

 **ERC Advanced Grant 2023**
Research proposal [Part B1]
*(Part B1 is evaluated both in Step 1 and Step 2,
Part B2 is evaluated in Step 2 only)*

AI STORIES: Narrative Archetypes of Artificial Intelligence

Cover Page:

- - Jill Walker Rettberg (PI)
- - University of Bergen
- - 60 months

AI STORIES is premised on the hypothesis that narrative archetypes fundamentally structure the output of contemporary artificial intelligence (AI). Large language models (LLMs) like GPT-4 are trained on vast quantities of text and images and generate new texts that are statistically similar to the training data. The scientific consensus acknowledges that LLMs replicate and sometimes exacerbate historical biases in their training data.

AI STORIES proposes that LLMs are also affected by a deeper bias: that of the narrative structures in the social media posts, news stories, marketing blurbs and novels the models are trained on. If this is the case it will deeply impact how we use and apply AI, and how we think about bias and cultural diversity in AI models. Currently available LLMs are largely trained on English-language texts, with a heavy weighting towards the United States. When they generate texts in non-English languages they may succeed in producing grammatically correct texts, but if my hypothesis is correct, their deeper content will be fundamentally structured by the stories that dominate in the training data. This is a threat to cultural diversity that goes well beyond the purely linguistic.

AI STORIES applies the humanities' deep knowledge of narrative to AI research by developing and testing this hypothesis. We will apply narratology to understand the narrative structures of LLM's training data. We test the hypothesis by training LLMs on specific kinds of narratives, then using prompt engineering and both qualitative and computational narratological analysis to reverse engineer the structures of AI-generated output. Three comparative case studies will look specifically at Scandinavian, Australian and either Indian or Nigerian stories.

The overall objective is to develop a narratology of AI, and to leverage the findings to ensure that policymakers, developers, educators and other stakeholders can use our research to direct the future of AI.

Section a: Extended Synopsis of the scientific proposal

“I’m Sydney, and I’m in love with you.” These words were generated by a large language model (LLM) released by Microsoft in January 2023 in a two-hour conversation with a journalist from the New York Times. The words stunned readers across the world: was this AI sentient, or just crazy? I argue that it is neither.

Current AI is trained on the stories humans tell. Sydney’s declaration of love is the result of a statistical analysis of vast quantities of human stories and cultural documents. “I’m Sydney, and I’m in love with you” is exactly the sort of thing you would expect an AI to say in a story.

The AI STORIES project will test and explore the hypothesis that narrative archetypes in the training data of LLMs like Sydney fundamentally affect the output of these AI models. The technical fields of machine learning (ML) and artificial intelligence (AI) have developed extremely rapidly in the last few years. With OpenAI’s release of ChatGPT in late November 2022, generative AI and LLMs have gone mainstream. Despite important critiques and analyses from the humanities and social sciences, almost all *development* of AI has been within computer science and related disciplines, such as data science, statistics and mathematics. But the current wave of AI is based on language, narratives and culture. We have reached a point where AI development needs the humanities. AI STORIES addresses this need.

AI Stories will provide the basis for a new research paradigm that combines qualitative and quantitative epistemologies for the future development of AI. This is crucial for all LLMs, as demonstrated by Sydney’s ill-fated declaration of love. It is particularly important for Europe, because our stories are not the same as the stories that OpenAI, Google and Baidu are training their LLMs on. Let me give you an example: Thorbjørn Egner’s *Folk og røvere i Kardemomme By* (*Folk and Robbers in Cardamom Town*) is a children’s book and musical well known to anyone who grew up in Norway or Denmark after 1950. The songs and audio play have been played on the radio and in homes and kindergartens for decades – there is even a theme park in Kristiansand. The story features three comical robbers who steal food because they are hungry and don’t understand that work is necessary. After being caught stealing sausages and chocolate they are rehabilitated by the kind police officer and townsfolk and end up saving the town from a fire. This story is not just a shared cultural reference, it is a cultural support for the Norwegian criminal justice system’s focus on rehabilitation above punishment. Think now of Hollywood stories about robbers or criminals: Bank heist and gangster movies glorify criminals; Disney movies have unambiguous villains who die at the end of the movie.

Narrative archetypes give us templates to live by. *Kardemomme By* and Hollywood movies are not just shared cultural references: they also structure the way we think about the world. Depending on the stories we hear, share and create, we shape possibilities for action and for understanding: Criminals can be rehabilitated, or villains deserve to come to a bad end. Deeper structures have also been identified by the structuralist narratologists of the 20th century: A promise or contract made at the start of a story is bound to be broken. A story goes from a balance to imbalance to a new balance. A hero goes on a journey and returns, but with a difference.

