

5-10-2022

# Artificial Intelligence as a Content Creator in the Publishing Industry

John Huston  
*Portland State University*

Follow this and additional works at: [https://pdxscholar.library.pdx.edu/eng\\_bookpubpaper](https://pdxscholar.library.pdx.edu/eng_bookpubpaper)



Part of the [Communication Technology and New Media Commons](#), and the [Publishing Commons](#)

**Let us know how access to this document benefits you.**

---

## Recommended Citation

Huston, John, "Artificial Intelligence as a Content Creator in the Publishing Industry" (2022). *Book Publishing Final Research Paper*. 65.

[https://pdxscholar.library.pdx.edu/eng\\_bookpubpaper/65](https://pdxscholar.library.pdx.edu/eng_bookpubpaper/65)

This Paper is brought to you for free and open access. It has been accepted for inclusion in Book Publishing Final Research Paper by an authorized administrator of PDXScholar. Please contact us if we can make this document more accessible: [pdxscholar@pdx.edu](mailto:pdxscholar@pdx.edu).

# **Artificial Intelligence as a Content Creator in the Publishing Industry**

John Huston

English Department - Book Publishing, Portland State University

Dr. Rachel Noorda

May 10, 2022

## **RESEARCH QUESTION**

Given that Artificial Intelligence (AI) writes differently than humans, what place might there be in publishing houses for AI as content creators or collaborators?

## **METHODS**

Several methods were used in my research. A literature review was used to define artificial intelligence (AI) in general and in relation to how it's used in literature today. I also relied heavily on Dr Leah Henrickson's work on Natural Language Generation (NLG), which discusses authorship and roles between authors and readers.

Several case studies were used in my research. Liam Porr's Guardian Op-ed was reviewed as his team used humans to edit GPT-3 output for a coherent article, much like I used it in a focus group. The Kimagure Artificial Intelligence Writer's Project was also evaluated. The writing for the project was composed by humans and AI subsequently restructured the novel. I also examined several projects Andreas Refsgaard is working on. Refsgaard is a coder who uses his skills to influence art, and he is pushing the boundaries of collaboration between human artists and machines. He has created an online bookstore that sells AI generated Science Fiction novels. This helped me explore the idea of AI generated writing as a genre of its own.

A focus group was used to evaluate reader opinions of AI-written/enhanced work vs. human-written work. A human-written short story of fiction of approximately 1300 words was used as a baseline. That story was then segmented, taking out critical portions of the narrative, and AI was allowed to fill in the missing information. I used of GPT-3, the same software used in the Guardian Op-ed by Liam Porr, for the AI narrative. I then recruited eleven volunteers for

the focus group, asked them to read both samples, and answer questions about both stories afterword.

Finally, interviews were conducted with those working in the field to use or advance AI in writing. Dr. Leah Henrickson, who is a professor of Digital Media at the University of Leeds, gives a framework for understanding how computers learn through algorithms. She also helped shape my thoughts around the ethical use of AI. An interview with Liam Porr focused on how he worked with GPT-3, and which ways he found most effective to work with the program. Allison Parrish, a professor at New York University who uses bots to explore language and poetry, was interviewed to discuss the role bots play in relation to language. She was also instrumental in helping me see other ethical factors when exploring machine-generated content.

## **WHAT IS ARTIFICIAL INTELLIGENCE?**

As with most complex research, terminology is key. Here we will lay the groundwork for understanding the research question.

What is Artificial Intelligence (AI)? There is no single agreed upon definition, but the generally accepted standard text for teaching AI is Stuart Russell's and Peter Norvig's *Artificial Intelligence: A Modern Approach*. In it they describe AI by considering four concepts:

### **1. Human approach:**

Systems that think like humans

Systems that act like humans

### **2. Ideal approach:**

Systems that think rationally

Systems that act rationally

Machines can solve problems using any one of these singular concepts, or a combination of any or all of them.<sup>i</sup>

Too this point, AI has been successful in singular tasks: chess machines play chess, self-driving cars drive cars. There is some overlap, but the goal of a general or “deep” AI that is a sort of human-mimicking super intelligence appears to be a long way off.

