019] [Ippolito et al. 2020]). Moreover, when literary theory is referenced, it is often in the context of techniques for plot modelling (e.g. [O’Neill and Riedl 2014] [Ansag and Gonzalez 2023]). I wish to suggest that engaging with other areas of literary theory, such as stylistics and genre studies, could also be beneficial for story generation research.

The textual analysis of the research material was conducted by first going systematically through the articles and then identifying recurring themes related to the research question: What or whom do the researchers view as the stories’ author and what does that imply on how authorship is viewed in the research? I observed that discussion related to authorship took place especially in connection with motivating the research, evaluating the outputs, and refining the style of the outputs. Themes that recurred throughout much of this material were reducing the impact of the human author on the outputs, juxtaposing human and computer-generated texts, and avoiding “machine-like” features in outputs. In the sections that follow, I examine how these topics were addressed in the articles and what kind of overall picture of the approach to authorship can be drawn from them.

Authorial freedom or authorial burden – the role of the human author in story generation

Authorship is arguably one of the most important topics in story generation research, especially when research motivations are examined. Developing a “computer author” that mimics human creativity is one of main objectives expressed in these articles. This goal is expressed, for example, by linking the research to the development of computational creativity, which is referred to simply as an imitation of human creativity (see e.g. [Jain et al. 2017, 2] [Alhussain and Azmi 2021, 1] [Ansag and Gonzalez 2023, 878]). Some researchers define the term “creative” further, asserting that a program is creative if the output it produces is different enough from the parameters and other information given to it [Alhussain and Azmi 2021, 25] or if the output is unpredictable to its programmers [Daza, Calvo, and Figueroa-Nazuno 2016, 14]. As [Yao et al. 2019, 7381] put it, the aim is often to make a program that passes the Turing test.^[2]

Over the time period during which the articles examined were published, the methods of story generation have altered greatly. At the beginning of the 1960s, in the early days of story generation research, the programmer played an integral part in shaping the information required for storytelling into a computer-readable form [Ryan 2017]. Methods such as simulation, planning, and story grammars have been and continue to be utilised in story generation.^[3] In recent decades, however, neural network methods, and even more recently LLMs such as generative pretrained transformers, have surged in popularity and are regularly used alongside other methods. With the use of neural networks, in which LLMs are also based on, it is possible to train the program to imitate a set of training data it has been given, for example by recognising which words typically appear together [Henrickson 2018a, 9–10] [Paaß and Giesselbach 2023, vii]. This has enabled greater levels of automation in the story generation process.

This shift in methodology has also affected how the authorship of the produced stories is discussed. With advances in natural language generation, such as the increased use of neural networks, previously dominant methods such as planning and story grammars are often criticised. [Yao et al. 2019, 7378] and [Tambwekar et al. 2019, 5982] criticise older approaches for relying heavily on extensive domain knowledge and annotation crafted by the programmers. “Hand-crafting” the information that is given to the program requires a lot of effort from the developers and significantly impacts the construction of the text, which is viewed as a shortcoming. In contrast, the use of neural models means there is less need for “hand-crafted” knowledge, an outcome that is regarded as desirable [Liu et al. 2020, 1725].

The researchers’ contributions are often contrasted with previous research by showcasing how much they have achieved “authorial reduction” (e.g. [Porteus and Lindsay 2019, 1074]) or lessened the “authorial burden” that goes into making a story generation program [Bottoni et al. 2020, 21] [Ansag and Gonzalez 2023, 878]. Correspondingly, previous work can be praised if generating the story requires “almost no human intervention” [Daza, Calvo, and Figueroa-Nazuno 2016, 10]. In other words, the aim is most often to minimise the role of the human(s) — especially the people who develop the program — in the creation of the text.

The goal of minimising the developer’s contribution is reflected in how the authorship of the stories is characterised. Notably, the researchers do not generally refer to themselves as authors or writers of the produced stories.^[4] Instead, the researchers typically name the program they have developed and refer to that program as the writer of the stories. Alternatively, the researchers often refer to the writer as a “model” or “system” and in some cases even as an “agent,” which implies some degree of authority over the output [Tambwekar et al. 2019] [Ammanabrolu et al. 2021].

As these findings show, authorship — or more narrowly, defining the extent to which a story is computer-generated — is closely tied to the goals of story generation research. However, the “authorship” term is not explicitly defined in any of the publications, as would be the case with literary research that focuses on authorship. Despite this, it is apparent the concept “authorship” is ascribed a meaning that is more specific than the everyday use of the word. Based on the research material, the developers of story generating algorithms seem to discuss authorship almost entirely by assessing how minimal their own role is in the text’s construction. This approach might partly stem from the generally desired features in software design: The less hand-crafted material there is in a program, the more flexible it becomes [Xu et al. 2018, 4307–4308] [Liu et al. 2020, 1725]. In such an approach, by focusing on the “authorial reductions,” the concept of authorship is reduced to a technical question: who or what has written each part in a story.

The effort to diminish the role of the human contributors as authors differs from many other works of electronic literature

that are produced using text-generation programs. Take, for example, Philippe Bootz's *La série des U*. As Bootz who is a renowned poet in the field of electronic literature, describes it, "The work generates a different text-that-is-seen [...] each time it is played through subtle variations in the timing at which the textual elements appear and the relation between the verbal text and the sonic component [...]" [Bootz 2005]. Like story generators, Bootz's work is not a single, stable text but a text generator that produces different iterations according to principles that he designed. The process of creating the program that produces these texts is similar to story generation. Yet, in contrast to the perceived computer authorship in story generation, Bootz considers himself the (co)author of *La série des U*. This is evident from how he refers to the work as "a typical example of adaptive generation that illustrates *my* present approach of e-poetry" [Bootz 2005, emphasis added]. Research on this work also refers to it as being authored by Bootz. For instance, Alexandra Saemmer considers *La série des U* as an illustrative case of "authors who [...] purposefully reveal details of their program" [Saemmer 2009, 485]. The comparison of story generation programs to similar works outside of the field highlights how differently the roles of the program and its developer can be viewed in terms of authorship.