The **objective of AI STORIES is to develop a narratology of AI:** a theoretical understanding of how narrative in the training data affects the models and outputs of large language models. I will develop this narratology by researching two main research questions:

- RQ1: How do narrative archetypes in training data shape the functioning and output of large language models?
- RQ2: To what extent do the narrative archetypes of different cultures make machine learning models culturally specific?

This is basic research that will have a significant scientific impact on narratology and the study of narratives in literary studies, media studies and internet studies. If my hypothesis is correct and I can communicate the results effectively AI STORIES will also have a significant scientific impact on research and development in machine learning and AI, as it will demonstrate that developing good AI requires cross-disciplinary collaborations between computer scientists, humanities scholars and social scientists who understand the narrative and cultural underpinnings of the training data. There is also a strong potential for societal impact, particularly on policies for regulating AI (e.g. to preserve cultural diversity) and on education, both to educate future developers of AI, and to educate the general public. Digital literacy will require understanding the cultural specificity of the large language

Key terms

LLM: Large Language Model - an AI model that can generate new text after being trained on a very large amount of text.
 Training dataset: the text that an LLM is trained on to build its model.
 GPT: Generative Pre-Trained Transformer. Introduced in 2017, a type of machine learning that can generate text or images that are statistically similar to the models training data.

models that are becoming part of everyday life.

State of the Art and beyond

Storytelling is core to human culture, and our ability to organise events as stories has been posited as integral to the development of language (Ferretti et al. 2017). Narratives are based on structures that are repeated and varied to express different ideas. In *Poetics* Aristotle described the structure of tragedies as having a beginning, a middle and an end, with the dramatic arc climaxing and ending in catharsis. Twentieth century narratologists went into more detail, inspired by the structuralist analysis of language. Vladimir Propp identified 27 elements in folk tales in 1928 (Propp 1968), and in the second half of the 20th century, the structuralist study of narratives became well-established (Levi-Strauss 1955; Greimas 1970; Genette 1980; Brooks 1992). With the “narrative turn” of the 1990s narrative became a category and even method used broadly across the humanities and social sciences (Brooks 2022; Czarniawska-Joerges 2004; Presser 2016; Brown et al. 1996; Russo 2021).

My hypothesis is in a sense structuralist: **I posit that there are archetypal narratives that are resurfaced by LLMs and applied to language generation outside of their original cultural context.** Such archetypal narratives would be similar to Lyotard’s *grand récit* (1984), to Lévi-Strauss’s mythological structures that lie above language (1955), Campbell’s *monomyth* (1949), Frye’s narrative archetypes (1957) Abbot’s masterplot (2002), Tosca’s archetypes (2023) or the “patternment” that linguist Benjamin Whorf argued lay at a scale beyond morphemes and syntax and that was *more* significant than vocabulary or grammar (Whorf 1952). However, I am **not arguing for a universal narrative archetype.** On the contrary, my hypothesis is that there are culturally specific and narrative structures that might be national or particular to another cultural grouping, and that LLMs threaten to flatten out this diversity by implementing just one set of cultural narratives in text generation. I am also not arguing for a hard determinism where these narrative archetypes determine how humans tell stories or that they are universal, although I do think they have an influence on human storytelling—we draw upon archetypes, but we also often rebel against them or alter them, remix them, counter them, parody them, reflect upon them and contextualise them. I presume that these archetypes are fluid and intersectional.

AI is not human. **LLMs are “stochastic parrots”** (Bender et al. 2021): they do not understand or interpret language, they simply identify statistical patterns in what we input and respond with an output of words that are statistically probable continuations of our input, based on the training data. As philosopher Luciano Floridi notes, parrots are actually quite intelligent; LLMs are not (2023). I find N. Katherine Hayles’s specification of *nonconscious cognition* and *technical cognition* (2017) and her analysis of cyber-semiotics (Hayles 2019) to be helpful in understanding what Floridi calls agency without intelligence, because it enables a non-anthropocentric understanding that avoids the vague notion of “intelligence”. In my first ERC project I built upon this to theorise the specific agency of AI in image recognition and image generation, which works very similarly to text generation (J. W. Rettberg 2023). This means that **while a purely structuralist approach to human storytelling is reductive, it may be appropriate to AI-generated stories.**

