

A social network for AI



Further progress in AI may require learning algorithms to generate their own data rather than assimilate static datasets. A Perspective in this issue proposes that they could do so by interacting with other learning agents in a socially structured way.

One of the main drivers behind the most important advances in artificial intelligence (AI) over the past two decades has been the availability of an increasing amount of training data. Notably, the collection and annotation of 14 million images in the ImageNet database for visual object recognition in 2012 resulted in a step change in the capabilities of deep learning algorithms. More recently, large datasets of varying modality – text, images, audio and whatever can be found on the internet – are fed into large language models.

These AI developments have produced impressive mainstream advances, such as in computer vision for self-driving cars and chatbot assistants that mimic the generation of human-like text and answer questions or prompts. However, the continuous scaling of AI models over the past decade, with performances increasing with model size and the amount of training data, is unsustainable for several reasons. There are substantial concerns about the energy consumption of large AI models, as discussed in a [Comment](#) in this issue. Furthermore, harvesting all possible data from the internet to train increasingly powerful generative AI models is problematic owing to several issues including those related to ethical¹ and copyright² concerns, as well as to the fact that the world could run out of high-quality data to train even bigger models in a few years³.

To break from the current trend of AI models getting massively larger and only incrementally better while eating up the world's energy and data resources, fresh ideas are needed in fundamental AI research. For example, learning algorithms could be designed that focus on seeking out or even creating new data rather than assimilating static data. An essential feature of natural intelligence is the drive to acquire novel information through exploration, a process which can be simulated by reinforcement learning approaches. However, a challenge is to do this efficiently and to circumvent a trade-off between exploration (gathering information) and exploitation (using previously gathered information to achieve a goal). Resolving this trade-off is a [challenge encountered by most organisms](#).

In a [Perspective](#) in this issue, Duéñez-Guzmán et al. present an intriguing vision for ongoing data generation by AI algorithms. They propose that such data generation arises when AI agents interact in a structured, social way as they compete and cooperate with each other, mimicking phenomena in biological and human evolution in which social interactions in various forms and scales are essential for societal transitions, emergent behaviour and innovations. A key concept that arises from such interactions, as highlighted by the authors, is 'compounding innovation', or innovation building on previous innovation, which arises as an environment continually changes due to specific types of social interactions between biological entities. The authors explain that this avoids the usual trade-off between exploration and exploitation, as exploitation continuously generates new data for learning opportunities.

The authors describe three forms of social structure in biological systems that drive compounding innovation, with unique

interactions in each form leading to changes in the data stream. The first form of social structure is collective living, in which the interactions of agents are anonymous and mediated by proximity, such as swarm intelligence. The second form is social relationships, in which the identity of individuals and their relationships matter during interactions, creating networks that facilitate cooperation and social learning. The third form is major transitions in evolution, in which larger-scale agents regulate the environments of smaller-scale ones.

Another, related approach to data generation with learning algorithms is to give autonomous agents an intrinsic motivation towards exploration and developing their own goals, with so-called autotelic agents. In a recent [Perspective](#), Colas et al.⁴ proposed to embed such autotelic, self-motivated AI agents in a socio-cultural environment where language is essential. The authors focussed on how the structure of language and cultural content can be leveraged by autotelic agents to turn socio-cultural interactions into internal cognitive tools, thereby becoming more human-like.

The translation of these and other ideas borrowed from cognitive science, social psychology and evolutionary biology into practical AI tools is at an early stage. However, they could provide fresh inspiration for new directions in AI developments, which need to move away from the current drive of continuous scaling of model sizes and datasets.

Published online: 17 November 2023

References

1. Preprint at <https://arxiv.org/abs/2112.04359> (2021).
2. Vincent, J. *The Verge* <https://go.nature.com/471Diho> (2022).
3. Pablo Villalobos, P. et al. Preprint at <https://arxiv.org/abs/2211.04325> (2022).
4. Colas, C., Karch, T., Moulin-Frier, C. & Oudeyer, P.-Y. *Nat. Mach. Intell.* **4**, 1068–1076 (2022).