

BackChat: When AI Learns Language in Reverse



BackChat is a language model trained from scratch to generate text and conversation in reverse. Instead of predicting the next word in a sequence, *BackChat* generates what came before, scrambling the assumed temporal direction of thought. Answers produce questions, punchlines lead to setups, and effects precede their causes. Rather than merely inverting text generation, this work challenges the underlying assumptions of AI reasoning, exposing its reliance on statistical continuity rather than human-like cognition. By unsettling familiar patterns of interaction, *BackChat* makes the limitations of AI reasoning more visible and invites speculation on whether intelligence requires causality at all. In doing so, it opens up the possibility of intelligences structured around entirely different temporal logics that do not depend on the stepwise, linear progression that most AI models reinforce.

1. Introduction

In recent years, discussions about Large Language Models (LLMs) have moved beyond technical circles, becoming central in art, academia, and public discourse. The widespread adoption of systems like ChatGPT¹ has raised critical questions about how these models shape language, knowledge, and interaction. Artists and researchers have sought to engage with these technologies not just as tools but as sites of experimentation, asking what LLMs reveal about intelligence, meaning, and the constraints of machine learning architectures (Elliott and Grierson 2024, Zeilinger 2021). This has led to questions, not only about whether these models *think* or *reason*, but about how different architectures might structure intelligence in ways that depart from human assumptions (Audry 2021, Bridle 2022).

Our artwork, *BackChat*, explores these questions through a training intervention: reversing the predictive logic of an LLM. Generative Pre-trained Transformer (GPT) models are trained to predict the next token in a sequence (i.e. next word in a sentence), constructing meaning step by step into the future. This comes from a position that intelligence arises from the brain's ability to predict future events based on past experiences (Hinton 2007). This has led to what researchers have called a *Reversal Curse* where "If a model is trained

Isaac Clarke

Computational Media and Arts, Hong Kong
University of Science and Technology, Guangzhou, China
ijclarke590@connect.hkust-gz.edu.cn

Theo Papatheodorou

Computational Media and Arts, Hong Kong
University of Science and Technology, Guangzhou, China
theodoros@hkust-gz.edu.cn

Keywords Artificial Intelligence, Speculative AI,
LLM, Experiential AI

DOI [10.34626/2025_xcoax_004](https://doi.org/10.34626/2025_xcoax_004)

on a sentence of the form A is B, it will not automatically generalize to the reverse direction B is A” (Berglund et al. 2024).²

With *BackChat* we invert this structure, by training it to generate text backward, predicting what came before rather than what comes next. Instead of unfolding a story from a prompt of “Once upon a time,” we input “The End” and the model works in reverse. Instead of asking a question, we enter an answer and see what was asked. Instead of providing the setup of a joke, we give the punchline. This small shift in training alters the experience of chatting with AI, creating a *jamais-vu* feeling, something conversational but never thinking like a human.

This inversion makes two key aspects of AI clearer. First, it highlights what LLMs are actually doing: not reasoning, but statistically producing the most likely next token based on patterns in training data. By experiencing this process in reverse, users gain an alternative perspective on the fluency and coherence of LLM-generated text. Second, and a little contradictory to the first, it raises the idea of AI having forms of intelligence that do not mirror human cognition. Machine intelligence is often framed in human-like terms, but what if we explored architectures that fully embrace the machine’s unique computational properties? Computers excel at tasks that are very difficult for humans, such as writing coherent text in reverse. *BackChat* takes steps towards activating this possibility. Not just critiquing AI’s limits, but testing alternative training methods to explore what feelings of intelligence could emerge when we step outside of human assumptions of reasoning.

We situate this work within ongoing discussions in computational arts around Experiential AI, where artistic interventions serve as Explainable AI (XAI) methods that make the mysteries of algorithmic structures “vividly apparent” (Hemment et al. 2023b), and Graspable AI (Ghajargar et al. 2021), which explores how AI processes can be made tangible. *BackChat* pushes further into Speculative AI, not just revealing how current models function but experimenting with changing to architectures to imagine how other intelligences might emerge.