Artificial intelligence has been used to generate stories since the 1950s (Strachey 1952; J. W. Rettberg 2021b). Before LLMs, narrative generation was done by explicitly coding story grammars into programs, for instance in systems like Talespin or MEXICA (S. Rettberg 2018; Sharples and Pérez y Pérez 2022). In contrast, LLMs like ChatGPT and Sydney use machine learning. In effect they use statistical linguistics to predict which words are likely to follow other words. Transformer models like GPT-4 use deep learning to infer patterns in vast quantities of texts and images. They no longer need the rules of grammar to be programmed into them: statistical analysis of enough sentences allow them to generate grammatically correct text in multiple languages. Although LLMs are still very new, authors are experimenting with creating literary work using LLMs (Johnston 2018; Strickland 2019; Amerika 2022; S. Rettberg et al. 2023). Literary scholars are likewise beginning to develop methodologies for analysing texts generated by LLMs (Alexander et al. 2021; S. Rettberg et al. 2023), and for using LLMs to analyse other texts (Querubín and Niederer 2022; Underwood 2023), but this is at a very early stage. There are also intriguing but untested ideas coming out of the AI enthusiast communities online, where discussions about adversarial prompts that “jailbreak” LLMs to go against their rules also inspire researchers (Wolf et al. 2023). Some of these play upon narrative tropes, which can be used to elicit certain kinds of response (Nardo 2023), as creative authors have also explored.

The humanities and social sciences have contributed significantly to research on AI over time. Science and technology studies (STS) have addressed AI and digital technologies for decades. One thread of AI criticism can be traced through analyses of quantification and concepts of accuracy (Mackenzie 1993) and classification (Bowker and Star 2008), which have been foundational for critical dataset studies today (Thylstrup 2022). In Internet Studies there have been many critical analyses of AI as implemented in social media platforms and recommendation systems (Seaver 2018; Bucher 2018; Amooore 2020) and more generally in society (Crawford 2021). Feminist scholars and Black scholars have analysed biases in AI (Losh and Wernimont 2018; Benjamin 2019; Chun 2022). Others also see hope and possibility in using AI (Bridle 2022). AI is seen to represent an “epistemic universality” and is a tool that has “been overgeneralized into epistemic destiny” (Hong 2022).

Narrative archetypes vs AI bias: AI STORIES is premised on a novel hypothesis: the output of LLMs is fundamentally affected by narrative archetypes in the datasets that the models are trained on. This goes beyond the current scientific consensus, which has found that biases in the datasets are replicated but sees this as primarily being about the proximity of two signs to each other. For example, in most datasets, words like *woman* or *she* are statistically more likely to appear close to *nurse* than to *doctor* (Bolukbasi et al. 2016). LLMs trained on such datasets will therefore predict that nurses are women, and doctors are men, thus replicating and even increasing the bias in the dataset. This type of AI bias is now broadly recognised with extensive research in the public and private sector attempting to mitigate it. My hypothesis goes beyond this by looking at the deeper narrative structures of the training data rather than just the syntagmatic sequences of signs.

Methodology

I have designed the project to have three partly overlapping stages, first providing the theoretical groundwork, then a more experimental stage where we test the theories, and finally the synthesis and communication of results, where we develop and communicate a narratology of AI. Three contrastive case studies will provide a basis for comparison with the “default” narratives generated by American LLMs. I will study Scandinavian, Australian and either Indian or Nigerian narratives. The first two are chosen due to my own cultural knowledge, coming from an Australian family but having mostly lived in Norway since I was eight. The third case is chosen because India and Nigeria have a large English-speaking culture but a very different narrative tradition than the United States. A final choice will be made based on the expertise of the postdocs hired.

Stage 1: Theorising

Objective: Review literature and develop a theoretical framework for understanding the narratology of the datasets LLMs are trained in. Provide basis for experimental phase in Stage 2.

Outputs: Literature reviews on sub-areas, special issue of journal, workshops, papers, and a co-authored monograph: *A Narratology of Datasets*. A summarising literature review will synthesize this work and provide a platform for stage 2 and 3.

First we need to build the groundwork. I will consolidate existing knowledge drawing upon narratological theory, anthropology, linguistics, internet studies and critical dataset studies, hiring two postdocs to review the literature and work with me on a cross-disciplinary theoretical synthesis. My extensive experience using narratology to analyse digital genres (Walker 1999; 2003; 2005; Corneliussen and Rettberg 2008; J. W. Rettberg 2009; 2014; 2021a) and my research on AI (J. W. Rettberg 2022; 2023) is an ideal basis for this work. Two postdocs will support me in this, one focusing on narratology and one focusing more on the cross-cultural aspects of narratives drawing upon fields such as cultural and linguistic anthropology and ethnographic research on legends and folklore. We will also survey publications in computer science, digital humanities, digital culture, critical algorithm studies and critical dataset studies to better understand the training data and the architecture of LLMs. Three workshops will structure this stage: one on dataset narratology, one on cross-cultural narratologies, and one on narratology and AI.