Some writers who use computer-programs to generate literary texts do not view the computer program(s) as author(s) in any sense. A computer-poet Allison Parrish, for example, approaches making poetry with text-generation as putting together a collage. For instance, her work *Two of Pentacles* is based on seventy-eight different snippets written by her and rearranged using different techniques. These include assigning each snippet a vector of words and letter sequences that represent it and a program that gives measurements on the snippets' similarities based on the vectors [Parrish 2025]. She views the different programs that she uses in the process as a medium of the assemblage, not as authors or collaborators. As she writes, "I never [...] attribute authorship to a program or a model. If I publish the results of a text generator, I am responsible for its content" [Parrish 2021]. It is apparent that in contrast to the way story generation researchers think about authorship, Parrish views authorship as a broader concept than merely estimating the ratio of human and computer effort.

21

The authorship of story generation programs could also be approached from the perspective of remixing and reusing material. Digital media theorist Leah Henrickson has proposed this kind of approach for another kind of work of algorithmic authorship: Twitter bot Pentametrone made by Ranjit Bhatnagar, a bot that gathered tweets in iambic pentameter in its feed and also arranged them into longer poems [Henrickson 2018b, 12–13]. Pentametrone inspired people to use the material in creative ways, and [Henrickson 2018b, 13] argued that instead of examining Pentametrone's retweets and poems as final literary texts, the chains of appropriation of texts are the interesting subject of analysis. Even though story generation programs are quite different from social media bots like Pentametrone in terms of their purpose and audience, they share a similar function of providing material for readers that can be used in subsequent texts. Story generation programs are mainly examined by other researchers that can utilise them in creating their own programs. Approaching story generation programs from this perspective could also bring a new element in discussion of their authorship. It highlights the collaboration of not only between the program and its developer, but between researchers.

22

One reason for the narrower approach to authorship in the research material is that in story generation producing the stories is only of an instrumental value, as the purpose of the research is technological advancement. While story generators are developed for purposes such as storytelling [Peng et al. 2018] [Porteus and Lindsay 2019] [Ansag and Gonzalez 2023], creating works of fiction [Daza, Calvo, and Figueroa-Nazuno 2016], interactive entertainment [Alhussain and Azmi 2021], and games [Cheong and Young 2015], these aims are described as the ultimate goals of story generation, not as direct products of the story generators in their current state. This is perhaps the most crucial difference between story generation and works such as *La série des U* or *Two of Pentacles*: The process and products of story generation are not viewed as literary works but rather as steps towards the goal of computational creativity. The developers seem to view their work in reducing the "authorial burden" as a technical task, not as an artistic endeavour. Thus, their view of authorship is also technical, and other dimensions of the term, such as the social role and responsibilities of the author, are not discussed.

23

This technical definition of authorship that leaves out so many aspects typically associated with literary authorship also leads one to ask whether text generators and the texts they produce should be approached as literature. This does not seem to be the framing in the surveyed articles, as the programs and the texts they produce are never referred to as

24

works of art or literature. However, this does not necessarily mean that the programs and their outputs should not be regarded as literature. In fact, this has been typical in older story generation programs that are regularly mentioned as part of electronic literature as well. Developers of older story generators such as Tale Spin or Mexica originally framed their work primarily as research of computational creativity and imitation of human cognition [Meehan 1976, 212–226] [Sharples and Pérez y Pérez 2022, 119–120]. Yet, examining these programs as literature has produced valuable discussion, for example on whether works of electronic literature should strive for simulating human writing [Aarseth 1997, 131–141] [Wardrip-Fruin 2009, 155].

Whether the framework of “literature” is useful in examining story generation programs depends also on the purpose and content of the produced texts. Older story generation programs tended to produce outputs in an identifiable literary genre, such as fables or fairy tales [Meehan 1976] [Riedl and Young 2010], detective stories [Klein et al. 1973], or short stories [Bringsjord and Ferrucci 2000]. In comparison, most stories presented in the research material are much shorter and not easily recognisable as belonging to a particular genre. Still, the researchers generally do have aesthetic objectives for the produced texts. These aims include creating suspense [O'Neill and Riedl 2014] [Cheong and Young 2015], literary language [Daza, Calvo, and Figueroa-Nazuno 2016], deepening the characters [Bottoni et al. 2020], following a narrative structure [Porteus and Lindsay 2019] [Liu et al. 2020], or arousing interest in the reader [Yao et al. 2019] [Peng et al. 2018] [Jain et al. 2017]. In this sense, it is worthwhile to analyse their outputs in literary terms.^[5]

25

The human-machine dichotomy

Developers of story generation programs not viewing themselves as authors or co-authors of the stories is reflected in the juxtapositions of “the human” and “the machine.” This is most evident in the evaluation processes of the programs. In studies that conduct a so-called “human evaluation” of the texts, it is common to include “human-written” texts alongside computer-generated ones to be read by the participants of the study (see e.g. [Cheong and Young 2015] [Daza, Calvo, and Figueroa-Nazuno 2016] [Guan et al. 2020] [Guan et al. 2021, 6396] [Chhun et al. 2022, 5797]).^[6] Since the goal of story generation is to mimic human intelligence and creativity, “human-written” stories are used as a yardstick against which the success of the program is measured. For example, the evaluation scores of the stories can be used as an indicator of the program’s effectiveness [Liu et al. 2020, 1731], “narrative intelligence” [Chhun et al. 2022, 5799], or the stories’ suspensefulness [Cheong and Young 2015]. By contrasting the “human-generated” and computer-generated texts in this way, the human and the machine are positioned as separate entities.

26

A study conducted by [Daza, Calvo, and Figueroa-Nazuno 2016] is an especially interesting example of this phenomenon in terms of authorship, because it juxtaposes experimental works by authors and texts that are generated with the program developed by the researchers. Image 1 shows a list of texts that participants were asked to evaluate without knowledge of their authorship.

27

Title	Author
Naked Lunch (fragment)	William Burroughs
A Big Man Existed	Machine
The Left Gets Threatened	Rachel Summer
The White Sky Met	Machine
A Beautiful Story	Machine
Finnegans Wake (fragment)	James Joyce
The First Man	Machine

Figure 1. List of texts evaluated by study participants in [Daza, Calvo, and Figueroa-Nazuno 2016, 17].

The picture depicts a dichotomy, where authors such as James Joyce are named but the texts generated by the program are put under the generic label of “machine.” This separation of the two is further repeated in the results, where the performance of “artificial texts” and “human texts” in the evaluation is presented. This division reflects the perception of authorship as a technical either-or question: A text is either is or is not a result of a text-generation process.