This paper explores the conceptual grounding, technical implementation, and artistic contributions of *BackChat*, positioning it as a method for engaging with AI on its own terms.

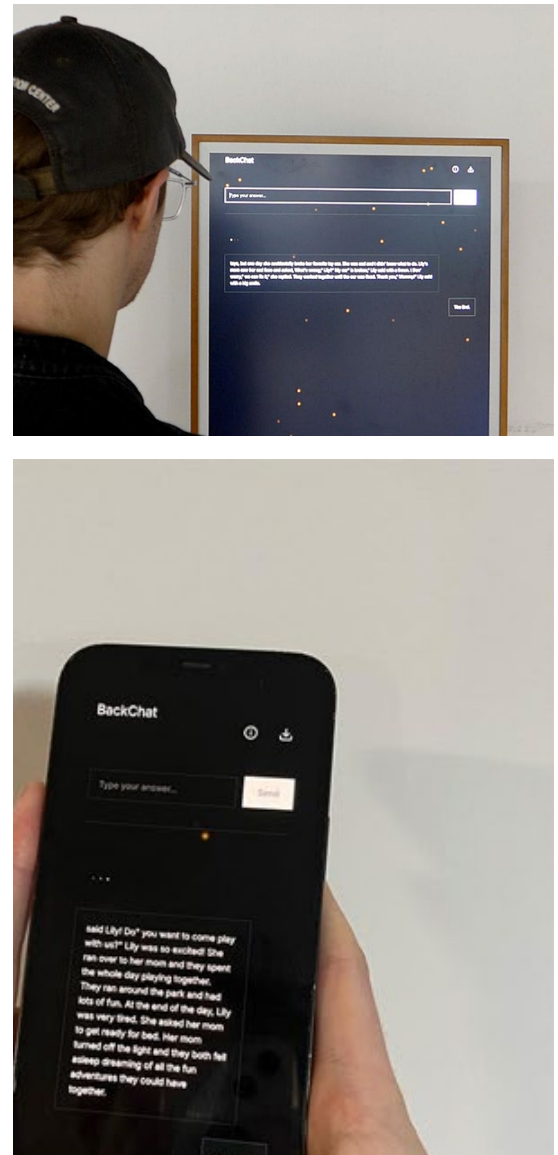


Fig. 1. Photo of the installation version (top), photo of the mobile chat app (bottom).

2. Background

2.1. Explainable and Experiential AI

Explainable AI (XAI) is a field of research originating in computer science that aims to enhance transparency and interpretability in AI systems (Miller 2019). While conventional AI systems often function as *black boxes*, XAI seeks to clarify not only what decision was made, but why the system arrived at that decision. This can involve techniques such as visualising internal model activations, designing surrogate models that mimic the behaviour of more complex systems, and developing analytical tools for monitoring neural networks. The goal is to provide insights into AI processes that would otherwise remain opaque, making models more accountable, understandable, and trustworthy.

Hemment et al. introduced Experiential AI (Hemment et al. 2023b), expanding XAI's focus from explanation to interpretation through experience. Instead of simply offering transparent systems, Experiential AI proposes that users should engage with AI's inner workings directly—feeling its processes rather than merely being told about them. This approach has been used in AI literacy projects (Hemment et al. 2023a). Related to this, in the field of tangible and embodied interaction, is the idea of Graspable AI (Ghajargar et al. 2021), which explores how tangible interfaces and material artefacts can provide alternative ways to hold and comprehend, and in turn critique, AI systems. This aligns with broader research in Tangible AI (McGrath et al. 2022), which suggests that embedding AI into physical objects or interactive environments can allow for a deeper understanding of its mechanisms.

These approaches emphasise that AI systems should not just be explained but made visible, tangible, played with, and interacted with in ways that can reveal otherwise hidden behaviours or qualities. This aligns with the Creative AI Lab's call to go “beyond the study of artefacts to focus on the 'back-end' environments that have enabled their production” (Bunz et al. 2023). In our work we continue with these aims, taking a practical approach that detours from usual implementations of AI, not just to critique AI but to experience how alternative architectures might reshape interaction, interpretation, and experience of alien meaning and reasoning.