Stage 2: Training and testing

Objectives: Test the hypothesis.

Outputs: Qualitative narratological analyses of generated output and analyses using computational narratological methods. A paper using statistical analysis of narrative archetypes in training data and in generated output to confirm or disprove the hypothesis. A monograph by the PI, *A Narratology of AI*. Documented and shared code (AI models) and datasets.

The experimental phase of the project begins after the first 18 months and aims to test the hypothesis that narrative archetypes are reproduced in the output of LLMs and to find to what extent culturally specific structures can be identified. Two main strategies will be used to generate data for analysis: 1) **Prompt engineering of mainstream LLMs.** In 2023 we would use publicly available LLMs like ChatGPT, Bing’s chat mode and Google’s Bard, but there will likely be other models by 2025. We will also consider using open source models like BLOOM. Since companies like OpenAI and Google do not share much information about the systems we would “reverse engineer” the LLMs by prompting them in multiple ways and analysing the outputs. 2) **Train new AI models on datasets selected to emphasise particular kinds of narratives.** We would start simply, by training models on datasets consisting only of one genre of narrative: e.g. Norwegian folktales, Indigenous Australian climate fiction, 19th century childrens’ literature, or fan fiction based on a culturally specific narrative like SKAM, aiming to find edge cases. We will also try counter-factual and speculative approaches, e.g. mixing genres or altering the training data as in Jake Elwes’ *Queering Datasets* project (2021). Fine-tuning or adding alignment layers to existing models is also a possibility.

We will try a range of methodologies for analysing the output models, designing precise strategies according to our findings in Stage 1 and the state of the art in 2025. One approach is **prompt engineering**, a rapidly developing field we will follow closely (White et al. 2023; Underwood 2023). This allows us to generate multiple outputs that can then be analysed using **qualitative narratological analysis, content analysis, data visualisation and computational narratological analysis**. Digital humanities methodologies for analysing narratives will be useful here (Antoniak, Mimno, and Levy 2019; McDowall, Antoniak, and Mimno 2023; Tangherlini et al. 2020) and we will also consider more **speculative methods** (Hamilton and Piper 2022; Querubín and Niederer 2022) and methods that use LLMs and other kinds of machine learning in the analysis (J. W. Rettberg 2022; Underwood 2023). I will hire a PhD student with a background in literary studies for the prompt engineering and analysis, and a postdoc with a background in digital humanities or computational linguistics to train the new models and work with computational methods. I have also budgeted for bringing in visiting researchers specialising in computational narratology.

This is the stage at which I hope to **confirm or disprove my hypothesis** that narrative archetypes influence the output of LLMs. Although the bulk of my research will be **qualitative** (because narrative is not an exact science) I will also test more **quantitative** approaches within this stage of the project. Based on findings from stage 1, we will identify narrative archetypes that are present in specific training data and look for corresponding archetypes appearing in the output. For this to be statistically valid we will need to develop methods for avoiding the analysis of the output being influenced by prior knowledge of the project, for instance by randomising output, using automated analysis or by using coders who are not familiar with the project. The research will also have a strong qualitative component more aligned with literary studies than statistical disciplines. In addition to testing the hypothesis, this stage of the project will contribute to **developing new methods for combining qualitative and quantitative methodologies in the study of LLMs**. Three workshops will structure this stage, the first two experimental and the third discussing results.

Stage 3: Synthesizing and leveraging: A narratology of AI

Objective: Synthesise findings to develop an overall narratology of AI. Communicate findings to developers (tech industry) and to stakeholders in education, cultural sector, journalism, government.

Outputs: Monograph: *A Narratology of AI*. Policy recommendations. Findings are communicated to stakeholders in formats that are relevant and accessible to different audiences.

The final stage involves synthesizing the results from the first two stages. Based on the findings from stage 1 and stage 2, we will develop a **narratology of AI**. In addition to papers on specific aspects of the narratology, I will publish a book, provisionally titled *A Narratology of AI*, to be written in a rigorous but accessible style making it suitable to be used as a university textbook, and to be read by both literature scholars and computer scientists. This accessibility is important if the research is to have impact across multiple disciplines.