Natural Language Generation (NLG) is the subfield of Artificial Intelligence research that turns non-readable (by humans) inputs into readable outputs: data into written language.<sup>ii</sup>

So, what is the difference between text editing and text generation? With something simple, like Microsoft Word, the degree of difference seems obvious. Word knows grammatical rules and points out when they are broken; it performs quality control, but it does not create. Grammarly, however, makes this distinction less clear. For example, Grammarly advertises the ability to make your writing feel more confident. In this way, Grammarly has learned some creative norms. It knows when a human generated passage doesn’t fit within a bucket it categorizes as confident. Here is one example Grammarly uses on its site:

Drafted text, “I think we should be able to solve this problem for you.”

Grammarly corrected text, “We can solve this problem for you.”

This example shows that the software knows more than basic grammatical rules, it knows how to change word choice to strike a particular tone. Because Grammarly knows how to define confident, then it will also know how to define its inverse, and in this way, it is beginning to examine human inputs and reflect on them based on its knowledge. This is the point at which a machine is influencing a human coauthor’s writing, but is it creating? What makes something like GPT-3 different from these systems is that it learns while it’s working. The human author or editor offers an input, the machine provides a response, I don’t like the response, so I change my

input, it offers a new response, I keep it, and it learns. Therefore, we say deep-learning machines are intelligent.

## **THE CENTAUR PROJECT**

Richard Beer is a Swedish journalist. He writes fiction novels—mainly crime thrillers—under the pen name Lars Amber. In an article for *Towards Data Science* he introduces the “Centaur concept.”<sup>iii</sup> He tells the story of Gary Kasparov, the chess grandmaster, who was beaten by IBM’s DeepBlue in 1997. Chess was one of the first areas AI researchers explored, and IBM developed DeepBlue to try to teach a machine to beat a human at one of humankind’s most strategic games. The next year (1998) Kasparov introduced “Centaur Chess,” a human playing chess and assisted by AI, which went on to beat AI players on their own. The concept: machines are great at process, humans have vision. The first attempts at directing AI to write novel-length works have born this out.

The first of these attempts was in 2014 when *The Day a Computer Wrote a Novel* was produced by the Kimagure Artificial Intelligence Writer Project. The Project was created by researchers at Nagoya University in Japan. They entered a submission for the Hoshi Shinichi Award, a Japanese writing contest that had the abstract rule that all writers—including non-human writers—could submit. The Hoshi Shinichi Award limits submissions to 10,000 Japanese characters, which is about 4000 English words.

The researchers divided the project into two teams. The first team used an existing AI program that was written to automate plot generation. AI gave them a plot, and the team of human authors wrote the story.<sup>iv</sup>

The second team used Natural Language Generation (NLG) software they created named Ghostwriter. Ghostwriter is made up of three modules: planner, configurator, and text generator.

The program requires three knowledge components: story grammar, a set of text fragments, and a set of configuration programs. These components can be thought of as data inputs. Ghostwriter is trained by these components via human-created text. A story is written then deconstructed into subsections, which creates a structural plan for Ghostwriter. Humans then program rules for those subsections and write small fragments of those subsections, which tells the NLG software how to create its version of the story. The software is further refined by adding more subsections and fragments to give Ghostwriter a larger database of information, and then controls are introduced to keep the AI author on track to write a coherent story.<sup>v</sup>

Of the two teams, the Ghostwriter team produced *The Day a Computer Wrote a Novel* which made it further through the rounds of judging for the Hoshi Shinichi Award than did the other Nagoya University submission. Neither, however, won the award, and subsequent interviews and articles indicate—although the accuracy of this is unclear—that most of the work was still done by human editors.

Andreas Refsgaard started [www.booksby.ai](http://www.booksby.ai), an online book shop for manuscripts completely written by machines. Many of the books in his webstore can be viewed in part on Amazon. It doesn't take long to see those novels written solely by AI are barely intelligible. Consider this opening passage from *The Imperfect Disaster*, by Barreast Wolf: "Institute and Sola had sheer if he had bent the seeing what aroused him. Now we have hidden these caverned goes if they entered a reindeer if I now." This is representative of AI writing unchecked. Sentences and paragraphs meld together into unintelligible garbles of words.