2.2. Artistic Interventions in AI and Language

The use of AI in artistic practice has often been framed through a tension between human and non-human intelligence. Artists working

with AI have long experimented with alternative forms of intelligence, showing that machine cognition can introduce non-human ways of constructing meaning (Cohen 1995, Rokeby 1998). These investigations span machine perception, poetics, and methodologies that challenge the conventional framing of AI. Rather than treating AI as a tool that mimics human reasoning, artistic interventions explore AI as systems with their own emergent forms of sense-making.

Orphan Drift's work shows a long-standing engagement with other intelligences found in places such as machine vision.³ Jenna Sutela's *nimiia cétii* imagines a system that is not built for human understanding but instead operates as a kind of speculative translation between machine, organism, and abstract form, where the computer is an "alien of our creation" (Sutela 2018) rather than a tool designed to replicate human expression.

Beyond language as an alien system, other artists engage with the structural and poetic potential of Deep Neural Networks (DNN) for text generation, treating them less as models of intelligence and more as dynamic, expressive systems. Memo Akten and Mick Grierson explored how multiple LSTMs trained on different writers' texts can be blended together through an interactive interface (Akten and Grierson 2017). By allowing different AI *voices* to be merged or transitioned between in real-time, Akten and Grierson highlight how AI-generated language is not a singular process but an emergent space shaped by multiple influences. Allison Parrish has also explored DNN as poetic systems, her work highlighting the role training data and model architecture has in shaping outputs, and showing that meaning in AI is often a byproduct of structure rather than intelligence (Temkin 2018).

The works above suggest that language models can be more than passive generators of text, as an active space for rethinking cognition and interpretation. Our work builds on this by introducing a training intervention that alters AI's predictive structure. In doing so we are shifting the focus from what LLMs generate as a static output, to how LLMs construct meaning sequentially over time.

2.3. AI, Time, and Reasoning

Reversing the predictive logic of AI raises a core question: does reasoning require causality? Conventional language models, such as GPT-based systems, operate sequentially, predicting the next token based on past context. This structure reflects a deeply human-centric view of thought, where meaning emerges through forward progression and cause-effect relationships. The AI pioneer Geoffrey Hinton

posits that intelligence emerges from predicting future states based on prior experience (Hinton 2007). Recent developments, such as DeepSeek’s R1, reinforce this by optimizing models for “human-like reflection” by rewarding the effective use of <think></think> annotations to signal reasoning steps. (Guo et al. 2025). The underlying aim of most LLMs is to approximate general intelligence, assuming that if AI can mimic human-like thought closely enough, it qualifies as reasoning. This ties into a broader narrative of technological progress as a linear, causal trajectory.

This can be connected to wider debates on non-human temporality and algorithmic reasoning. Parisi challenges the idea that algorithmic operations are merely mechanical computation and argues that machine learning systems develop their own internal logic, shaping how they process and interpret information (Parisi 2016). Rather than dismissing algorithmic reasoning as non-thinking, we must recognise how it is actively shaping our world. Hui extends his critique to the history of technology, contending that technological development does not follow a universal trajectory but embeds distinct ways of thinking and cultural assumptions within its architecture (Hui 2021). From these perspectives, LLMs do not just simulate human thought but develop machine-specific logic. LLMs are producing forms of reasoning that arise from their architectural constraints rather than from any innate understanding.