In addition to synthesising the results, this stage of the project will focus on communicating and leveraging the findings. Communicating research is time-consuming but essential, especially when trying to impact multiple disciplines in addition to policymakers. To ensure **strategic and effective communication**, I will hire a full-time administrative project manager with expertise in research communication. They will develop strategies for reaching stakeholders including not only scientists in humanities fields but also AI developers, computer scientists, policymakers and educators. Targeted public talks and media contributions will be combined with one-to-one networking and other strategies to reach key actors in AI development and policy. I aim to have a high-profile paper in *Nature* or *Science*, for the research to be discussed in mainstream international media such as *The Guardian*, *Le Monde* and the *Washington Post* and to present at major venues for AI research.

My capacity to lead the project

I am uniquely positioned to research the narrative basis for generative AI, both due to my personal research background as a literary scholar who researched AI-driven technologies in my first ERC project *Machine Vision in Everyday Life: Playful Interactions with Visual Technologies in Digital Art, Games, Narratives and Social Media* (2018-2024, grant agreement number 771800), and as the Co-Director of the Center for Digital Narrative (CDN) at the University of Bergen. My colleague Scott Rettberg and I won a prestigious grant from the Norwegian Research Council to build this world-leading Center of Excellence from 2023-2033. There is excellent potential for synergies with other projects at CDN, but there is no replication or direct overlap between AI STORIES and CDN. When we applied for funding for CDN, LLMs were at a far less advanced stage, and so our focus was on what we termed *algorithmic narratives*, that is, how AI can be used in collaboration with humans to generate narratives, how recommendation and moderation algorithms promote or censor certain stories in social media, and how narratives are developing in video games, electronic literature and new genres like virtual reality or the augmented reality promised in the hyped ideas about the Metaverse. AI STORIES comes at digital narratives from the opposite perspective. While **CDN investigates how digital**

media are changing the way we tell stories, AI STORIES asks how the stories we have told for centuries are shaping AI. AI STORIES will also draw upon the UiB AI network, of which I am a steering committee member, and which fosters cross-disciplinary research on AI across the University of Bergen. Both CDN and UiB AI provide AI STORIES with a strong platform for communicating our research broadly.

Table 1: Timeline of the project, showing the three stages and the people who will make up the research team.

	2024	2025	2026	2027	2028	2029
Stages	THEORISING					
		TRAINING AND TESTING				
				SYNTHESIZING AND LEVERAGING		
People	PI: 40% time commitment, fully funded by ERC					20% time comm.
	Postdoc 1 (4 years: 75% ERC-funded, UiB pays additional 25% for teaching)					
	Postdoc 2 (4 years: 75% ERC-funded, UiB pays additional 25% for teaching)					
	Postdoc 3 (4 years: 75% ERC-funded, UiB pays additional 25% for teaching)					
	PhD (4 years, UiB-funded, 25% teaching)					
	Project and Communication Manager 100% - administrative position fully funded by ERC					
	Data management expert 20%					
	Visiting researcher					
	Research assistant 20%					
Events						
	Kick-off conference		Workshops every half year			Synthesising conference

High Risk, High Gain

This is high risk research. First, **AI technologies are changing so rapidly** that this project could not have been proposed just a year ago. It is probable that AI will make further unexpected leaps that alter the premises for my research. Stage 2 of the project, the training and testing, is most at risk from this, so I will use methods from agile development and design thinking to develop iterative processes where we produce minimal products that can be tested early. This enables us to change strategy rapidly if an approach doesn't work or when technology changes. Secondly, **the hypothesis is bold and testing it requires new methodologies**. Preliminary evidence is compelling: we know there are strong narrative elements in the training data, the outputs of LLMs clearly have narrative elements, and if the hypothesis that models that replicate narrative structures in the training data is correct that would also help explain AI bias. There is however a risk that my research will disprove my hypothesis. In this case, the project would still develop narratological tools to analyse the web texts typical of the training data and to analyse outputs of LLMs, which would be of great value to the study of digital narratives and automatically generated texts. The methodological innovation is also a risk, but I have the appropriate experience and project design to succeed. Even if the computational analysis and model training were to fail, we could still use conventional literary analysis on output from LLMs.

