

Editorial: What Have Large-Language Models and Generative AI Got to Do With Artificial Life?

Alan Dorin

Monash University
Computational and Collective
Intelligence Group
Department of Data Science and AI
Faculty of Information Technology
alan.dorin@monash.edu

Susan Stepney

University of York
Department of Computer Science
susan.stepney@york.ac.uk

Accessible generative artificial intelligence (AI) tools like large-language models (LLMs) (e.g., ChatGPT,¹ Minerva²) are raising a flurry of questions about the potential and implications of generative algorithms and the ethical use of AI-generated text in a variety of contexts, including open science (Bugbee & Ramachandran, 2023), student assessment (Heidt, 2023), and medicine (Harrer, 2023). Similarly, among the graphic and visual arts communities, the use of generative image synthesis algorithms (e.g., DALL-E,³ Midjourney,⁴ Stable Diffusion⁵) that take text prompts as input and produce works in the style of a particular human artist, or no artist who ever lived, are causing consternation and posing challenging questions (Murphy, 2022; Plunkett, 2022). The use of generative AI to create deep fakes has also been in the spotlight (Ruiter, 2021), as has its role in answering scientific research questions directly (Castelvecchi, 2023).

To our minds, the questions these technologies are raising do not seem to be of a fundamentally different character to questions asked about AI for many years. They largely concern (a) what is possible, (b) what is right, and (c) the implications of the technology's use. For instance,

1. Can AI generate documentary “evidence” that is indistinguishable from reality? Can AI generate artifacts that are competitive with (or superior to) those made by a human?
2. How is the concept of “truth” confused or undermined by the output of these technologies? Is it ethical to load examples of a human's art to generate a model for replicating their style?
3. Who owns the intellectual property rights of AI-generated artifacts? Who is accountable if an AI-generated artifact causes harm? Will I lose my job to an AI; does AI make creative writers or artists redundant?

The sense of urgency and public interest in these questions has increased as generative models become practical and ubiquitous tools with levels of skill or expertise that appear to be of above average human competency.

1 <https://openai.com/blog/chatgpt>.

2 <https://ai.googleblog.com/2022/06/minerva-solving-quantitative-reasoning.html>.

3 <https://openai.com/product/dall-e-2>.

4 <https://www.midjourney.com>.

5 <https://stablediffusionweb.com>.

Of general relevance to academic publishers, including *Artificial Life*, is that LLMs can compose text structured in the form of research paper sections or even whole articles. And these texts are often (currently) hard to distinguish from human-written text, even by software trained to detect text generated by language models (LMs). When questioned (on February 20, 2023) “What should a journal editor do about LMs?” You.com’s chatbot suggested,

A journal editor should ensure that all submitted manuscripts adhere to the journal’s style guide, and that the language used is appropriate for the journal’s audience. They should also ensure that any language models (LMs) used to assist with the writing process are properly cited and that any generated text is labeled as such. Additionally, the journal editor should ensure that the LMs are used in accordance with the journal’s policies and ethical considerations.

That’s a generic response that begins tangentially to the questioner’s intended focus. As any human reader can immediately see, the chatbot is not aware of the social context of the question, nor has it had the opportunity to gauge the questioner’s level of expertise. But the text is “good enough”—it does address relevant issues, and it’s an answer that probably took you a few seconds to read, interpret, and critique. It may take you more time to read and analyze the text and this paragraph than it took for us and the chatbot to generate it. Was this therefore a waste of your time, or ours? Is the chatbot wasting your time? Are we? Some journals and publishers have drafted formal policies that require the use of LLMs for writing submissions to be explicitly acknowledged (e.g., at Springer-Nature; “Tools Such as ChatGPT,” 2023). In practice, their use (or misuse) may be very difficult to detect.

Artificial Life, and its publisher MIT Press generally, is also adopting the policy that any use of generative AI, for any part of a submitted work, including but not limited to text, images, sound, data, mathematics, logic, reasoning, programming code, or algorithms, must be prominently, explicitly, and unambiguously labeled and its source formally cited (e.g., via a name, manufacturer, URL, version number, or access date).

Journals and publishers have also moved to prevent LLMs from being listed as authors on articles. For instance, Springer-Nature’s policy was online earlier this year, and, although it has now seemingly been removed from its original location, variants of it have been incorporated into the authorship policies of some journals:

Large Language Models (LLMs), such as ChatGPT, do not currently satisfy our authorship criteria. Notably an attribution of authorship carries with it accountability for the work, which cannot be effectively applied to LLMs. Use of an LLM should be properly documented in the Methods section (and if a Methods section is not available, in a suitable alternative part) of the manuscript. (Nature, 2023)

In a policy that remains online at time of writing, the journal *Science* states,

Text generated from AI, machine learning, or similar algorithmic tools cannot be used in papers published in *Science* journals, nor can the accompanying figures, images, or graphics be the products of such tools, without explicit permission from the editors. In addition, an AI program cannot be an author of a *Science* journal paper. A violation of this policy constitutes scientific misconduct. (Science, 2023)

Artificial Life and MIT Press are taking an approach in alignment with those of the editorial boards (and publishing house legal teams) of such journals: that authorship is associated with responsibility and accountability for an article. However, for *Artificial Life*, the issue doesn’t stop there.