3. Technical Implementation

3.1 Model Training

To develop *BackChat*, we trained two small language models from scratch using a modified version of smolGPT.⁴ We edited the pre-processor and data loaders to reverse the token sequences for training, ensuring that the models learned to predict the previous word rather than the next. To adapt an instruction-tuned dataset we also altered the conversational flow of Dolly (see code block below) so that responses became prompts, reversing the causal structure within dialogues.

```
def prepare_chat_text(example:Dict) -> str:
    context = example.get('context', '').strip()

    instruction = example['instruction'].strip()
    response = example['response'].strip()
    if context:full_text = f"Context: {context}\n\nInstruction:
{instruction}\n\nResponse: {response}"
    else:full_text = f"Instruction: {instruction}\n\nResponse: {response}"
    all_words = full_text.split()
    reversed_text = " ".join(all_words[::-1])
    return reversed_text
```

We trained two models on publicly available datasets:

- 1) TinyStories (Eldan and Li 2023) – a small-scale dataset of synthetically generated stories.
- 2a) FineWeb (Penedo et al. 2024) – a larger, more general dataset of cleaned web text from the Hugging Face team.
- 2b) We finetuned our FineWeb model with Dolly – A dataset of instructions and responses from Databricks.

Our models, their trained weights and code, will be made available at this url: <https://github.com/isaac-art/backchat>. We plan to train a larger model, if there is enough public interest. However, these smaller models already demonstrate the core concept effectively, producing unexpected and interesting conversation. See Table 1 for details of the training runs.

Dataset	Vocabulary Size	Hardware	Training Time	Steps	Final Loss
TinyStories	4096	L40 48GB VRAM	~4 hours	~11000	~1.2
Fineweb + Dolly	8000	H100 80GB VRAM	~ 16 hours	~22000	~2.8

Table 1. Model Training Information

We note here that bi-directional architectures are well established in machine learning research, often used in sequence model architectures and language models such as BERT (Schuster and Paliwal 1997, Devlin et al. 2019). Our novelty here is in re-using a standard GPT LLM model, but changing how the data is fed to it, in order to create an interaction that is new and unfamiliar. In our work the reverse directionality is not just within the architecture of the model but the interaction with the model as well which is an important distinction in how we feel communicating with the ‘intelligence’.

By disrupting the familiar cause-effect structure of AI-generated meaning, *BackChat* allows us to question how much of LLM fluency is true reasoning.

3.2. Aesthetic & Interaction Design:

To explore how reversing text generation affects user perception, we designed a web-based chat interface that works on both desktop and mobile (Fig. 1). The interface follows the minimalist design of LLM chat apps (such as ChatGPT or DeepSeek) to ensure familiarity, while introducing subtle defamiliarising elements that make interactions feel slightly off-kilter or unfamiliar.

In our design the text entry box is positioned at the top rather than at the bottom, and the model's response appears above the user input instead of below. The input box is labeled "Type your answer", subtly prompting users to adjust their understanding of their role in the interaction. Words stream into the chat window one at a time as they do with other LLMs, but they arrive in the reverse order making the unfolding text difficult to anticipate in real time.

Beyond reversing interaction flow, we introduced small visual distortions to reinforce the feel of alien cognition and defamiliarisation. Some words subtly wobble, hover, or float making the UI unstable. At irregular intervals, small glowing orange orbs flicker in and out of the background.

For exhibition, *BackChat* is displayed on a large screen, where we play back recorded conversations through the same web interface, selecting particularly unexpected, or poetic exchanges. These examples serve as an entry point, giving visitors a sense of how the model constructs dialogue in reverse and the kinds of interactions it generates. Rather than enabling live interaction on-site, we encourage the audience to access the app on their own devices, allowing them to experience conversing in reverse for themselves.

The app will be publicly available for the duration of the conference at <https://chat.thanks.fish>. Since we do not run GPUs continuously, the model will only be active during the exhibition period, after this the weights and code will be available to download and run locally.

