

Editorial: Generative artificial intelligence and the ecology of human development

There is no escaping the news about Large Generative AI Models (LGAIMs). Commercial applications of artificial intelligence in the form of Large Language Models (LLMs) and Generative AI (e.g., chatGPT; Dall-E) have taken centre stage in the media sphere, business, public policy, and education, quickly gaining traction outside the original scholarly field of computer science. Leading developmental researchers such as Frank (2023) and Gopnik (2023) have written on prominent academic platforms about how the science of child development may contribute to the understanding and even the education of LGAIMs. Furthermore, with the hybridization of online and offline social interactions, applications of generative AI are rapidly becoming part of young people's private sphere – transitioning us from the era defined by apps optimised to gain people's attention to a new era that will focus on apps' ability to form direct relationships with their users. This is all unfolding at exponential rates, so much so that scholars, opinion makers, and policymakers are sounding the alarm and debating a moratorium to allow research, reflection, and regulation (Clarke, 2023).

The widespread introduction of generative AI is taking place at the heels of the COVID-19 pandemic, which was treated as a global emergency by the World Health Organization, leading to drastic public health measures and a redirection of scientific research programs on a massive scale (Ioannidis, Bendavid, Salholz-Hillel, Boyack, & Baas, 2022). Families, practitioners, and policymakers faced many urgent questions about this novel threat and scientists responded beyond the obvious fields of immunology, infection prevention, and public health, aided by the quick reallocation of research funding and expedited translation into practice. These events have led to scientific breakthroughs in some fields (e.g., mRNA; infection control), while the impact in other fields remains to be seen.

Now again, the proliferation of LGAIMs and the prospect of artificial general intelligence generate many new questions and possibilities for breakthroughs, spurring governments to announce large investments in research and development and computing infrastructure. Once again, we may ask whether interests of youth and families are best served by trying not to become too distracted by this

shiny new toy or whether in the public interest, we should engage early and comprehensively with new powerful tools.

Developmental psychological research on generative AI

Developmental research conducted on LGAIMs has captured the imagination of many, with preprints of articles being widely circulated and debated. For example, Kosinski posted in February 2023 a preprint article reporting on a series of experiments suggesting that LGAIMs perform better on false beliefs tasks if these are trained on more language data (i.e., successive versions of GPT-3 and GPT-4), which he interpreted as demonstrating how Theory of Mind is an ability that can emerge from learning to put one word after the other. Another study quickly followed that showed that the performance of LGAIMs on false belief tasks was not robust against variations in vignettes and prompts, supporting a more sceptic stance towards the emergence of higher cognitive functions in these models (Ullman, 2023). The issue is far from settled, however, and researchers may have started to consider how LLMs perform on other psychological tests and assessments and to ponder the implications of their initial observations of LGAIMs' performance.

The current issue includes fascinating empirical work on aggressive behaviour in response to simulated rejection (Quarmley, Vafiadis, & Jarcho, 2023) and on parents' ability to see the world from the point of view of their child with autism (Oppenheim et al., 2023). Both studies relate to the models that youth and adults construct of their social world and their place in it and both studies grapple with the problem that those models are opaque to social partners and can only be indirectly studied by researchers. The allure of research on LGAIMs is that, in principle, the models that underly the complex social-cognitive abilities ascribed to them may be directly described and probed (Frank, 2023). However, such work becomes only relevant for students of human behaviour and development outside linguistics, such as the readers of this journal, if LGAIMs can be demonstrated to show functional competencies (pairing linguistic knowledge with formal reasoning,

We may ask whether interests of youth and families are best served by trying not to become too distracted by this shiny new toy or whether in the public interest, we should engage early and comprehensively with new powerful tools.

Conflict of interest statement: See Acknowledgements for full disclosures.

world knowledge, situational modelling, and communicative intent) over merely formal competence (as a product of linguistic knowledge only; Mahowald et al., 2023). Still, it is likely that someone somewhere is working to overcome the considerable challenges in combining technologies so that LGAIMs attain functional competencies.

Additional problems may need to be tackled as well if generative AI is to catalyse our science in ways akin to the automation of gene sequencing and the arrival of affordable neuroimaging apparatus. For one, science has learned the hard way that reproducibility and replicability are key aspects of scientific rigour. This requires that either the commercial operators of LGAIMs implement and facilitate open science practices, such as persistent identifiers to the versions of their models and continued access, or that research and investment should focus on open-source LLMs (Spirling, 2023). Another set of problems relates to novel regulatory, legal, and ethical dilemmas that need to be resolved across the research value chain from institutes, to funders, to publication outlets.

Generative AI in a hybrid social world

It may provide some comfort to parents, professionals, historians, and policymakers that developmental and cognitive research has not yet found definitive evidence of psychological properties such as thought and intent and that it is an open question whether such properties may ever emerge. However, commercial LGAIMs are capable of advanced linguistic tasks such as holding a conversation, summarising text, and extracting information from data. Interfaces to LGAIMs have been built in the form of chatbots that allow anyone, including children, to interact with AI. Some of the ramifications are not unique to generative AI and may be gauged based on what we are learning for example, about the effects of displacement of offline activities and social interactions to online ones, about the impact of disinformation, or about the effectiveness of chatbot interventions. Other ramifications might be unique, however. As soon as they have learned to talk, children can engage with artificial conversational agents, such as those that are currently on the market as smart speakers and trust these relatively more than humans for factual information (Xu, 2023). With generative AI, the range of topics for conversation becomes as broad as the internet. Furthermore, the sophisticated and adaptive conversational skills of applications that use LLMs may reduce the advantage that human social partners still enjoy on the trustworthiness of personal information.