It's in shorter passages, and in collaboration with human authors, where machine-written literature begins to resemble works that a human would produce. In September of 2020 the Guardian newspaper, a 200-year-old publication, ran an op-ed written by a machine. Liam Porr,

a Berkley undergraduate student at the time, got permission from OpenAI to use their language-generating software, GPT-3, to see if it could write the piece.<sup>vi</sup> The op-ed is clear and succinct, although sounds as if a slightly angst-filled teen wrote it. After publishing the op-ed, Porr discussed the process: “All the content in the op-ed was taken from output of GPT-3, but not verbatim. It generated several outputs. And then the *Guardian* editors took the best outputs and spliced them together into this one large op-ed.”<sup>vii</sup>

When speaking with Liam Porr about the ways he found most effective to interact with GPT-3 to get the outcomes he sought, Porr explained that he found GPT-3 liked structured information, “bullet points, clearly established sections like ‘title,’ ‘subject,’ and ‘full text.’”<sup>viii</sup> This led our discussion to creativity and whether software engineers think of being creative when writing code. Initially, Liam said no, that algorithm design was about “task formulation and scope.”<sup>ix</sup> That it was about utility. This was an interesting area to consider.

Western tech culture may be the most results-driven culture on the planet. It was interesting listening to Liam—a software engineer himself—explain how GPT-3 worked. It was clear that the notion of a machine “creating” was not a part of the thought process when a language-generation engine was created. The algorithm was designed to solve a problem, and not to create a body of work. In literature, that problem is finishing a prescribed-length manuscript.

This is what I call the soulless algorithm: the boiling down of artistic creation to task, purpose, and scope. Liam and I explored the idea of creativity, which he initially thought of in aesthetic terms, and said it was something software engineers did not consider unless they were creating visual imagery. But, as we talked more, we agreed that engineers were creative in finding solutions, and in that way, there was creativity in the design of machine-learning algorithms. But does this task-driven mindset skirt more dire issues?



## **ETHICS**

Ethics, like the humanities, weave their way into our lives and decision making, which makes our training in ethics hard to spot on any given day. More scientific pursuits, like math or engineering, have fairly linear signs of success or failure. Our ideas around ethics and AI in literature can be evaluated in a few different ways: bias, homogeneity, and content ownership. They are all different.

### **Bias**

Analogies help when thinking of intelligent machines. Artificial Intelligence is like a brain. It has a database of information to draw on, which we may correlate to our memory. For the machine its database is the internet, or intranets that it has access to.

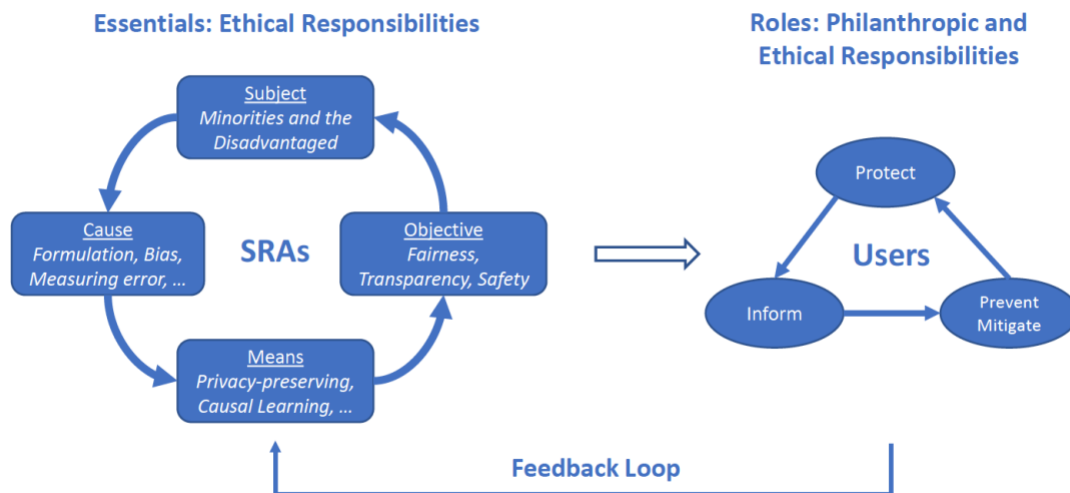
We have behaviors which are generally products of our experiences. We use these behaviors to make decisions based on the information we have stored or at hand. AI also has behaviors, but in its case, they are named algorithms. Algorithms allow machines to learn, to decipher context, and to adapt to unfamiliar environments.<sup>x</sup>