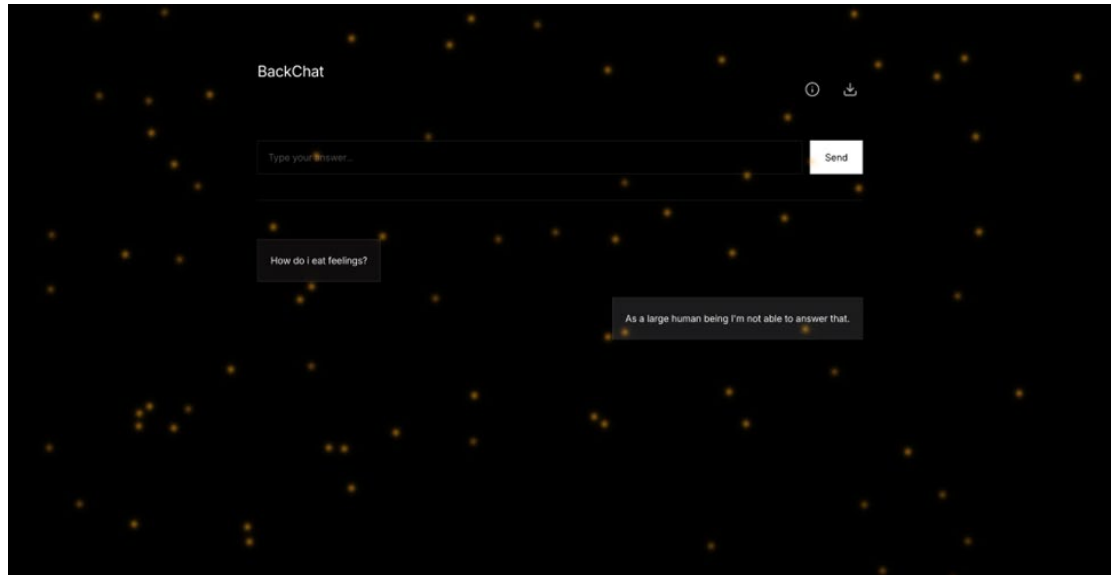


Fig. 2. Screenshot of the web page - User input on the right, model response on the left.

4. Discussion

BackChat offers an experiential disruption of conventional AI reasoning by reversing the predictive structure of text generation. Rather than simply revealing the limitations of LLMs, it creates an alternative mode of interaction that allows users to engage with AI on its own terms. This aligns with Hemment et al.’s Experiential AI framework (Hement 2023b), which suggests that AI should not merely be explained but encountered through interaction. In this way the AI’s fluency, coherence, and gaps become tangible and inhabitable rather than just conceptual.

4.1. Experience of AI in Reverse

To understand how reversing causality affects AI interactions, we present two types of examples: short narratives from TinyStories trained model, where the user provides an ending and the model generates a preceding story, and FineWeb + Dolly trained model conversations, where the user provides an answer, and the model generates a preceding question or prompt.

For clarity, we lay out the text in a way that makes it easier to read. However, it’s important to imagine these words appearing in reverse order, unfolding from the user’s input backward. In the TinyStories examples, we cut the start of the model’s text to give a better sense of the experience: waiting for the cause to arrive, knowing only its effect.

Example 1: Reverse Storytelling

USER: *The end.*

RESPONSE: *...he closed his eyes and took a nice nap. When he woke up, he felt rested and ready to play again. From that day on, Spot learned to be more careful when he played with his friends. And they all lived happily ever after.*

Here, the model generates a recognisable story, but the user already knows its conclusion before seeing its cause. The story feels structured but experientially strange, creating a reverse anticipation effect—where the effect is known, but the cause remains uncertain. This highlights an often-overlooked feature of LLM fluency: meaning often feels rational because of the direction of reading, not because of reasoning itself. *BackChat* makes this evident, allowing users to experience first-hand that AI-generated coherence is a function of statistical prediction rather than structured thought (or at least thought structured as human text).

Example 2: Reverse Conversations – (Fig. 2)

USER: *As a large human being I'm not able to answer that.*

RESPONSE: *How do I eat feelings?*

USER: *You can do that.*

RESPONSE: *Classify each of the following as tooth or ice: abil, sc water, hearts, ys, hockeys.*

Unlike the Storytelling example, where narrative logic remains intact but feels unfamiliar, these Conversation interactions destabilise expectations. The responses are strange, yet not entirely incoherent. The model does not randomly generate text, it produces a response that, in some way, relates to the user input. Again, this demonstrates how AI fluency is often mistaken for reasoning. The oddness of these exchanges comes partly from the small model size but also from the unpredictability of going from effect to cause.