The picture also shows that multiple experimental works are included in the comparison. The authors of the article state that this is because they want to stress that the texts they produced are “at least as experimental” as the included human written works [Daza, Calvo, and Figueroa-Nazuno 2016, 18]. However, the article does not discuss in more detail how the experimental texts in question relate to authorship. William Burroughs, for instance, famously aimed to obfuscate authorial authority in his works: He has described writing *Naked Lunch* without conscious control [Heal 2016], and suggested that he used the cut-up method in writing it [Gontarski 2015, 173]. In this way, *Naked Lunch* and Burrough's attempts to minimise his own conscious role in writing it provides a kind of precursor for similar pursuits in story generation. This nuance is lost in the human-machine dichotomy.

Furthermore, the line between “the human” and “the machine” is often blurrier than the dichotomy suggests. For one, the “computer-generated” works are often largely shaped by the developers of the program. For example, in the case of *Suspenser*, the developers have formulated a list of fifty-seven events that the program can pick from to include in a storyline, which has an effect on the produced stories. To evaluate the results, Cheong & Young have also asked a “human author” to construct stories based on the same pool of events, and these stories are then read and evaluated by study participants [Cheong and Young 2015, 46]. The categories “human-generated” and “machine-generated” do not accurately describe the process of creating these stories. Both types of stories are shaped by the programmers, who have chosen what events are possible in a story, regardless of whether the storyline itself is constructed by the human author or the *Suspenser*.

Even with programs where the role of the developers is less central than with *Suspenser*, they shape the produced stories significantly. The programs need some sort of rules to function, for example concerning the length and genre of the stories, the kinds of texts that the program should imitate, and whether the program should focus mainly on coherence, suspense, or literariness. These rules are designed and put in place by the developers of the programs. As an example, [Daza, Calvo, and Figueroa-Nazuno 2016, 19] state that their program produces output based on a single word that the user of the program gives to it as a prompt, and that they do not have any control over the substance of the stories produced. In a limited technical sense — who or what has constructed the text at the sentence level — this

28

29

30

might be the case. However, if authorship is understood more broadly and includes coming up with the concept of the work, the developers can be regarded as co-authors of the texts. The researchers undeniably shape the production of the stories by setting the program goals such as coherence and literary style, and by choosing the training data.

Furthermore, “purely” human-written texts are becoming an increasingly difficult category to identify, given the recently risen popularity of LLM based programs. As LLMs have become more widely known and available, creating partly computer-generated literary works has been possible for a larger number of people. This is evident, for example, on book platforms such as Amazon that have experienced acceleration of publishing rates [Cabezas-Clavijo et al. 2024].^[7] The fact that human and machine writing are increasingly merging outside of the research context is making it increasingly difficult to distinguish between “human-written” and “computer-generated” texts. Especially when story generation programs are examined as a form of literature, a more accurate approach could be to acknowledge that the outputs are not void of human influence, and that whether they are written by a human or a machine is a sliding scale.

31

Writing with a machine — Approaches to “mechanical” features in texts

As demonstrated, the main goal of story generation programs is to imitate human intelligence, and “human-written” stories are used as a measure of the success of computer-generated stories. This aim is reflected in attempts to reduce the “mechanical” feel of the generated texts. These attempts are most clearly seen in evaluations of the stories, where the desirable and undesirable features of the produced stories are defined.

32

Story features that are deemed “undesirable” are discussed especially in the evaluations of the quality of the stories. These features include, most notably, repetition and incoherence in stories (see e.g. [Yao et al. 2019, 7383–7384] [Tambwekar et al. 2019, 5986] [Guan et al. 2020, 93] [Herrera-González, Gelbukh, and Calvo 2020, 82] [Guan et al. 2021, 6396]). These qualities are often discussed in contrast to the stories written by people and thus are deemed too machine-like. Bottoni et al. describe how stylistic features such as repetition can make the text sound “mechanical”:

33

Sentences in stories are affected by the sentences around them, and most readers can quickly notice if they don’t flow together “correctly”: e.g. “Noah went to the store. He went to the store to buy eggs.” or “Sara saw a spaceship. She boarded the spaceship.” There are quotation marks around “correctly” because prescriptively these sentences have correct grammar. However, they sound robotic and inhuman — “incorrect” in a rhetorical sense. [Bottoni et al. 2020, 24]

The style of the text — in this case, the monotone syntax of the sentences — is deemed incorrect because of its “mechanicalness.” Repetition in story generation can mean repeating the same syntactic structure, as in the case of the program by [Bottoni et al. 2020], but also what developers characterise as repetitiveness of plot, where the same event is depicted repeatedly.

Incoherence in stories is another feature that is often considered undesirable. What researchers regard as incoherence can manifest in the stories in various ways. For example, [Xu et al. 2018, 4313] note that some of the stories produced have a “chaotic timeline,” meaning that the stories go backwards in time. Conflicts of logic or impossible events can also result in incoherence. Guan et al. give the following extract of a story with conflicting logic:

34

He noticed a car in the road. He decided to stop. He got out of his car. He drove for half an hour.
[Guan et al. 2020, 103]

In this example, beginning to drive the car after exiting it is what creates incoherence. Researchers can also define incoherence as general “chaoticness” in stories. For example, Liu et al. characterise one of the stories generated by their program as “inferior” and “chaotic,” because it lacks a clear main character [Liu et al. 2020, 1731]. As in the case of repetition, the researchers often attribute the incoherence to the computational origins of the text. For example, when describing types of incoherence in stories, Guan et al. point out that “unrelated events and conflicting logic, or globally chaotic scenes” are common in natural language generation models [Guan et al. 2021, 6396]. Because coherence is regarded by the researchers as an integral feature of successful stories, they try to avoid all instances of incoherence.

This approach, which emphasises the necessity of reducing features that supposedly give away the text's computational origins, is different from those of many other genres of electronic literature. Many writers specifically use computational methods because of the features that they can bring into the text. For example, random generation of text can be used to create surprising associations [Hayles 2008, 26]. In fact, "randomness," which can be computationally simulated in many ways, has been an important part of the formation of electronic literature and other forms of digital artistry [Montfort et al. 2013]. In these cases, the "chaoticness" is integral to the works, which is in some ways an opposite approach to story generation.