The implications of generative AI are relevant to a broad spectrum of society. But an interest in generative computational processes is arguably at the center of Artificial Life research. How might LLMs be specifically relevant to Artificial Life, as opposed to the subdiscipline (yes, that's ironic) of AI? Here are a few ideas.

The production of novelty can be explored through the use of an LM that continually takes as its input text composed by humans, other LMs, and its own output. Such a system is relevant to our field's interests in feedback loops, open-endedness, and the emergence of complexity. Is this system engaged in language acquisition through "social" interactions?

LMs might be used to explore questions related to the emergence of meaning in language. Can meaning be generated by an LM, or is it specific to living things? Can LMs evolve to be better interpreters and writers? How does the text LMs generate change the way humans produce and use language?

If any work responding to these questions was presented in the form of a formal research experiment documented in an article, then this would naturally fall within the scope of human-authored research. However, an LLM-generated poem, song, or essay can be of value to researchers in Artificial Life exploring these topics (even if it isn't very good; Cave, 2023). The coauthorship of such a work by an LM and a human as a way of communicating ideas about Artificial Life would be interesting to consider. In this case, the text contributed by the LM would need to be quoted as an "example" within the text of a submission made by a human author who determined that it was worthy of submission. Even though the text itself is a direct, self-referential, and, we would hope, revealing exploration of an LM system's quirks, capabilities, or limitations, the determination of its relevance must, for the time being at least, be made by and attributed to a human.

There is a precedent for such work, much of which has been explored under the banners of cybernetic, generative, and Artificial Life art (e.g., see many historical examples in Benthall, 1972; Ohlenschläger, 2012; Reichardt, 1968; Whitelaw, 2004). In such contexts, the art is published. Associated commentary and/or explanations may come later, and these might not be authored by the same system or person who made the original work.

An interview, debate, discussion, or duet between a human and an AI, or between several computer programs, can also challenge our ideas about living systems and their exchange of information, the use of language, or the production of improvised movement and sound. The inclusion of extracts from discussions with computer chatbots dates back at least to the advent of ELIZA in the 1960s: "Men are all alike—IN WHAT WAY—They're always bugging us about something or other" (Weizenbaum, 1966, p. 36). Likewise collaborative improvisations performed by robots, algorithms, and humans have an established place in music (Bown, 2011; Eldridge, 2005). As far as we know, the journal hasn't published such works previously. But we could.

For something along these lines to be published today, as with any contribution, it would of course need to provide novel perspective or insight. However, the main point here is that in these scenarios, even though we might intuitively feel that the generative AI system warrants the status of contributor at the level of coauthor, we have to insist on a human author of the submission. They would have ultimate responsibility for the work produced by the generative algorithm so that, for instance, if the LLM's poetry influenced thought in a positive and productive way, or if it incited violence, we would have somebody accountable to thank or blame. If the article's publication required the payment of an open access fee, we would also have somebody from whom to extract the payment!

We haven't yet received any submissions made by generative AI (as far as we know). But this issue contains novel work by human authors. In fact, we have recently published a spate of varied special issues reporting on the research presented at human gatherings, some in person, some online. These have covered a wide range of exciting activity in the Artificial Life community: Issue 28:2 has extended versions of selected papers from the 2019 Artificial Life conference; 28:3 is a collection of articles on embodied intelligence; 28:4 is the Artificial Life 2021 conference special issue; 29:1 explores agent-based models of human behavior. We extend our thanks to all the guest editors for their hard work in handling the selection and review of articles for their issues. New ideas for

special issues in the subdomains of Artificial Life are always welcome. If you have an idea, please contact us.

After that run of special issues, we welcome you back to a general issue of contributed research articles that cover a wide range of Artificial Life topics—including distributed control, emergence, dynamical systems, self-organization, game theory, artificial chemistry, and biocomputing—addressed through theory, models, simulations, and physical experiments.

We start with a letter from Bull and Liu, on “A Generalised Dropout Mechanism for Distributed Systems.” They use a modified NK model to sharpen the criteria for determining when local control is more beneficial than global control. Next, we have an article from Gershenson, on “Emergence in Artificial Life.” He uses the difference in information present at different levels of a system as the basis for a new definition of emergence, one of the fundamental components of ALife.

The article from Howison et al., “On the Stability and Behavioral Diversity of Single and Collective Bernoulli Balls,” describes a platform for investigating how dynamical systems may be used as the basis for designing a variety of agent behaviors. This platform, both in simulation and as a physical system, comprises a collection of “Bernoulli balls” in an airflow, interacting with each other and with the flow. The aim is to develop a dynamical system with a diverse set of possible behaviors.