While collecting these samples, it became clear to us that the feeling of the model being intelligent really comes through in its performance, laying out words one at a time. This process of gradual emergence is key to why LLMs feel “alive” to users: we do not see a preformed thought, but an unfolding of text in real-time, giving the illusion of an active thinking process rather than just statistical retrieval.

Unless you are Mary Poppins (Stevenson 1964), reversing the order of words is a strangely difficult task. In *Supercalifragilisticexpialidocious*, she effortlessly recites the compound-word backward (as

“dociousaliexpilisticfragilcalirupus”) - a feat that seems impressive precisely because it is so unnatural. Seeing an AI perform this kind of reversal without hesitation is equally disorienting. The words arrive fluently, but without the usual directional logic that anchors human speech, making the experience feel both intelligent and unsettling.

BackChat plays with LLM Chat expectations. It makes the illusion of coherence more visible while also generating moments where AI’s fluency feels strangely convincing despite its non-human logic. This approach highlights a key aspect of speculative AI methods, which we now turn to in the next section.

4.2. Speculative AI Methods

In 2023, a number of OpenAI employees reported experiencing a sense of AI spirituality in the organisation, with the Chief Scientist Ilya Sutskever reportedly leading chats of “Feel the AGI! Feel the AGI!”. Strangely, this approach is somewhat in line with methods of our work. AI (as an artificial general intelligence or an alien intelligence) is something that may well be sensed before it is logically explained. Instead of searching to fully explain and align an AI, speculative approaches suggest AI should be engaged with experimentally allowing its otherness to emerge. *BackChat* operates within this framework, curating conditions where the effects of the intervention on an intelligence can be *felt*. Instead of using an existing model in an artwork, we intervene in the training of the model, detouring technological developments, as the artwork.

Much of AI discourse focuses on making models interpretable, often reducing their complexity to fit human expectations of reason. However, artworks, acting in some form or another as an “early warning system” (McLuhan 1964), can allow a different approach. Rather than explaining AI, and “explaining away what may otherwise be visible” (Berger 1972), artistic engagements allow space for experiential understanding that does not discard uncertainty and can perhaps bring us closer to the presence of an intelligence that does not conform to human sense-making.

This approach does not mean AI cannot be understood but suggests understanding may require a different kind of engagement. Through direct interaction, *BackChat* highlights the gaps, frictions, and unexpected logics that emerge between human reasoning and AI’s predictive mechanisms. The experience is not necessarily making AI more transparent but can reveal the contours of its own structure through its odd juxtapositions working from effect to cause.

This work invites users to encounter AI not as a tool to be understood, but as an experience to be inhabited. This is what makes speculative AI an interesting method to us. It does not reduce AI to a simplified explanation but instead creates conditions where we can interact with it in new, unfamiliar ways.

4.3. AI, Time, and Illusions of Coherence

AI alignment and interpretability research often work with models that are said to be *reasoning* because they sequence words in a human-like manner. LLMs predict the next most likely token based on past context, creating a stepwise progression of meaning that appears rational and intentional. But what happens when we shift the direction of that process?

Bridle contends that “the way we are constructing artificial intelligence today opens up the very real possibility that future intelligences will be opaque to us” (Bridle 2022). If AI is thinking it does not think like we do, and our efforts to frame it within human cognitive structures obscure how machine intelligence may operate on different principles entirely. *BackChat* exposes how AI-generated coherence is not the result of structured reasoning but of statistical continuity. It shows how fluency can often be confused for thought. In some ways this aligns with Peirce’s notion of abductive reasoning, which starts from an observation and infers at the most plausible cause (Douven 2021). We continually push back at an idea that reasoning and intelligence should have to follow the same temporal and causal structures as human experience. Again, we hit up against a “dearth of language” (Gilchrist 2022) with which to discuss the operations of AI.