35

Furthermore, in computer-generated literature outside of story generation research, the fact that the text is a result of a joint effort of human and computer can be embraced, as the interplay between the two is often what makes the work interesting. The interaction between the programmer and the program can be intentionally made visible to the reader as an integral part of the work, either by revealing the code that produces the text [Marino 2020, 2015–220] [Hongisto 2023], or by letting the reader uncover the underlying principles of the program by observing multiple iterations of texts produced by the same program [Schoenbeck 2013]. The value that the human-computer interaction brings to the text can also be partly unintentional. Pressman et al. note that in digital culture, surprising results or glitches are often appreciated either as an aesthetic quality of the work or as an opportunity to catch a glimpse of the structure of the program [Pressman, Marino, and Douglass 2015, 41]. Whereas with these kinds of works the ambiguity of authorship is something to be examined and pondered by the reader, the goal of story generation is to eliminate such considerations. As with the Turing test, the aim can be that of deception [Natale 2021, 28] — to trick the evaluators into thinking a text is written by a human (see e.g. [Daza, Calvo, and Figueroa-Nazuno 2016]). Contrasted with the approaches that highlight the human-machine interaction, the aim to imitate human writing by avoiding "mechanicalness" in texts can be distinguished as a feature that characterises story generation as a subtype of electronic literature, at least in the form that the research material represents.

36

Considering the aim of mimicking human writing in story generation, avoiding features that might lead the reader to believe that a text is computer-generated is understandable. However, categorically avoiding features such as repetition and incoherence may not always serve this goal, as it does not take into account their value as stylistic devices. The value of the unexpectedness of computer-generated outputs can be observed through stories that have been labelled failures. The examples of "unsuccessful" stories are often intriguing, sometimes arguably more so than their "successful" counterparts. Take for example these two short stories generated by the model by [Guan et al. 2020] below:

37

[MALE] was driving around in the snow. Suddenly his car broke down on the side of the road.
[MALE] had to call a tow truck. The tow truck came and took [MALE] home. [MALE] was happy he was able to get home. [Guan et al. 2020, 103]

[MALE] was on thin ice with his job. He had a friend over to help him. [MALE] was able to hold his breath the entire time. he was so cold that he froze in his tracks. [MALE] finally felt good about himself. [Guan et al. 2020, 104]

Out of these examples, the former is the "successful" story and the latter the "unsuccessful" one, as Guan et al. aim to produce coherent stories and have labelled the latter story as "chaotic" [Guan et al. 2020, 103–104]. However, there is something enthralling in the story's use of metaphors that have to do with freezing and that merge with the protagonist feeling cold, and there is an element of surprise in the protagonist's contentment at the end of the story. It is not a coherent story by the standards of [Guan et al. 2020], but it arguably generates more interest and questions, as well as providing more room for interpretation than the "successful" story.

Another well-known case of similar "unsuccessful" stories turning out more amusing than the "successful" ones are Tale-spin's failed stories [Meehan 1976]. The so-called "mis-spun" fables generated by the program first raise genre expectations only to soon drop them, which creates a humorous effect. These "unsuccessful" stories have been circulated much more widely than the successful ones, partly because of their entertainment value [Bolter 1991, 180] [Aarseth 1997, 131] [Wardrip-Fruin 2009, 130]. The popularity of these stories is an example that there can be value in

38

the stories that the programs' creators have deemed unfit.

Similarly, avoiding all instances of repetition may not effectively serve the aim of producing texts that imitate human writing. Yao et al. give the following example of an unsuccessful story with repetition on the level of plot:

39

Anna was cutting her nails. She cut her finger and cut her finger. Then she cut her finger. It was bleeding! Anna had to bandage her finger. [Yao et al. 2019, 7380]

It is easy to understand why this story is considered a failure, since Yao et al.) aim to describe “a sensible sequence of events” [Yao et al. 2019, 7378]. However, as they also want to “mimic human practice in real world story writing” it is worth pointing out that by forbidding repetition, they dismiss its potential as a stylistic device [Yao et al. 2019, 7379]. In the above story, for example, the repeated finger cutting brings a chilling tone to the text, implying self-harm.

One could argue that pointing out the humour and tellability of these stories is beside the point of story generation. The difficulty of producing long, coherent texts without repetition using natural language processing models is well established [Fan, Lewis, and Dauphin 2018] [Ammanabrolu et al. 2021] [Peng et al. 2022]. From this perspective, it is understandable that the researchers want to develop these capabilities in story generation models. More generally, the critique of ignoring common literary devices in the development of story generation programs applies best to research where the intent is to produce a literary output. Still, avoiding incoherence and repetition entirely and dismissing them as “mechanical” features does not seem consistent with the goal of mimicking human writing. Although computers are prone to producing texts with such features, they are not inherently machine-like but common stylistic devices in literature. Even when randomness is not used as a foregrounded literary device, is some amount of randomness generally desirable even in computer-generated creative texts [Hua and Raley 2020]. If these features are left out of writing, part of the expressive power of fiction is lost.

40

Conclusion

The approach to authorship remains somewhat consistent throughout the material presented here. While the articles do not explicitly define the concept of authorship, a closer examination suggests that the term is used in a way that differs from its everyday use. The aim is to automatise the creation process of texts and to minimise the role of the human authors: a challenge that seems to be viewed as a technical one by the researchers. This view is reflected in the technical use of the term *authorship*, mostly referring to who (or what) has composed which parts of the text, while the aspect of who (or what) is responsible for the work is not discussed.

41

This narrow perspective on authorship affects how the human-machine relations in the creative process of the texts are discussed. The stories are most often referred to as computer-generated and contrasted with texts written by humans. This view differs from a large part of computer-generated literature, where authorship is often regarded as a broader concept, and the role of the designer of the program as an author is often empathised more.

42

This difference in approach to authorship is also reflected on the content of the produced texts. One consequence of the dichotomy of human and machine authorship is that the stories can often be evaluated based on the absence of features that are regarded as “mechanical,” mainly incoherence and repetition. Developing programs that are less prone to producing texts with such features undoubtedly improves the understandability and fluency of texts, but without them texts mimicking human writing would remain lacking in stylistic versatility.

