
Building AI Without Reproducing the Brain

Odysseus Kaloidas
UC Berkeley

Abstract

How much useful intelligence can we build without reproducing the way humans remember, reason, and plan? Aircraft made flight useful at scale without reproducing birds; AI may likewise achieve substantial work through different mechanisms. An external document store can support memory, while search and checking can improve answers. The question is how far such arrangements can progress before their lack of human-like integration becomes a practical limit. The analogy does not establish that current methods will suffice. It makes the task, its costs, and its requirements the basis for judging a machine, rather than its resemblance to a person.

1 Flight Without a Bird

Useful AI need not reproduce human cognition in every detail to perform important work. An aircraft cannot hatch, heal its wings, or land on a twig. These omissions hardly diminish its value for crossing an ocean. AI discussions can lose this distinction when the remarkable breadth of human cognition becomes the implicit threshold for meaningful progress. A machine can fall short of that breadth while exceeding us on consequential work.

LeCun’s 2024 interview identifies limitations in reasoning, planning, and working memory, the temporary retention of information during a task, while acknowledging language models’ usefulness [6]. His May 2026 interview again separates commercial usefulness from his rejection of a path to human-level intelligence [2]. The distinction deserves to survive the public argument: failure to reproduce the whole human package is not a ceiling on useful machine capability.

We can therefore ask two questions separately. How much useful work can these systems perform as they improve? And will their improvements eventually produce the broad adaptability of human intelligence? A limited answer to the second need not imply a limited answer to the first.

2 The Aircraft Analogy, Precisely

This is an established engineering argument. Russell and Norvig use artificial flight to distinguish intelligent performance from human imitation [7]. LeCun himself rejects detailed brain copying through a bird-and-aircraft comparison [3]. It would therefore misrepresent him to suggest that his world-model program seeks a replica of neural tissue.

The harder question is what follows when a system lacks a human capability. A need for reliable planning does not tell us which algorithm must provide it. Failure today also leaves open what retrieval, search, different training, or more computation could achieve. Yet a list of possible additions is only a list: each still has to work, and the assembled system has to work too.

Marr’s separation of computational problems, algorithms, and implementations makes the distinction explicit [5]. Preserving a function need not preserve its biological realization. The analogy earns its place only when the function is specified: aircraft escape the need for feathers, not the need to generate lift and remain controllable.

3 Memory, Reasoning, and Planning Can Be Engineered Differently

An archive assistant does not need to recollect a document as a person would. Retrieval-augmented generation searches an external index and supplies relevant documents to the model when it answers [4]. This lets information remain outside the model's parameters, the numerical values adjusted during training. The arrangement can still retrieve an obsolete record, miss the relevant one, or misread what it finds. Its success depends on recovering and using the right evidence, whatever resemblance the process bears to human memory.

Reasoning performance can also depend on a system larger than one uninterrupted generation. Cobbe and colleagues improve mathematical problem solving by generating candidate solutions and using a verifier, a model trained to assess solutions, to select among them [1]. That result does not demonstrate general reasoning or guarantee correct verification. It demonstrates that a weakness of the initial generator need not be the final capability of the assembled system.

For planning, MuZero learns to predict quantities useful for choosing actions and searches possible continuations in Atari, Go, chess, and shogi [8]. Its successes do not establish competence in unfamiliar physical environments. They do show that consequential sequential decisions can emerge from a computational arrangement without a claim to reproduce human deliberation.

The external index still has to support memory; the verifier still has to judge; the planner still has to choose actions. Engineering these functions differently does not remove them. Nor do successes on separate tasks establish that the functions will work together in an unfamiliar situation.

4 Useful Capability Has a Different Boundary

For the archive assistant, users need accurate answers, traceable sources, acceptable response time, and affordable operation. Human-like episodic memory is one possible route to some of those functions, but it is not the service being requested. An external index can meet the need on its own terms. A robot expected to recover from unfamiliar physical failures has a more demanding specification.