The gain, especially if my hypothesis is correct, is extremely high and **will enable significant scientific advances across multiple disciplines**. Understanding the link between narrative and AI is a necessary but very challenging leap for research on narratives in literature and other media, as current methods were developed to analyse human-crafted stories and must now contend with stories generated by AI. The findings will also be significant for intercultural research, linguistics and computer science, and will have profound impact on the development of AI in computer science and the technology sector, and on policy. Understanding how narratives in the training data affect narrative outputs of LLMs will have implications for how AI is designed and regulated in areas ranging from the preservation of cultural heritage and diversity to managing the spread of AI-generated propaganda and disinformation. A narratology of AI will also be needed in other disciplines that rely on narrative either as a methodology, as in psychology, history or the "case method" in law. As these fields adopt AI (whether in clinical practice, courts of law or scientific research) they will need to understand how the AI works: what stories is it telling and why? Is it more likely to tell certain kinds of story? Which ones?

AI STORIES connects the humanities to emerging technologies to develop **urgently needed knowledge** about the AI-generated texts that will be increasingly integrated in jobs, education and media across society. This is exactly the sort of research that ERC grants were designed to support: I have **a bold and ambitious hypothesis that could not have been formulated even a year ago**. I have the knowledge and experience needed to test my hypothesis, which could have profound impact on multiple scientific disciplines as well as significant societal impact.

If this research is not done now, we risk losing aspects of our cultural heritage to AI models that flatten cultural differences without anyone even understanding what is happening.

References

- Abbott, H. Porter. 2002. *The Cambridge Introduction to Narrative*. Cambridge, UK ; New York, NY, USA: Cambridge University Press.
- Alexander, Anne, Caroline Bassett, Alan Blackwell, and Jo Lindsay Walton. 2021. *Ghosts, Robots, Automatic Writing: An AI Level Study Guide*. Cambridge, UK: Cambridge Digital Humanities.
- Amerika, Mark. 2022. *My Life as an Artificial Creative Intelligence*. Sensing Media: Aesthetics, Philosophy, and Cultures of Media. Stanford, California: Stanford University Press.
- Amoore, Louise. 2020. *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Durham: Duke University Press.
- Antoniak, Maria, David Mimno, and Karen Levy. 2019. 'Narrative Paths and Negotiation of Power in Birth Stories'. *Proceedings of the ACM on Human-Computer Interaction* 3 (CSCW): 1–27. <https://doi.org/10.1145/3359190>.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Margaret Mitchell. 2021. 'On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🐦'. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–23. Virtual Event Canada: ACM. <https://doi.org/10.1145/3442188.3445922>.
- Benjamin, Ruha. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. Cambridge, UK: Polity.
- Bolukbasi, Tolga, Kai-Wei Chang, James Zou, Venkatesh Saligrama, and Adam Kalai. 2016. 'Man Is to Computer Programmer as Woman Is to Homemaker? Debiasing Word Embeddings'. *ArXiv*, July. <http://arxiv.org/abs/1607.06520>.
- Bowker, Geoffrey C., and Susan Leigh Star. 2008. *Sorting Things out: Classification and Its Consequences*. 1. paperback ed., 8. print. Inside Technology. Cambridge, Mass.: MIT Press.
- Bridle, James. 2022. *Ways of Being: Animals, Plants, Machines: The Search for a Planetary Intelligence*. New York: Farrar, Straus and Giroux.