43

Examining the approach to these “mechanical” features in texts reveals a larger issue within the evaluation of computer-generated texts. Many researchers stress the difficulty of reliably evaluating the quality of the stories (e.g. [van der Lee et al. 2019, 355] [Herrera-González, Gelbukh, and Calvo 2020, 86] [Chhun et al. 2022, 5794]). Indeed, the task of objectively and quantifiably evaluating stories is fundamentally challenging since it enters the territory of personal taste and aesthetic judgements — an area where absolute claims of quality are impossible. However, incorporating more contextual information in the process of both development and the evaluation of the program (Whom are the stories aimed at? What types of reading situations are they meant for?) could help in formulating the evaluation criteria. Different media, text types, genres, and reading contexts all affect which stylistic and narrative features are desirable,

44

and thorough knowledge of these topics could help in designing programs for different purposes. Developing evaluation criteria for the stories might be an area where story generation could benefit from an interdisciplinary approach.

The authorship of the produced texts can be discussed differently depending on whether story generation is placed in the realm of science or art. On one hand, in the sample of articles analysed in this article, the programs or produced texts are regarded not as literature but as technical improvements of natural language generation technology. This view explains the rather technical outlook on authorship. On the other hand, story generation can be approached as a genre of electronic literature. This approach evokes a more diverse image of authorship than can be inferred from story generation research: It refers not only to how the text technically came to be, but to who (or what) is accountable for it. Outlining the extent to which story generation aims at literary expression is an interesting question from the perspective of both computer science and electronic literature studies. In computer science, considering the question would mean defining the intended use of the programs more precisely, and if the aim is literary expression, taking into account the wide range of features of narrative and fiction. In the context of electronic literature research, on the other hand, story generation research must be viewed as a broad and varied field, some parts of which can be more justifiably included in the scope of electronic literature than others. Navigating the research output of story generation research and its relations to electronic literature — and literature more broadly — deserve close examination also in the future.

45

Funding Acknowledgement

This research was funded by Emil Aaltonen Foundation (Emil Aaltosen Säätiö).

46

Notes

[1] <https://eliterature.org/about/>

[2] The Turing test is a method frequently referenced in artificial intelligence research, which tests whether a program can exhibit behaviour that leads its users to believe they are interacting with a human [Natale 2021, 2].

[3] Simulation is a technique in which the program is given the characters and objects of the storyworld, and the story progresses with the interactions between these entities. In the planning method, goals are set for the characters and/or the story, and the story is constructed to fulfil these goals. Story grammars are overarching structures for the story events, which the program follows. For more extensive definitions of these methods, see [Kybartas and Bidarra 2017].

[4] The researchers not perceiving themselves as writers or authors of the works is not in itself unique to story generation. Scott Rettberg notes that authors of electronic literature often elude disciplinary categorisations such as poet, engineer, and designer, and instead these different roles can merge [Rettberg 2018, 35].

[5] There are, however, a few articles [Li, Ding, and Liu 2018] [Ippolito et al. 2020] in the research material where such aims are not mentioned. This means that the produced outputs are neither recognisable as literary texts nor have a context where they would be offered to readers as literature. In these cases, it would not seem fruitful to insist on interpreting and critiquing them as literature. The fact that there are these kinds of programs in the research material shows that story generation is a versatile field that cannot be regarded as a unified whole in relation to electronic literature.

[6] Most of the papers conducted both automatic evaluation and evaluation by human participants [O'Neill and Riedl 2014] [Fan, Lewis, and Dauphin 2018] [Li, Ding, and Liu 2018] [Peng et al. 2018] [Xu et al. 2018] [Porteus and Lindsay 2019] [Yao et al. 2019] [Guan et al. 2020] [Liu et al. 2020], while some relied only on automatic evaluation [Jain et al. 2017] [Ippolito et al. 2020], and some only on human evaluation [Cheong and Young 2015] [Daza, Calvo, and Figueroa-Nazuno 2016] [Tambwekar et al. 2019] [Ammanabrolu et al. 2021] [Bottoni et al. 2020] [Peng et al. 2022].

[7] This development has also brought about the question of authorship of not only the developers, but also the users of the program. However, there is little discussion on the role of the reader in relation to authorship in the research material. This is not only because the research material largely precedes the widespread use of LLMs, but also because a significant number of the story generation programs are not designed to assist in writing or to be used interactively like chatbots (at least not at the state that they are presented in the publications). Instead, the aim is to create stories whose content the readers cannot influence, based on parameters given at the beginning.

Works Cited

- Aarseth 1997** Aarseth, E.J. (1997) *Cybertext: Perspectives on ergodic literature*. Baltimore: Johns Hopkins University Press.
- Alhussain and Azmi 2021** Alhussain, A.I. and Azmi, A.M. (2021) "Automatic story generation: A survey of approaches", *ACM Computing Surveys*, 54(5), pp. 1–38. Available at: <https://doi.org/10.1145/3453156> (Accessed: April 2024).
- Ammanabrolu et al. 2021** Ammanabrolu, P. et al. (2021) "Automated storytelling via causal, commonsense plot ordering", in *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7), pp. 5859–67. Available at: <https://doi.org/10.1609/aaai.v35i7.16733> (Accessed: April 2024).