Generative AI in the real world


Having taught machines to generate language indistinguishable from our own is a profound moment in

human history. While the evidence is inconclusive, with indications that overfitting may be responsible for some of the remarkable abilities reported as emerging from LGAIMs, the real-world implication remains historic. In 2022, one of Google's engineers tasked with testing their LaMDA large language models was fired after claiming it to be sentient. What is important here is not whether the model was conscious or not, but that a person was willing to give up their job in defence of his beliefs about the model. So much so, people went as far as to question his mental health. This is likely the first of many such stories to come.

Conclusions

Understanding the impact of the rapidly digitalising ecology of human development and testing ways to leverage technological advances to support families and young people are increasingly important themes in our field and journal. The current issue, for example, features the trial by Werner-Seidler and colleagues (2023), showing how digital technology may not only disrupt sleep but can also be used to reduce insomnia. It may be a little too early to also turn to generative AI as a source of novel models of human mental health and functioning. However, people use language to create myth and legend, to create art and science, and to create friendships and loving relationships. Now that LGAIMs can effortlessly produce such language if prompted well, new opportunities for redressing the social order may present themselves, for better or worse.

Carlo Schuengel^{1,2} 

Alastair van Heerden^{3,4} 

¹Educational and Family Studies, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands;

²Amsterdam Public Health Research Institute, Amsterdam UMC and Vrije Universiteit Amsterdam, Amsterdam, The Netherlands; ³Center for Community Based Research, Human Sciences Research Council, Pietermaritzburg, South Africa and ⁴SAMRC/WITS Developmental Pathways for Health Research Unit, Department of Paediatrics, School of Clinical Medicine, Faculty of Health Sciences, University of the Witwatersrand, Johannesburg, South Africa
E-mail: c.schuengel@vu.nl

Acknowledgements

C.S. is joint editor of the *Journal of Child Psychology and Psychiatry* for which he receives an honorarium. His research is supported by grants from The Netherlands Organisation for Health Research and Development (ZonMw) and from the Scientific Research Foundation (SWO) 's Heeren Loo. A.H. is supported in his research by grants from the National Institutes of Health (NIH), the Bill and Melinda Gates Foundation (BMGF), and the European and Developing Countries Clinical Trials Partnership (EDCTP).

References

- Clarke, L. (2023). Call for AI pause highlights potential dangers. *Science*, 380, 120–121.
- Frank, M.C. (2023). Baby steps in evaluating the capacities of large language models. *Nature Reviews Psychology*, 1–2. doi: 10.1038/s44159-023-00211-x
- Gopnik, A. (2023). A very human answer to one of AI's deepest dilemmas. *Observer*. Available from: <https://www.psychologicalscience.org/observer/ai-deepest-dilemmas> [last accessed 1 July 2023].
- Ioannidis, J.P.A., Bendavid, E., Salholz-Hillel, M., Boyack, K.W., & Baas, J. (2022). Massive covidization of research citations and the citation elite. *Proceedings of the National Academy of Sciences of the United States of America*, 119, e2204074119.
- Kosinski, M. (2023). Theory of mind may have spontaneously emerged in large language models. *ArXiv*, abs/2302.02083.
- Mahowald, K., Ivanova, A.A., Blank, I.A., Kanwisher, N.G., Tenenbaum, J.B., & Fedorenko, E. (2023). Dissociating language and thought in large language models: A cognitive perspective. *ArXiv*, abs/2301.06627.
- Oppenheim, D., Koren-Karie, N., Hamburger, L., Maccabi, Y., Slonim, M., & Yirmiya, N. (2023). Parental insightfulness is associated with mother–father–child interactions among families of preschoolers with an Autism Spectrum Disorder diagnosis. *Journal of Child Psychology and Psychiatry*, 64, 1359–1368.
- Quarmley, M., Vafiadis, A., & Jarcho, J.M. (2023). Irritability and rejection-elicited aggression in adolescents and young adults. *Journal of Child Psychology and Psychiatry*, 64, 1346–1358.
- Spirling, A. (2023). Why open-source generative AI models are an ethical way forward for science. *Nature*, 616, 413.
- Ullman, T.D. (2023). Large language models fail on trivial alterations to theory-of-mind tasks. *ArXiv*, abs/2302.08399.
- Werner-Seidler, A., Li, S. H., Spanos, S., Johnston, L., O'Dea, B., Torok, M., Ritterband, L., Newby, J. M., Mackinnon, A. J., & Christensen, H. (2023). The effects of a sleep-focused smartphone application on insomnia and depressive symptoms: A randomised controlled trial and mediation analysis. *Journal of Child Psychology and Psychiatry*, 64, 1324–1335.
- Xu, Y. (2023). Talking with machines: Can conversational technologies serve as children's social partners? *Child Development Perspectives*, 17, 53–58.

Accepted for publication: 4 July 2023