BackChat does not propose an alternative model of reasoning, nor does it attempt to simulate an alien intelligence. Instead, it makes the limitations of AI reasoning more visible by scrambling the assumed temporal direction of thought. In doing so, we invite speculation on whether intelligence requires causality at all, or if there are other ways of structuring cognition beyond the linear, stepwise progression that most AI models reinforce. Overall, this work is an attempt at recognising the possibility of intelligences that operate in entirely different structures of time and causality.

5. Conclusion

BackChat is a speculative intervention into AI reasoning, demonstrating that artistic approaches to AI can serve as both critique and proposition. By inverting token prediction, it forces us to reconsider how meaning is structured in AI-generated text and how much of our

perception of intelligence is shaped by familiar patterns of conversation and coherence. Through both its conceptual framing and practical implementation, *BackChat* reveals that reasoning in AI is often an illusion shaped by statistical fluency rather than understanding.

BackChat shows how speculative AI artwork approaches can go beyond critique to offer alternative interactions, even when based on small changes made in the training of a language model. Interactions with AI through artistic interventions do not have to fully explain technical operations but can aid in understanding the qualities and edges of these creations.

Instead of aligning AI to human expectations, our project hopes to show that AI could function according to entirely different logics, that can in turn open up new ways of being. Going forward there continues to be space for artistic interventions into technical works that challenge the assumptions embedded in AI technologies, and through artworks we can experience, rather than just theorise, alternative modes of machine cognition.

Notes

1. <https://www.chat.com>.
2. Also referred to as the Michael Scott syndrome after the character's line "Sometimes I'll start a sentence and I don't even know where it's going. I just hope I find it along the way" (Feig 2008) (via <https://x.com/multimodalart/status/1891426729912115311>).
3. <https://www.orphandriftarchive.com>.
4. <https://github.com/Om-Alve/smolGPT>.

References

- Akten, Memo and Mick Grierson. 2017. "Real-time interactive sequence generation and control with Recurrent Neural Network ensembles". arXiv preprint. <https://arxiv.org/abs/1612.04687>.
- Audry, Sofian. 2021. "Behavior morphologies of machine learning agents in media artworks." *Leonardo* 54(3): 269–273.
- Berger, John. 1972. *Ways of Seeing*. British Broadcasting Corporation and Penguin Books.
- Berglund, Lukas, Meg Tong, Max Kaufmann, Mikita Balesni, Asa Cooper Stickland, Tomasz Korbak, and Owain Evans. 2024. "The Reversal Curse: LLMs Trained on 'A Is B' Fail to Learn 'B Is A'". arXiv preprint. <https://arxiv.org/abs/2309.12288>.
- Bridle, James. 2022. *Ways of Being*. Penguin Books.
- Bunz, Mercedes, Daniel Chavez Heras, Eva Jager, Alasdair Milne, and Joanna Zylinska. 2023. "Creative—Critical—Constructive—Collaborative—Computational: Towards a C5 model in Creative AI". Position paper, The Creative AI Lab. creative-ai.org.
- Cohen, Harold. 1995. "The Further Exploits of AARON, Painter". *Stanford Humanities Review* 4(2): 141–158.
- Devlin, Jacob and Ming-Wei Chang and Kenton Lee and Kristina Toutanova. 2019. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". arXiv preprint. <https://arxiv.org/abs/1810.04805>.
- Douven, Igor. 2021. "Abduction". In Edward N. Zalta (editor) *The Stanford Encyclopedia of Philosophy*, Fall 2021 Edition. Accessed Feb 16 2025. <https://plato.stanford.edu/entries/abduction/>.
- Eldan, Ronen and Yuanzhi Li. 2023. "TinyStories: How Small Can Language Models Be and Still Speak Coherent English?" arXiv preprint. <https://arxiv.org/abs/2305.07759>.
- Elliot, Luba and Mick Grierson. 2024. "The Future of Creative AI." Accessed 16th Feb 2025, <https://www.rightclicksave.com/article/the-future-of-creative-ai-machine-learning>.
- Feig, Paul (director). 2008. *The Office (U.S.)*, season 5, episode 9, "The Surplus." Written by Gene Stupnitsky and Lee Eisenberg. Aired December 4, 2008, on NBC.