- Ansag and Gonzalez 2023** Ansag, R.A. and Gonzalez, A.J. (2023) "State-of-the-art in automated story generation systems research", *Journal of Experimental & Theoretical Artificial Intelligence*, 35(6), pp. 877–931. Available at: <https://doi.org/10.1080/0952813X.2021.1971777> (Accessed: April 2024).
- Bök 2002** Bök, C. (2002) "The piecemeal bard is deconstructed: Notes toward a potential robopoetics", *Object 10: Cyberpoetics*, pp. 10–18. Available at: <http://www.ubu.com/papers/object.html> (Accessed: April 2024).
- Bolter 1991** Bolter, J.D. (1991) *Writing space: The computer, hypertext, and the history of writing*. Hillsdale: Lawrence Erlbaum Associates.
- Booth 1983** Booth, W.C. (1983) *The rhetoric of fiction*, 2nd edn. Chicago: University of Chicago Press.
- Bootz 2005** Bootz, P. (2004) *The set of U / La série des U*. Available at: https://collection.eliterature.org/1/works/bootz_fremiot_the_set_of_u.html (Accessed: April 2024).
- Bottoni et al. 2020** Bottoni, B. et al. (2020) "Character depth and sentence diversification in automated narrative generation", in *Proceedings of the 33rd Annual Florida Artificial Intelligence Research Society Conference (FLAIRS-2020)*. AAAI Press, pp. 21–26. Available at: <https://cdn.aaai.org/ocs/18401/18401-79301-1-PB.pdf> (Accessed: April 2024).
- Bringsjord and Ferrucci 2000** Bringsjord, S. and Ferrucci, D.A. (2000) *Artificial intelligence and literary creativity: Inside the mind of BRUTUS, a storytelling machine*. Mahwah: L. Erlbaum Associates.
- Cabezas-Clavijo et al. 2024** Cabezas-Clavijo, Á. et al. (2024) "This book is written by ChatGPT: A quantitative analysis of ChatGPT authorships through Amazon.com", *Publishing Research Quarterly*, 40(2), pp. 147–63. Available at: <https://doi.org/10.1007/s12109-024-09998-w> (Accessed: April 2024).
- Chatman 1994** Chatman, S. (1994) *Story and discourse: Narrative structure in fiction and film*, new edn, 6th pr. Ithaca: Cornell University Press.
- Cheong and Young 2015** Cheong, Y. and Young, R.M. (2015) "Suspenser: A story generation system for suspense", *IEEE Transactions on Computational Intelligence and AI in Games*, 7(1), pp. 39–52. Available at: <https://doi.org/10.1109/TCIAIG.2014.2323894> (Accessed: April 2024).
- Chhun et al. 2022** Chhun, C. et al. (2022) "Of human criteria and automatic metrics: A benchmark of the evaluation of story generation", in *Proceedings of the 29th International Conference on Computational Linguistics*, pp. 5794–836. Available at: <https://aclanthology.org/2022.coling-1.509> (Accessed: April 2024).
- Daza, Calvo, and Figueroa-Nazuno 2016** Daza, A., Calvo, H. and Figueroa-Nazuno, J. (2016) "Automatic text generation by learning from literary structures", in *Proceedings of the Fifth Workshop on Computational Linguistics for Literature*, pp. 9–19. Available at: <https://doi.org/10.18653/v1/W16-0202> (Accessed: April 2024).
- Fan, Lewis, and Dauphin 2018** Fan, A., Lewis, M. and Dauphin, Y. (2018) "Hierarchical neural story generation", in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 889–98. Available at: <https://doi.org/10.18653/v1/P18-1082> (Accessed: April 2024).
- Febvre and Martin 1984** Febvre, L. and Martin, H. (1984) *The coming of the book: The impact of printing 1450–1800*, trans. by D. Gerard. London: Verso.
- Finkelstein and McCleery 2005** Finkelstein, D. and McCleery, A. (2005) *An introduction to book history*, 2nd edn. London: Routledge.
- Funkhouser 2007** Funkhouser, C. (2007) *Prehistoric digital poetry: An archaeology of forms, 1959–1995*. Tuscaloosa: University of Alabama Press.
- Gaggi 1997** Gaggi, S. (1997) *From text to hypertext: Decentering the subject in fiction, film, the visual arts, and electronic media*. Philadelphia: University of Pennsylvania Press.
- Gervás 2013** Gervás, P. (2013) "Propp's morphology of the folk tale as a grammar for generation", in M.A. Finlayson et al. (eds) *2013 Workshop on Computational Models of Narrative*, Vol. 32, pp. 106–22. Available at: <https://doi.org/10.4230/OASlcs.CMN.2013.106> (Accessed: April 2024).
- Gervás 2012** Gervás, P. (2012) "Story generator algorithms", in *The living handbook of narratology*. Hamburg: Hamburg University. Available at: <http://www.lhn.uni-hamburg.de/article/story-generator-algorithms> (Accessed: April 2024).
- Goldsmith 2011** Goldsmith, K. (2011) *Uncreative writing: Managing language in the digital age*. New York: Columbia University Press.

- Gontarski 2015** Gontarski, S.E. (2015) *Creative involution: Bergson, Beckett, Deleuze*. Edinburgh: Edinburgh University Press.
- Guan et al. 2020** Guan, J. et al. (2020) "A knowledge-enhanced pretraining model for commonsense story generation", *Transactions of the Association for Computational Linguistics*, 8, pp. 93–108. Available at: https://doi.org/10.1162/tacl_a_00302 (Accessed: April 2024).
- Guan et al. 2021** Guan, J. et al. (2021) "OpenMEVA: A benchmark for evaluating open-ended story generation metrics", in *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 6394–407. Available at: <https://doi.org/10.18653/v1/2021.acl-long.500> (Accessed: April 2024).
- Hayles 2008** Hayles, N.K. (2008) *Electronic literature: New horizons for the literary*. Notre Dame: University of Notre Dame Press.
- Heal 2016** Heal, B.J. (2016) "Authorship in Burroughs's Red Night trilogy and Bowles's translation of Moroccan storytellers", *CLCWeb: Comparative Literature and Culture*, 18(5), pp. 1–9. Available at: <https://doi.org/10.7771/1481-4374.2966> (Accessed: January 2026).
- Henrickson 2018a** Henrickson, L. (2018) "Butterflies, busy weekends, and chicken salad: Genetic criticism and the output of @Pentametrone", *Authorship*, 7(1). Available at: <https://doi.org/10.21825/aj.v7i1.8619> (Accessed: April 2024).
- Henrickson 2018b** Henrickson, L. (2018) "Tool vs. agent: Attributing agency to natural language generation systems", *Digital Creativity*, 29(2–3), pp. 182–90. Available at: <https://doi.org/10.1080/14626268.2018.1482924> (Accessed: April 2024).
- Henrickson 2021** Henrickson, L. (2021) *Reading computer-generated texts*. Cambridge: Cambridge University Press.
- Herrera-González, Gelbukh, and Calvo 2020** Herrera-González, B.D., Gelbukh, A. and Calvo, H. (2020) "Automatic story generation: State of the art and recent trends", *Advances in Computational Intelligence*, pp. 81–91. Available at: https://doi.org/10.1007/978-3-030-60887-3_8 (Accessed: April 2024).
- Hongisto 2023** Hongisto, T. (2023) "Reading computer-generated texts: Examining code as a reading strategy", *Digital Creativity*, 34(4), pp. 296–310. Available at: <https://doi.org/10.1080/14626268.2023.2274451> (Accessed: April 2024).
- Hua and Raley 2020** Hua, M. and Raley, R. (2020) "Playing with unicorns: AI Dungeon and citizen NLP", *Digital Humanities Quarterly*, 14(4).
- Ippolito et al. 2020** Ippolito, D. et al. (2020) "Toward better storylines with sentence-level language models", in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7472–78. Available at: <https://doi.org/10.18653/v1/2020.acl-main.666> (Accessed: April 2024).
- Jain et al. 2017** Jain, P. et al. (2017) "Story generation from sequence of independent short descriptions", in *SIGKDD Workshop on Machine Learning for Creativity (ML4Creativity)*, pp. 1–7. Available at: <https://doi.org/10.48550/arXiv.1707.05501> (Accessed: April 2024).
- Jefferson 2014** Jefferson, A. (2014) *Genius in France: An idea and its uses*. Princeton: Princeton University Press.
- Klein et al. 1973** Klein, S. et al. (1973) *Automatic novel writing: A status report*. Tech. Report 186. Madison: University of Wisconsin, Computer Science Department.
- Kybartas and Bidarra 2017** Kybartas, B. and Bidarra, R. (2017) "A survey on story generation techniques for authoring computational narratives", *IEEE Transactions on Computational Intelligence and AI in Games*, 9(3), pp. 239–53. Available at: <https://doi.org/10.1109/TCIAIG.2016.2546063> (Accessed: April 2024).
- Lamarque 2009** Lamarque, P. (2009) *The philosophy of literature*. Oxford: Blackwell.
- Landow 1992** Landow, G.P. (1992) *Hypertext: The convergence of contemporary critical theory and technology*. Baltimore: The Johns Hopkins University Press.
- Lebrun 2017** Lebrun, T. (2017) "Who is the artificial author?", in *Advances in Artificial Intelligence*, Vol. 10233. Cham: Springer International Publishing, pp. 411–15. Available at: https://doi.org/10.1007/978-3-319-57351-9_47 (Accessed: April 2024).
- Li, Ding, and Liu 2018** Li, Z., Ding, X. and Liu, T. (2018) "Generating reasonable and diversified story ending using sequence to sequence model with adversarial training", in *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 1033–43. Available at: <https://aclanthology.org/C18-1088> (Accessed: April 2024).