Ichinose et al. present “How Lévy Flights Triggered by the Presence of Defectors Affect Evolution of Cooperation in Spatial Games.” Lévy flights model a kind of random motion with both small and big displacements. Here Lévy flights are combined with game theory concepts in an agent-based model. The authors investigate how the presence of defectors changes the optimal behaviors.

Next, Scott and Pitt investigate “Interdependent Self-Organizing Mechanisms for Cooperative Survival.” Complex survival games, where cooperation is needed to survive intermittent catastrophes, need complex strategies. Here the authors look at social self-organization, which, as any complex domain, has aspects that can make the situation better in some cases and, in other cases, worse. They conclude that such systems need to be able to reflect on their own operation through some kind of self-model.

Sienkiewicz and Jedruch tell us about “DigiHive: Artificial Chemistry Environment for Modeling of Self-Organization Phenomena.” This two-dimensional continuous space simulation environment supports experiments with the goal of facilitating open-ended simulations. It steers more toward natural physical and biological systems (e.g., it includes energy conservation), rather than toward the more abstract operation of some other artificial chemistries. The authors describe the rationale and operation of the system and use it to investigate aspects of self-organization and self-replication in cellular-like systems.

Finally, Svahn and Prokopenko examine “An Ansatz for Computational Undecidability in RNA Automata.” An *ansatz* is an “educated guess” about the form of the solution to a problem that can be used to provide a stepping-stone to finding the solution. Here the approach uses the known computational power of a set of automaton models as the form of solution and shows how RNA behaviors map to these models to demonstrate the computational power of this biological form of computing.

References

- Benthall, J. (1972). *Science and technology in art today*. Thames and Hudson.
- Bown, O. (2011). Experiments in modular design for the creative composition of live algorithms. *Computer Music Journal*, 35(3), 73–85. https://doi.org/10.1162/COMJ_a_00070
- Bugbee, K., & Ramachandran, R. (2023). *The ethics of large language models: Who controls the future of open science?* <https://impactunofficial.medium.com/the-ethics-of-large-language-models-who-controls-the-future-of-open-science-43cca235401d>
- Castelvecchi, D. (2023). How will AI change mathematics? Rise of chatbots highlights discussion. *Nature*, 615, 15–16. <https://doi.org/10.1038/d41586-023-00487-2>, PubMed: 36808415

- Cave, N. (2023, January). I asked Chat GPT to write a song in the style of Nick Cave and this is what it produced. What do you think? *Red Hand Files*, no. 218. <https://www.theredhandfiles.com/chat-gpt-what-do-you-think/>
- Eldridge, A. (2005). Cyborg dancing: Generative systems for man-machine musical improvisation. In T. C. Innocent (Ed.), *Third iteration: Third international conference on generative systems in the electronic arts* (pp. 129–141). Monash University Publishing.
- Harrer, S. (2023). Attention is not all you need: The complicated case of ethically using large language models in healthcare and medicine. *eBioMedicine*, 90, 104512. <https://doi.org/10.1016/j.ebiom.2023.104512>, PubMed: 36924620
- Heidt, A. (2023). Arms race with automation [Technology feature]. *Nature*. <https://doi.org/10.1038/d41586-023-00204-z>, PubMed: 36693972
- Murphy, B. P. (2022). No, the Lensa AI app technically isn't stealing artists' work—but it will majorly shake up the art world. *The Conversation*. <https://theconversation.com/no-the-lensa-ai-app-technically-isnt-stealing-artists-work-but-it-will-majorly-shake-up-the-art-world-196480>
- Nature. (2023). Authorship. *Nature Portfolio*. <https://www.nature.com/nature-portfolio/editorial-policies/authorship>
- Ohlenschläger, K. (Ed.). (2012). *VIDA art and artificial life 1999–2012*. Fundación Telefónica.
- Plunkett, L. (2022). AI creating “art” is an ethical and copyright nightmare. *Kotaku*. <https://kotaku.com/ai-art-dall-e-midjourney-stable-diffusion-copyright-1849388060>
- Reichardt, J. (Ed.). (1968). *Cybernetic serendipity: The computer and the arts*. Studio International.
- Ruiter, A. D. (2021). The distinct wrong of deepfakes. *Philosophy and Technology*, 34, 1311–1332. <https://doi.org/10.1007/s13347-021-00459-2>
- Science. (2023). Authorship. *Science*. <https://www.science.org/content/page/science-journals-editorial-policies#authorship>
- Tools such as ChatGPT threaten transparent science; here are our ground rules for their use [Editorial]. (2023). *Nature*, 613, 612. <https://doi.org/10.1038/d41586-023-00191-1>, PubMed: 36694020
- Weizenbaum, J. (1966). ELIZA—A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45. <https://doi.org/10.1145/365153.365168>
- Whitelaw, M. (2004). *Metacreation: Art and artificial life*. MIT Press.