Scale adds another reason to distinguish the job from its biological implementation. At comparable acceptable quality, a system that costs more to develop can become cheaper overall if each use costs less. With fixed development and per-use costs, repeated savings eventually repay the larger initial investment. Checking and correction belong in that calculation. The relevant advantage is reliable work delivered at lower total cost; biological resemblance matters where it changes that performance or cost.

Sutton's Bitter Lesson supplies a historical argument for methods that exploit expanding computation, especially search and learning [9]. It is not a theorem of unlimited improvement. A method that looks unlike human thought may nevertheless improve quickly because additional computation makes it more capable. Whether it does is an empirical question. Meanwhile, an imperfect system can already be worth building while a more comprehensive design remains unfinished.

5 When the Job Requires Integrated Adaptability

Human cognition combines mutually supporting capacities. Memory informs planning; physical understanding constrains plausible consequences; experience changes future decisions. An external index and a verifier may work on separate benchmarks yet fail together in an unfamiliar situation.

Zador argues that animal competence reflects inherited structure as well as lifetime learning [10]. Animals arrive with useful structure shaped before their own learning begins. That makes their efficiency a valuable source of ideas, even when an engineered system uses those ideas differently.

For a robot operating under tight energy limits with sparse feedback, such integrated adaptability may be central to the task itself. An aircraft analogy cannot explain away the deficit. For an archive-analysis system, matching animal autonomy may add little to the intended service. The difference lies in the job and its constraints.

The useful question is how far these alternative arrangements can take us. Some tasks may reward systems that store vastly more than a person, check thousands of candidates, or repeat a narrow procedure cheaply. Others may expose the cost of weak integration. The aircraft analogy invites attention to both: what work becomes possible through a different design, and what essential capability remains missing?

References

- [1] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training Verifiers to Solve Math Word Problems. arXiv:2110.14168, 2021. URL <https://arxiv.org/abs/2110.14168>.
- [2] Annelies Gamble. AI for the real world: A conversation with Yann LeCun. Mixture of Experts, May 12, 2026. URL <https://www.mixtureofexperts.co/p/ai-for-the-real-world-a-conversation>. Firsthand interview.
- [3] Lee Gomes. Facebook AI director Yann LeCun on his quest to unleash deep learning and make machines smarter. IEEE Spectrum, February 18, 2015. URL <https://spectrum.ieee.org/facebook-ai-director-yann-lecun-on-deep-learning>. Interview with Yann LeCun.
- [4] Patrick Lewis, Ethan Perez, Aleksandra Piktus, et al. Retrieval-augmented generation for knowledge-intensive NLP tasks, 2020. URL <https://arxiv.org/abs/2005.11401>. arXiv:2005.11401.
- [5] David Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. MIT Press, 2010. doi: 10.7551/mitpress/9780262514620.001.0001. Originally published in 1982.
- [6] Billy Perrigo. Meta’s AI Chief Yann LeCun on AGI, Open-Source, and AI Risk. TIME, February 13, 2024. URL <https://time.com/6694432/yann-lecun-meta-ai-interview/>. Interview with Yann LeCun, conducted January 26, 2024.
- [7] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2 edition, 2003. URL <https://aima.cs.berkeley.edu/4th-ed/pdfs/newchap01.pdf>. Chapter 1.
- [8] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, et al. Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 588:604–609, 2020. doi: 10.1038/s41586-020-03051-4. URL <https://www.nature.com/articles/s41586-020-03051-4>.
- [9] Richard S. Sutton. The bitter lesson. Author essay, March 13, 2019. URL <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>.
- [10] Anthony M. Zador. A critique of pure learning and what artificial neural networks can learn from animal brains. *Nature Communications*, 10:3770, 2019. doi: 10.1038/s41467-019-11786-6. URL <https://www.nature.com/articles/s41467-019-11786-6>.