Algorithmic bias would seem our biggest concern when training machines to sift through reams of literature, then produce competent works based on what it has reviewed. Deep learning machines continue to learn on their own, so it can become unclear—over time—how their decisions are made, and how biases have influenced those decisions.

We do know that these algorithms are susceptible to bias. Several AI systems used to predict criminal recidivism have identified higher levels of re-offense among BIPOC individuals.

When specific algorithms have been evaluated to look for the root of the bias, none have been found, which would indicate a broader problem in the data provided to the machine.<sup>xi</sup>

Lu Cheng, Kush Varshney, and Huan Liu look at this in their work on AI bias, and they focus on algorithm construction as the solution.<sup>xii</sup> They point to numerous areas where bias can be injected into an AI system—often in supervised (as in human supervised) machine learning—and acknowledge that transparency is often not achievable. Their solution is socially responsible AI algorithms or SRA's. Here they begin a discussion of how to write trustworthy algorithms for machines, and they give us this framework to start with:



Imke Van Heerden and Emil Bas expand on this idea when it applies directly to literature. They argue in their paper on machine learning and literary theory that in order to move AI to the next level when writing literature, it needs to be trained in more than just the technical components of literary theory, but also the machine must understand things like creativity and imagination.<sup>xiii</sup> They promote including literary experts—professors and authors—to accomplish this.

But this is another gateway for bias if we don't watch for it. Whom do we choose for this role? What do we acknowledge as literature? Even if the whole of English-language literature was at a machine's disposal, that is still only the minds of those who have written in English. A comprehensive and inclusive definition of literature will be our first step. Within it, consideration of language, culture, and form should all be considered. It is here that we confront the problem of homogeneity.

### **Homogeneity**

There are a few problems with machines and writing when it comes to homogeneity. The first is the corpus it draws from: the internet. As Dr. Hendrickson pointed out in our interview, the unofficial language of coding is English.<sup>xiv</sup> The fact that the backbone of the internet has been created on a language used by Western culture most certainly injects bias from the beginning. Taken a step further, if we contemplate how much is not on the internet, such as stories in spoken languages such as Sarsi (a Native American language that has no known written form), it becomes clear that using AI to create literature based on a huge, but still limited body of work like the internet is fraught with areas where many cultures, and their literary works, may be overlooked.

The second may be the idea that we can predict what word should come next according to Allison Parrish. The word "should" does a lot of work in this context. Allison explained that although we have conventions for language, much of how language is interpreted is unseen; contained within the experience of the reader.<sup>xv</sup> That means machines will have a tough time choosing the next "right" word for the most eloquent output. In a most basic example, I will lean on a blogpost that Allison wrote about Firefox's spell-check function. She looked up the word "transphobia" and Firefox put a little-red warning line under the word, one she describes as,

“Perhaps,” the algorithm says to me, speaking through this red underline, “You meant to type something else.”<sup>xvi</sup> What this means is the further we move the human away from the machine, the more likely that “the next right” word is based on a solution instead of an interpretation by the reader.

### **Content Ownership**

Content ownership is peripheral to the machine, but important in that it may prioritize machines over human content creators. Large corporations fund research in the arts to legitimize the technologies they are using. Allison was clear that she worried that she may be over-judging corporate money in research against an ethical instinct that it may influence her work. But those misgivings don’t minimize the impact that money has on machine-manipulated language.

The ethics of machines in art—in this case literature—is when machines kidnap ownership of content from its creators. Machines are extensions of their corporate creators. In talking with Allison Parrish, she points out that, “Companies take content from the web, compress it with machines, then transfers ownership from authors to companies.”<sup>xvii</sup> By this she means AI uses the internet as a corpus of material. It learns from that body of work, internalizes it to the machine’s respective servers, then those companies create content with it that they own and protect with resources exponentially higher than any author can access.