- Brooks, Peter. 1992. *Reading for the Plot: Design and Intention in Narrative*. Harvard University Press.
- . 2022. *Seduced by Story: The Use and Abuse of Narrative*. New York: New York Review Books.
- Brown, Brian, Peter Nolan, Paul Crawford, and Alison Lewis. 1996. 'Interaction, Language and the "Narrative Turn" in Psychotherapy and Psychiatry'. *Social Science & Medicine* 43 (11): 1569–78. [https://doi.org/10.1016/S0277-9536\(96\)00053-6](https://doi.org/10.1016/S0277-9536(96)00053-6).
- Bucher, Taina. 2018. *If... Then: Algorithmic Power and Politics*. New York: Oxford University Press.
- Campbell, Joseph. 1949. *The Hero With a Thousand Faces*. New York: Pantheon Books.
- Chun, Wendy Hui Kyong. 2022. *Discriminating Data: Correlation, Neighborhoods, and the New Politics of Recognition*. Cambridge, MA: MIT Press.
- Corneliussen, Hilde, and Jill Walker Rettberg. 2008. *Digital Culture, Play, and Identity: A World of Warcraft Reader*. MIT Press.
- Crawford, Kate. 2021. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven: Yale University Press.
- Czarniawska-Joerges, Barbara. 2004. *Narratives in Social Science Research*. London: Sage Publications.
- Elwes, Jake. 2021. 'Queering Datasets'. In . Cambridge, UK. <https://youtu.be/5ukrHDnm6rQ>.
- Ferretti, F., I. Adornetti, A. Chiera, S. Nicchiarelli, R. Magni, G. Valeri, and A. Marini. 2017. 'Mental Time Travel and Language Evolution: A Narrative Account of the Origins of Human Communication'. *Language Sciences* 63 (September): 105–18. <https://doi.org/10.1016/j.langsci.2017.01.002>.
- Floridi, Luciano. 2023. 'AI as Agency Without Intelligence: On ChatGPT, Large Language Models, and Other Generative Models'. *Philosophy & Technology* 36 (1): 15, s13347-023-00621-y. <https://doi.org/10.1007/s13347-023-00621-y>.
- Frye, Northrop. 1957. *Anatomy of Criticism: Four Essays*. Princeton: Princeton UP.
- Genette, Gérard. 1980. *Narrative Discourse: An Essay in Method*. Translated by Jane E. Lewin. Ithaca NY: Cornell UP.
- Greimas, A. J. 1970. *Du Sens: Essais Sémiotiques*. Paris: Seuil.
- Hamilton, Sil, and Andrew Piper. 2022. 'The COVID That Wasn't: Counterfactual Journalism Using GPT'. <https://doi.org/10.48550/ARXIV.2210.06644>.
- Hayles, N. Katherine. 2017. *Unthought: The Power of the Cognitive Nonconscious*. Chicago: University Of Chicago Press.
- . 2019. 'Can Computers Create Meanings? A Cyber/Bio/Semiotic Perspective'. *Critical Inquiry* 46 (1): 32–55. <https://doi.org/10.1086/705303>.
- Hong, Sun-ha. 2022. 'Predictions without Futures'. *History and Theory* 61 (3): 371–90. <https://doi.org/10.1111/hith.12269>.
- Johnston, David Jhave. 2018. *ReRites: Human + A.I. Poetry*. <http://glia.ca/rerites/>.
- Levi-Strauss, Claude. 1955. 'The Structural Study of Myth'. *The Journal of American Folklore* 68 (270): 428. <https://doi.org/10.2307/536768>.
- Losh, Elizabeth M., and Jacqueline Wernimont, eds. 2018. *Bodies of Information: Intersectional Feminism and Digital Humanities*. Debates in the Digital Humanities. Minneapolis: University of Minnesota Press.
- Lyotard, Jean-François. 1984. *The Postmodern Condition: A Report on Knowledge*. Theory and History of Literature, v. 10. Minneapolis: University of Minnesota Press.
- Mackenzie, Donald. 1993. *Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance*. Cambridge, MA: MIT Press.
- McDowall, LeAnn, Maria Antoniak, and David Mimno. 2023. 'Sensemaking About Contraceptive Methods Across Online Platforms'. <https://doi.org/10.48550/ARXIV.2301.09295>.
- Nardo, Cleo. 2023. 'The Waluigi Effect (Mega-Post)'. *LessWrong.Com* (blog). 3 March 2023. <https://www.lesswrong.com/posts/D7PumeYTDPfBTp3i7/the-waluigi-effect-mega-post>.