- Liu et al. 2020** Liu, D. et al. (2020) "A character-centric neural model for automated story generation", in *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(02), pp. 1725–32. Available at: <https://doi.org/10.1609/aaai.v34i02.5536> (Accessed: April 2024).
- Marino 2020** Marino, M. (2020) *Critical code studies*. Cambridge, Massachusetts: The MIT Press.
- Meehan 1976** Meehan, J. (1976) "The metanovel: Writing stories by computer". PhD dissertation. New Haven: Yale University, Computer Science.
- Montfort et al. 2013** Montfort, N. et al. (2013) *10 PRINT CHR\$(205.5+RND(1)); : GOTO 10*. Cambridge: The MIT Press.
- Montfort and Pérez y Pérez 2023** Montfort, N. and Pérez y Pérez, R. (2023) "Computational models for understanding narrative", *Revista de Comunicação e Linguagens*, 58, pp. 97–117. Available at: <https://doi.org/10.34619/gnzq-r7ri> (Accessed: April 2024).
- Natale 2021** Natale, S. (2021) *Deceitful media: Artificial intelligence and social life after the Turing test*. Oxford: Oxford University Press.
- O'Neill and Riedl 2014** O'Neill, B. and Riedl, M. (2014) "Dramatis: A computational model of suspense", in *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pp. 944–50. Available at: <https://doi.org/10.1609/aaai.v28i1.8836> (Accessed: April 2024).
- Paaß and Giesselbach 2023** Paaß, G. and Giesselbach, S. (2023) *Foundation models for natural language processing: Pre-trained language models integrating media*. Cham: Springer International Publishing.
- Parrish 2021** Parrish, A. (2021) *Language models can only write poetry*. Available at: <https://posts.decontextualize.com/language-models-poetry/> (Accessed: January 2026).
- Parrish 2025** Parrish, A. (2025) *Writing two of pentacles*. Available at: <https://posts.decontextualize.com/writing-two-of-pentacles/> (Accessed: January 2026).
- Peinado and Gervás 2006** Peinado, F. and Gervás, P. (2006) "Evaluation of automatic generation of basic stories", *New Generation Computing*, 24(3), pp. 289–302. Available at: <https://doi.org/10.1007/BF03037336> (Accessed: April 2024).
- Peng et al. 2018** Peng, N. et al. (2018) "Towards controllable story generation", in *Proceedings of the First Workshop on Storytelling*, June 2018, pp. 43–49. Available at: <https://doi.org/10.18653/v1/2022.findings-emnlp.520> (Accessed: April 2024).
- Peng et al. 2022** Peng, X. et al. (2022) "Inferring the reader: Guiding automated story generation with commonsense reasoning", in *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 7008–29. Available at: <https://doi.org/10.18653/v1/2022.findings-emnlp.520> (Accessed: April 2024).
- Porteous and Lindsay 2019** Porteous, J. and Lindsay, A. (2019) "Protagonist vs antagonist PROVANT: Narrative generation as counter planning", in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1069–77. Available at: <https://doi.org/10.5555/3306127.3331805> (Accessed: April 2024).
- Porteus and Lindsay 2019** Porteus, J. and Lindsay, A. (2019) "Protagonist vs antagonist PROVANT: Narrative Ggeneration as counter planning", in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 1069–1077.
- Pressman, Marino, and Douglass 2015** Pressman, J., Marino, M.C. and Douglass, J. (2015) *Reading project: A collaborative analysis of William Poundstone's Project for Tachistoscope {Bottomless Pit}*. Iowa City: University of Iowa Press.
- Propp 1968** Propp, V.Y. (1968) *Morphology of the folktale*, trans. by L. Scott. Austin: University of Texas Press.
- Rettberg 2018** Rettberg, S. (2018) *Electronic literature*. Cambridge: Polity.
- Riedl and Young 2010** Riedl, M.O. and Young, R.M. (2010) "Narrative planning: Balancing plot and character", *Journal of Artificial Intelligence Research*, 39, pp. 217–68. Available at: <https://doi.org/10.1613/jair.2989> (Accessed: April 2024).
- Ryan 2017** Ryan, J. (2017) "Grimes' fairy tales: A 1960s story generator", in N. Nunes, I. Oakley, and V. Nisi (eds) *Interactive Storytelling*, Vol. 39, pp. 89–103. Available at: http://dx.doi.org/10.1007/978-3-319-71027-3_8 (Accessed: April 2024).
- Ryan 1991** Ryan, M. (1991) *Possible worlds, artificial intelligence, and narrative theory*. Bloomington: Indiana University Press.