For all of these reasons—bias, homogeneity, and content ownership—creators of AI language generation programs, and those that hope for their success—we must oversee the overseers. But readers have a vote in all of this as well.

## READERS ACCEPTING AI AUTHORS


In my focus group I presented two versions of a short fiction story.\* Sample H1 was completely written by a human author. The second version, labeled G3, started with the same story (H1), but it was dissected and key portions of it were fed to GPT-3, which would then contribute its input based on where it thought the story was or should be going.<sup>1</sup> I asked focus group participants the following questions:

- a. “Which sample better immerses you in the setting of the story?”
- b. “Which sample better connects you to the protagonist of the story?”
- c. “Which sample feels better paced to you as a reader?”
- d. “Do you have any personal observations about either story you would like to elaborate on?”

This is what I found:

### 1. “Which sample better immerses you in the setting of the story?”

11 responses




---

\* The recruitment question as well as more detailed focus group information can be found in Appendix 1.

<sup>1</sup> A note on working with GPT-3 is relevant here. I found while working with the software that the smaller the prompt given the machine, the more the human author, or editor in my case, could control its tangential tendencies. Fed large blocks of texts—paragraphs of three or more sentences—sent GPT-3 in entirely new directions. These directions were not uninteresting, but they did get increasingly disconnected from the story’s premise. I have to acknowledge the challenge of this task: I did not want to defeat the purpose of my own research by massaging the AI collaboration too much but knowing where that line existed was an imperfect balance.


2. "Which sample better connects you to the protagonist of the story?"

11 responses



3. "Which sample feels better paced to you as a reader?"

11 responses



I was surprised by the answers. Readers were almost evenly split in how immersed they were in the story, and they also felt the pacing was comparable. There was a greater connection to the protagonist in the AI-generated story, which was more surprising. This would indicate a reading public that may accept AI contribution to literature, but my focus group readers weren't told that one example was co-authored by AI. They were asked, "I am studying the reception of short fiction in publishing. If you would be interested in helping, please respond to this email and I will send you two, short pieces of fiction and the link to a (very short) survey." So, in this case, anti-AI bias was not a factor.

In her paper *Computer-Generated Fiction in a Literary Lineage: Breaking the Hermeneutic Contract* Dr. Leah Henrickson says, “Modern readers appear unable to separate the work from its creator: the two are conceptually bound, as we search for the author embedded within the text to contextualize its production.”<sup>xxviii</sup> How then will readers see a machine as co-author? Dr. Henrickson explains the relationship between reader and author via the Hermeneutic Contract. Which she describes as, “that reading is accepted as a social act wherein the reader receives an interesting and understandable text written by an author motivated by an intent-directed agency.”<sup>xxix</sup> Does this indicate an anti-AI bias that readers may not be able to overcome?

She goes on to say, “Computer-generated texts in their current state, however, bring the hermeneutic contract into question. The hermeneutic contract’s communication principle rests on the assumption that readers believe that authors want their texts to be both interesting and understandable.”<sup>xxx</sup> How can a machine *want* their text to be anything? Does this break the compact with the reader, and thereby turn them away from works crafted by AI?

But in her book *Reading Computer-Generated Texts* Henrickson says, “Texts have no inherent meaning, but are assigned meaning by human writers and readers. Further, computational agency complements human agency...”<sup>xxxi</sup> This is in line with Allison Parrish’s view that, “...contained within the experience of the reader.”<sup>xxxii</sup> It is also not far from our Grammarly example: that “computational agency compliments human agency,” but on much large scale of course.<sup>xxxiii</sup> We will need a larger body of coherent, machine-generated literature in the world to fully understand how well it will be accepted.

## CONCLUSION

Is there a place in publishing houses for machines as content creators? Unequivocally yes. For now, the technology isn't proficient enough to rely on AI to write sole-source stories, let alone novel-length manuscripts. Only the most cutting-edge publishers may integrate machine-authored works, much like some producers promote extremely experimental music. But we have found that readers are open to reading the works of AI, and, the fact is, we are already too late.