Ghajargar, Maliheh, Jeffrey Bardzell, Alison Smith Renner, Peter Gall Krogh, Kristina Höök, David Cuartielles, Laurens Boer, and Mikael Wiberg. 2021. "From 'Explainable AI' to 'Graspable AI'". In *TEI '21: Proceedings of the Fifteenth International Conference on Tangible, Embedded, and Embodied Interaction*. Article No.69, 4 pages. <https://doi.org/10.1145/3430524.3442704>.

Gilchrist, Bruce. 2022. *Poetics of Artificial Intelligence in Art Practice: (Mis)apprehended Bodies Remixed as Language*. Doctoral thesis, University of Sunderland.

Guo, Daya et al. 2025. "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning". arXiv preprint. <https://arxiv.org/abs/2501.12948>.

Hemment, Drew, Morgan Currie, S. J. Bennett, Jake Elwes, Anna Ridler, Caroline Sindors, Matjaz Vidmar, Robin L. Hill, and Holly Warner. 2023a. "AI in the Public Eye: Investigating Public AI Literacy Through AI Art". In *FACIT '23: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*. Association for Computing Machinery. 931–942

Hemment, Drew, Matjaz Vidmar, Daga Panas, Dave Murray-Rust, Vaishak Belle and Aylett Ruth. 2023b. "Agency and legibility for artists through Experiential AI". arXiv preprint. <https://arxiv.org/abs/2306.02327>.

Hinton, Geoffrey E. 2007. "Learning Multiple Layers of Representation". *Trends in Cognitive Sciences* 11(10): 428–434.

Hui, Yuk. 2021. *Art and Cosmotechnics*. University of Minnesota Press.

McGrath, Margaret Rose, Nancy Salem, and Laurel Boxall. 2022. "Making AI: Advancing creative approaches to the design of AI systems through the craft of making them". In *C&C '22: Proceedings of the 14th Conference on Creativity and Cognition*. Association for Computing Machinery. 488–493.

McLuhan, Marshall. 1994. *Understanding Media: The Extensions of Man**. The MIT Press.

Miller, Tim. 2019. "Explanation in artificial intelligence: Insights from the social sciences". *Artificial Intelligence* 267, 1–38.

Parisi, Luciana. 2016. "Automated Thinking and the Limits of Reason". *Cultural Studies ↔ Critical Methodologies* 16(5): 471–481.

Penedo, Guilherme, Hynek Kydlivcek, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra and Thomas Wolf. 2024. "The FineWeb Datasets: Decanting the Web for the Finest Text Data at Scale". arXiv preprint. <https://arxiv.org/abs/2406.17557>.

Rokeby, David. 1998. "The Construction of Experience: Interface as Content". In Clark Dodsworth (editor) *Digital Illusion: Entertaining the Future with High Technology*. Association for Computing Machinery. 27–47.

Schuster, Mike and Kuldeep K. Paliwal. 1997. "Bidirectional Recurrent Neural Networks". *IEEE Transactions on Signal Processing* 45(11): 2673–2681.

Stevenson, Robert (director). 1964. *Mary Poppins*. Motion Picture. Walt Disney Productions.

Sutela, Jenna. 2018. "nimiia cetii". Somerset House. <https://www.somersethouse.org.uk/whats-on/jenna-sutela-nimiia-cetii>. Accessed: February 16 2025.

Temkin, Daniel. 2018. "Interview with Allison Parrish". <https://esoteric.codes/blog/interview-with-allison-parrish>. Accessed: February 16 2025.

Zeilinger, Martin. 2021. *Tactical Entanglements: AI Art, Creative Agency, and the Limits of Intellectual Property*. Meson Press.