- Saemmer 2009** Saemmer, A. (2009) "Aesthetics of surface, ephemeral, re-enchantment and mimetic approaches in digital literature", *Neohelicon*, 36(2), pp. 477–88. Available at: <https://doi.org/10.1007/s11059-009-0016-2> (Accessed: April 2024).
- Schäfer 2007** Schäfer, J. (2007) "Gutenberg galaxy revis(it)ed: Brief history of combinatory, hypertextual and collaborative literature from the Baroque period to the present", in P. Gendolla and J. Schäfer (eds) *Writing, reading and playing in programmable media*. Bielefeld: Transcript Verlag, pp. 121–160.
- Schmidt 2010** Schmid, W. (2010) *Narratology: An introduction*. Berlin: de Gruyter.
- Schoenbeck 2013** Schoenbeck, R. (2013) "Playing with chance: On random generation in playable media and electronic literature", *Digital Humanities Quarterly*, 7(3). Available at: <https://dhq.digitalhumanities.org/vol/7/3/000165/000165.html> (Accessed: April 2024).
- Sharples and Pérez y Pérez 2022** Sharples, M. and Pérez y Pérez, R. (2022) *Story machines: How computers have become creative writers*. London: Routledge.
- Tambwekar et al. 2019** Tambwekar, P. et al. (2019) "Controllable neural story plot generation via reward shaping", in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI-19)*, pp. 5982–88. Available at: <https://doi.org/10.24963/ijcai.2019/829> (Accessed: April 2024).
- Turner 1994** Turner, S.R. (1994) *The creative process: A computer model of storytelling and creativity*. Hillsdale: Lawrence Erlbaum Associates.
- Veale 2014** Veale, T. (2014) "Coming good and breaking bad: Generating transformative character arcs for use in compelling stories", in *Proceedings of the 5th International Conference on Computational Creativity*, pp. 1–8. Available at: <https://computationalcreativity.net/iccc2014/proceedings/> (Accessed: April 2024).
- Wardrip-Fruin 2009** Wardrip-Fruin, N. (2009) *Expressive processing: Digital fictions, computer games, and software studies*. Cambridge, Massachusetts: MIT Press.
- Woodmansee 1994** Woodmansee, M. (1994) *The author, art, and the market: Rereading the history of aesthetics*. New York: Columbia University Press.
- Xu et al. 2018** Xu, J. et al. (2018) "A skeleton-based model for promoting coherence among sentences in narrative story generation", in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 4306–15. Available at: <https://doi.org/10.18653/v1/D18-1462> (Accessed: April 2024).
- Yao et al. 2019** Yao, L. et al. (2019) "Plan-and-write: Towards better automatic storytelling", in *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, pp. 7378–85. Available at: <https://doi.org/10.1609/aaai.v33i01.33017378> (Accessed: April 2024).
- van der Lee et al. 2019** van der Lee, C. et al. (2019) "Best practices for the human evaluation of automatically generated text", in *Proceedings of the 12th International Conference on Natural Language Generation*, pp. 355–68. Available at: <https://doi.org/10.18653/v1/W19-8643> (Accessed: April 2024).

Recommendations

DHQ is testing out three new article recommendation methods! Please explore the links below to find articles that are related in different ways to the one you just read. We are interested in how these methods work for readers—if you would like to share feedback with us, please complete our short evaluation survey. You can also visit our documentation for these recommendation methods to learn more.

SPECTER Recommendations

Below are article recommendations generated by the SPECTER model:

1. The Ends of Editing, 2009, Peter M. W. Robinson, University of Birmingham
2. Writing to be Found and Writing Readers, 2011, John Cayley, Brown University
3. Reading Note G: Ada Lovelace and the Clerical Labor of Codework, 2026, Zachary Mann, University of Southern California
4. Introduction: Reading Code in the Age of AI, 2026, Jeremy Douglass, University of California, Santa Barbara; Mark C. Marino, University of Southern California
5. The Literary And/As the Digital Humanities, 2013, Jessica Pressman, ACLS Fellow; Lisa Swanstrom, Florida Atlantic University

DHQ Keyword Recommendations

Below are article recommendations generated by DHQ Keywords:

1. To Visualize Past Communities: A Solution from Contemporary Practices in the Industry for the Digital Humanities, 2017, G  rald P  oux, Universit   Paris Ouest Nanterre La D  fense & Institut d'Histoire Moderne et Contemporaine (CNRS); Jean-Roch Houllier, Thales University
2. SpotiBot — Turing Testing Spotify, 2018, Pelle Snickars, Ume   University; Roger M  hler, Ume   University
3. Computational Linguistics and Classical Lexicography, 2009, David Bamman, Tufts University; Gregory Crane, Tufts University
4. Sounding for Meaning: Using Theories of Knowledge Representation to Analyze Aural Patterns in Texts, 2013, Tanya Clement, University of Texas, Austin; David Tch  ng, University of Illinois, Urbana-Champaign; Loretta Auvil, University of Illinois, Urbana-Champaign; Boris Capitanu, University of Illinois, Urbana-Champaign; Megan Monroe, University of Maryland, College Park
5. Cross-codex Learning for Reliable Scribe Identification in Medieval Manuscripts, 2024, Julius Wei  mann, Media and Digital Technologies, St. P  lten University of Applied Sciences; Markus Seidl, Media and Digital Technologies, St. P  lten University of Applied Sciences; Anya Dietrich, MEG Unit, Brain Imaging Center; Martin Haltrich, Library, Klosterneuburg Abbey

TF-IDF Recommendations

Below are article recommendations generated by the TF-IDF Model:

1. Playing with Chance: On Random Generation in Playable Media and Electronic Literature, 2013, Robert Schoenbeck, University of California, Irvine
2. The Machine in the Text, and the Text in the Machine, 2010, Manuel Portela, University of Coimbra
3. Stretched Skulls: Anamorphic Games and the memento mortem mortis, 2012, Stephanie Boluk, Vassar College; Patrick LeMieux, Duke University
4. Applied Digital Humanities and the Creative Industries in the United Kingdom, 2022, James Smithies, King's College London; Sarah Atkinson, King's College London; Elliott Hall, King's Digital Lab
5. Digital Editions and Version Numbering, 2020, Paul A. Broyles, North Carolina State University



This work is licensed under a Creative Commons Attribution-NoDerivatives 4.0 International License.