AI is a fact. The wealthiest governments and corporations in the world are racing to create the most sophisticated versions of machine-learning technologies. As an example, in 2017 AlphaZero, a deep-learning machine created by DeepMind, was taught the rules of chess, but, unlike its predecessors, did not play other players or bespoke chess-playing machines at all in its training. It was only allowed to play itself, or in other words, it was only allowed to learn from itself. AlphaZero went on to beat the previous chess-master program, and it did it in just four hours.<sup>xxiv</sup>

So, AI is here to stay in all domains. Language is unique in that it is a medium of creation and also a conveyer of information, therefore it is especially important that in publishing we make sure it matures in a representative way. Dr. Leah Henrickson said it masterfully, "Words are the only way we can express what is in our minds to others."<sup>xxv</sup> When machines begin to use their words to pass on information, knowledge, perhaps even wisdom, we must ensure those words are not just the words of the powerful. Perhaps the final paragraph of *The Day a Computer Wrote a Novel* sums it up best:

*"I writhed with joy, which I experienced for the first time, and kept writing with excitement. The day a computer wrote a novel. The computer, placing priority on the pursuit of its own joy, stopped working for humans."*



- 
- <sup>i</sup> Russell, Stuart, and Peter Norvig. "Artificial intelligence: a modern approach." (2002).
- <sup>ii</sup> Reiter, Ehud, and Robert Dale. "Building applied natural language generation systems." *Natural Language Engineering* 3, no. 1 (1997): 57-87.
- <sup>iii</sup> Beer, Richard. 2021. "Are You Going to Read a 'Centaur Book' Soon?" Medium. February 16, 2021. <https://towardsdatascience.com/are-you-going-to-read-a-centaur-book-soon-44368fed8829>.
- <sup>iv</sup> Cassel, D. (2016, April 16). The Day a Computer Wrote a Novel That Almost Won a Literary Award. Retrieved May 3, 2022, from <https://thenewstack.io/day-computer-wrote-novel-almost-won-literary-competition/>
- <sup>v</sup> Sato, Satoshi. 2016. "A Challenge to the Third Hoshi Shinichi Award." <https://aclanthology.org/W16-5505.pdf>
- <sup>vi</sup> GPT-3. 2020. "A Robot Wrote This Entire Article. Are You Scared Yet, Human? | GPT-3." *The Guardian*, September 8, 2020, sec. Opinion. <https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3>.
- <sup>vii</sup> Smith, Laura and Worthington Leah, "The Edge Episode 7: Hey Siri, Write Me a Poem." The Edge, May 19, 2021. Podcast, website, 46:56. <https://alumni.berkeley.edu/california-magazine/just-in/2020-12-03/the-edge-episode-7-robot-poetry>
- <sup>viii</sup> Porr, Liam. Interview by John Huston. Video Interview. Zoom, January 27, 2022.
- <sup>ix</sup> Porr, Liam. Interview by John Huston. Video Interview. Zoom, January 27, 2022.
- <sup>x</sup> Cheng, Lu, Kush R. Varshney, and Huan Liu. "Socially responsible ai algorithms: Issues, purposes, and challenges." *arXiv preprint arXiv:2101.02032* (2021).
- <sup>xi</sup> Danks, David, and Alex John London. "Algorithmic Bias in Autonomous Systems." In *IJCAI*, vol. 17, pp. 4691-4697. 2017.
- <sup>xii</sup> Cheng, Lu, Kush R. Varshney, and Huan Liu. "Socially responsible ai algorithms: Issues, purposes, and challenges." *arXiv preprint arXiv:2101.02032* (2021).
- <sup>xiii</sup> van Heerden, Imke, and Anil Bas. "AI as Author—Bridging the Gap Between Machine Learning and Literary Theory." *Journal of Artificial Intelligence Research* 71 (2021): 175-189. <https://doi.org/10.1613/jair.1.12593>.
- <sup>xiv</sup> Henrickson, Leah. Interview by John Huston. Video Interview. Zoom, January 6, 2022.
- <sup>xv</sup> Parrish, Allison. Interview by John Huston. Video Interview. Zoom, February 1, 2022.
- <sup>xvi</sup> Parrish, A. (2021, August 14). Desire (under)lines: Notes toward a queer phenomenology of spell. Retrieved May 3, 2022, from <https://posts.decontextualize.com/queer-in-ai-2021/>
- <sup>xvii</sup> Parrish, Allison. Interview by John Huston. Video Interview. Zoom, February 1, 2022.
- <sup>xviii</sup> Henrickson, Leah. "Computer-generated fiction in a literary lineage: Breaking the Hermeneutic contract." *Logos* 29, no. 2-3 (2018): 54-63.
- <sup>xix</sup> Henrickson, Leah. "Computer-generated fiction in a literary lineage: Breaking the Hermeneutic contract." *Logos* 29, no. 2-3 (2018): 54-63.
- <sup>xx</sup> Henrickson, L. (2021). *Reading Computer-Generated Texts*. Cambridge University Press.
- <sup>xxi</sup> Henrickson, L. (2021). *Reading Computer-Generated Texts*. Cambridge University Press.
- <sup>xxii</sup> Parrish, Allison. Interview by John Huston. Video Interview. Zoom, February 1, 2022.
- <sup>xxiii</sup> Henrickson, L. (2021). *Reading Computer-Generated Texts*. Cambridge University Press.
- <sup>xxiv</sup> Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K., & Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419), 1140-1144. <https://doi.org/10.1126/science.aar6404>
- <sup>xxv</sup> Henrickson, Leah. Interview by John Huston. Video Interview. Zoom, January 6, 2022.