- Presser, Lois. 2016. 'Criminology and the Narrative Turn'. *Crime, Media, Culture: An International Journal* 12 (2): 137–51. <https://doi.org/10.1177/1741659015626203>.
- Propp, Vladimir. 1968. *Morphology of the Folktale*. Translated by Laurence Scott. 2nd ed. Austin, Texas: Indiana University.
- Querubin, Natalia Sánchez, and Sabine Niederer. 2022. 'Climate Futures: Machine Learning from Cli-Fi'. *Convergence: The International Journal of Research into New Media Technologies*, October. <https://doi.org/10.1177/13548565221135715>.
- Rettberg, Jill Walker. 2009. "'Freshly Generated for You, and Barack Obama": How Social Media Represent Your Life'. *European Journal of Communication* 24 (4): 451–66. <https://doi.org/10.1177/0267323109345715>.
- . 2014. *Blogging*. 2nd ed. Cambridge: Polity Press.
- . 2021a. "'Nobody Is Ever Alone": The Use of Social Media Narrative to Include the Viewer in SKAM'. *The Journal of Popular Culture* 54 (2): 232–56. <https://doi.org/10.1111/jpcu.13015>.
- . 2021b. 'Speculative Interfaces: How Electronic Literature Uses the Interface to Make Us Think about Technology'. *Electronic Book Review*, October. <https://electronicbookreview.com/essay/speculative-interfaces-how-electronic-literature-uses-the-interface-to-make-us-think-about-technology/>.
- . 2022. 'Algorithmic Failure as a Humanities Methodology: Machine Learning's Mispredictions Identify Rich Cases for Qualitative Analysis'. *Big Data & Society* 9 (2). <https://doi.org/10.1177/20539517221131290>.
- . 2023. *Machine Vision: How Algorithms Are Changing the Way We See the World*. Basingstoke: Polity Press.
- Rettberg, Scott. 2018. *Electronic Literature*. Cambridge: Polity.
- Rettberg, Scott, Talan Memmott, Jill Walker Rettberg, Jason Nelson, and Patrick Lichty. 2023. 'AIwriting: Relations Between Image Generation and Digital Writing'. <https://doi.org/10.48550/ARXIV.2305.10834>.
- Russo, Maria Teresa. 2021. 'Ricoeur's Hermeneutic Arc and the "Narrative Turn" in the Ethics of Care'. *Medicine, Health Care and Philosophy* 24 (3): 443–52. <https://doi.org/10.1007/s11019-021-10020-9>.
- Seaver, Nick. 2018. 'Captivating Algorithms: Recommender Systems as Traps'. *Journal of Material Culture*, December, 1359183518820366. <https://doi.org/10.1177/1359183518820366>.
- Sharples, Mike, and Rafael Pérez y Pérez. 2022. *Story Machines: How Computers Have Become Creative Writers*. Abingdon, Oxon ; New York: Routledge.
- Strachey, Christopher. 1952. *M.U.C. Love Letter Generator*.
- Strickland, Stephanie, ed. 2019. *ReRites: Human + A.I. Poetry: Raw Output: A.I. Trained on Custom Poetry Corpus: Responses: 8 Essays about Poetry and A.I.* Montreal, Quebec: Anteism Books.
- Tangherlini, Timothy R., Shadi Shahsavari, Behnam Shahbazi, Ehsan Ebrahimzadeh, and Vwani Roychowdhury. 2020. 'An Automated Pipeline for the Discovery of Conspiracy and Conspiracy Theory Narrative Frameworks: Bridgegate, Pizzagate and Storytelling on the Web'. Edited by Yu-Ru Lin. *PLOS ONE* 15 (6): e0233879. <https://doi.org/10.1371/journal.pone.0233879>.
- Thylstrup, Nanna Bonde. 2022. 'The Ethics and Politics of Data Sets in the Age of Machine Learning: Deleting Traces and Encountering Remains'. *Media, Culture & Society* 44 (4): 655–71. <https://doi.org/10.1177/01634437211060226>.
- Tosca, Susana. 2023. *Sameness and Repetition in Contemporary Media Cultures*. S.l.: Emerald.
- Underwood, Ted. 2023. 'Using GPT-4 to Measure the Passage of Time in Fiction'. *Using Large Digital Libraries to Advance Literary History* (blog). 19 March 2023. <https://tedunderwood.com/2023/03/19/using-gpt-4-to-measure-the-passage-of-time-in-fiction/>.
- Walker, Jill. 1999. 'Piecing Together and Tearing Apart: Finding the Story in "Afternoon".' In , edited by Klaus Tochtermann, Jörg Westbomke, Uffe K. Wiil, and John J. Leggett, 111–17. ACM Press.
- . 2003. 'Fiction and Interaction: How Clicking a Mouse Can Make You Part of a Fictional World'. Dr. art. thesis, University of Bergen. <http://www.ub.uib.no/elpub/2003/d/517001>.
- . 2005. 'Distributed Narrative: Telling Stories Across Networks'. In *The 2005 Association of Internet Researchers Annual*, edited by Mia Consalvo, Jeremy Hunsinger, and Nancy Baym, 91–102. New York: Peter Lang.
- White, Jules, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C. Schmidt. 2023. 'A Prompt Pattern Catalog to Enhance Prompt Engineering with ChatGPT'. <https://doi.org/10.48550/ARXIV.2302.11382>.
- Whorf, Benjamin Lee. 1952. 'Language, Mind, and Reality'. *ETC: A Review of General Semantics* 9 (2): 167–88.
- Wolf, Yotam, Noam Wies, Yoav Levine, and Amnon Shashua. 2023. 'Fundamental Limitations of Alignment in Large Language Models'. <https://doi.org/10.48550/ARXIV.2304.11082>.