## Bibliography

- Beer, Richard. 2021. "Are You Going to Read a 'Centaur Book' Soon?" Medium. February 16, 2021. <https://towardsdatascience.com/are-you-going-to-read-a-centaur-book-soon-44368fed8829>.
- Cassel, D. (2016, April 16). The Day a Computer Wrote a Novel That Almost Won a Literary Award. Retrieved May 3, 2022, from <https://thenewstack.io/day-computer-wrote-novel-almost-won-literary-competition/>
- Cheng, Lu, Kush R. Varshney, and Huan Liu. "Socially responsible ai algorithms: Issues, purposes, and challenges." *arXiv preprint arXiv:2101.02032* (2021).
- Danks, David, and Alex John London. "Algorithmic Bias in Autonomous Systems." In *IJCAI*, vol. 17, pp. 4691-4697. 2017.
- Gross, Sabine. "Cognitive readings; or, the disappearance of literature in the mind." (1997): 271-297.
- Hall, Stefan. 2018. "AI Will Write a Best-Seller by 2049, Experts Predict." World Economic Forum. March 2018. <https://www.weforum.org/agenda/2018/03/timeline-of-creative-ai/>.
- Henrickson, Leah. "Computer-generated fiction in a literary lineage: Breaking the Hermeneutic contract." *Logos* 29, no. 2-3 (2018): 54-63.
- Henrickson, Leah. *Reading Computer-Generated Texts*. Cambridge University Press, 2021.
- Henrickson, Leah. Interview by John Huston. Video Interview. Zoom, January 6, 2022.
- Lombardo, Giorgia. 2019. "What Happens When AI Writes a Novel?" DeMagSign. December 3, 2019. <https://medium.com/demagsign/what-happens-when-ai-writes-a-novel-173d407af567>.
- Margolis, Eric. n.d. "The Day a Computer Wrote a Novel." Ideas.newsrx.com. Accessed November 9, 2021. <https://ideas.newsrx.com/blog/the-day-a-computer-wrote-a-novel>.
- Parrish, Allison. Interview by John Huston. Video Interview. Zoom, February 1, 2022.
- Porr, Liam. Interview by John Huston. Video Interview. Zoom, January 27, 2022.
- Russell, Stuart, and Peter Norvig. "Artificial intelligence: a modern approach." (2002).
- Reiter, Ehud, and Robert Dale. "Building applied natural language generation systems." *Natural Language Engineering* 3, no. 1 (1997): 57-87.
- Samuel, Sigal. 2019. "How I'm Using AI to Write My next Novel." Vox. August 30, 2019. <https://www.vox.com/future-perfect/2019/8/30/20840194/ai-art-fiction-writing-language-gpt-2>.

Sato, Satoshi. 2016. "A Challenge to the Third Hoshi Shinichi Award." <https://aclanthology.org/W16-5505.pdf>

Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., Graepel, T., Lillicrap, T., Simonyan, K., & Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419), 1140-1144. <https://doi.org/10.1126/science.aar6404>

Smith, Laura and Worthington Leah, "The Edge Episode 7: Hey Siri, Write Me a Poem." The Edge, May 19, 2021. Podcast, website, 46:56. <https://alumni.berkeley.edu/california-magazine/just-in/2020-12-03/the-edge-episode-7-robot-poetry>

van Heerden, Imke, and Anil Bas. "AI as Author—Bridging the Gap Between Machine Learning and Literary Theory." *Journal of Artificial Intelligence Research* 71 (2021): 175-189

Are You Scared Yet, Human? | GPT-3." *The Guardian*, September 8, 2020, sec. Opinion. <https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3>.

"Will AI Write the next Great American Novel?" 2021. Cal Alumni Association. March 15, 2021. <https://alumni.berkeley.edu/california-magazine/spring-2021/will-ai-write-the-next-great-american-novel>.

## Appendix 1 (Focus Group Comments)

1. G3 is quicker paced dispersed with dialogue, imagery, and setting.
2. There were more sensory details in G3 that helped establish the setting and the characters. The formatting of G3 sometimes made it harder to understand than H1.
3. It is unclear whether these stories are complete. If they are meant to be whole, short stories, there are so many questions that do not get answered, namely why they live in these pods, why they have done so much "following," and what exactly the children are doing when they are sent away.
4. I liked the ambiguous ending of H1 more than the ending of G3. G3 was a smidge more confusing, especially the end when it seems like The Waste is being personified, then we switch back to Frank's POV.
5. I liked the dialogue in G3. I felt more connected to the protagonist in H1 and felt that it showed more than told. And to me, G3 felt more like it leaned on exposition than H1.
6. I liked the added details that were added to G3 that I did not see in H1. Especially the ending
7. I read each story once and thought I was drawn to one more than the other. But then confusion set in and I read them one more times. I feel equally somewhat confused and my thoughts are surrounded by suspicious feeling, not knowing what it is that I should have noticed, etc. I am going to go with my first instinct, which is that I think I might like the writing of the first one better, but cannot tell you exactly why. There are slight differences in the beginning of the stories and the children are referred to in the second one right away, but not in the first one, but overall I think they are very similar.
8. I personally felt like there wasn't enough space between the dialog and narration in G3, and that's why I felt disconnected.
9. G3 drew me in faster than H1, and it felt more authentic to me. I felt more connected to the characters in this story and the pacing felt less disjointed.
10. Both stories featured an engaging setting and humanized the protagonist effectively, but H1 included more details, leaving me less confused. G3 was clean and quick to read, making it more fun, but perhaps less impactful.
11. I would say the immediacy of the story is more apparent in G3. This might be to the dialogue and the subtle hints in characterization in this dialogue. With such a short story, it is important to immediately capture the reader, to get to the point faster, to try to do away with all of the excess/background noise in order to get at the core of the story. G3 get us there quicker. It might be different in novels and in longer fiction. But with short stories, I think starting off the conflict as soon as possible, as in G3, is the better option.

## Appendix 2 (Abstract)

### **ABSTRACT**

This paper explores the current capabilities of Artificial Intelligence (AI) in generating literature, and where AI authorship may fit into publishing. AI researchers are pushing machines into almost all cognitive realms, and those machines practicing literary writing are trying to achieve the machine-written novel. That capability appears to be a distant target, but machines are writing, and they are doing it with the help of human co-authors. Publishers and consumers have long accepted new genres and their authors in literature. Now may be the time for publishing houses to explore the idea of adding AI into their houses as content creators. To examine this, I interviewed Leah Hendrickson whose research focuses on NLG and authorship, conducted a case study that assessed AI co-authorship and its appeal to readers, and interviewed the two publishing houses owners who seek innovative stories to understand their willingness or resistance to publishing AI co-authored texts.