Fostering Epistemic Diversity in the AI Era

Sofia G. Gallego[1]

[1]CosmoVentures, Paris, France, with Apart Research

Apart Research AI × Epistemics Research Sprint, November 2026

Abstract

Creativity, innovation and improvement often happen in the space between ideas. As more of our thinking passes through AI, that space can quietly shrink, or it can grow. Evidence from science, art, politics, law, culture and work points the same way: each person gains, and the whole loses range. This position paper outlines directions in which the community can foster epistemic diversity, and proposes one concrete step: a public map of ideas, where concepts are seen from many perspectives and the space between them stays in view, tested against the range of ideas real people reach.

1 Introduction

A word like *water* reads differently to a chemist, a farmer, a historian and a poet. Each reading is true, and new ideas often appear where two of them meet. Groups that think in different ways find what groups of the individually best problem solvers miss [1], and the most cited science tends to join familiar ideas in unusual combinations [2]. The world is more connected than ever, and so are its questions: climate, health, technology and justice cross every border between fields.

How can people and AI think together in a way that keeps the space between ideas open, and even widens it? A harmonic interaction would keep what works in each and enhance what can be enhanced. Section 2 gathers what is at stake across fields, Section 3 what has been tried, Section 4 where the community could go, and Section 5 proposes a map of ideas.

2 Background

Across many fields the same pattern appears: each person gains, and the whole loses range.

Science and creativity. Researchers who use AI publish more and are cited more, while the range of topics studied together shrinks [3], and questions that suit AI can crowd out others [4]. Writers and artists helped by AI produce work judged better, which becomes more alike over time [5, 6, 7].

Politics, law and institutions. People who write with a model that leans toward one view move their own opinions toward it [8]. When many decision makers use the same model, their errors stop being independent, and the same people are wrongly rejected everywhere [10, 11].

Culture and representation. Models answer most like people in Western, rich societies [30, 31], image generators amplify stereotypes [32], most languages have almost no place in language technology [33], and writing suggestions pull other cultures toward Western styles [12]. Bias and homogenization are one problem seen from two sides: whose perspective becomes the default.

Work, learning and power. The first jobs to thin out are those where young people start a career [26, 27], and fewer people entering a field means fewer perspectives growing inside it.

Students who write with AI write more alike and remember less [13]. A few providers build the models most people think with [29], and whether they replace people or extend them is a choice [28].

The models. Chat training reduces variety [14], models are more alike than people [15, 16], and rare ideas can fade from use, a process called knowledge collapse [17, 18].

3 Related Work

Measuring diversity. Benchmarks ask whether models give humanlike variety [19] and how much of the world's knowledge they cover [18].

Widening models. Personas and backstories make the answers of models more varied [20, 21].

Tools for groups. Polis maps opinions so groups see where they differ [22], Community Notes shows a note only when people who usually disagree both find it helpful [23], and the Habermas Machine drafts statements a divided group can share [9]. Each works on a question already asked. The space between ideas, where new questions come from, has fewer tools.

Earlier work by the author. Two earlier studies led to this proposal. In False epistemic redundancy, with a collaborator, AI agents given different reasoning personas sounded different and still missed the same cause [24]. In Epistemic Bridges, made during this sprint, people and models link the same pairs of words, and each model is placed on the map that people's answers draw [25]: the newer chat models reached less of it than people, even with their shared voice removed, and answering as an ordinary person drawn at random was the one change that widened every model's ideas. Steering a small model's activations toward that perspective brought it to the range of people while its knowledge and safety held.

4 Directions

The community can foster epistemic diversity by measuring it against what real people do, alongside accuracy and safety; by giving models perspectives and lives alongside better reasoning, and keeping several independent models in use where decisions affect many people; by building tools that show what lies around an idea before an answer settles it; by bringing fields and cultures together in debates and interdisciplinary residencies; by protecting the languages, cultures and entry-level work the market leaves out; and by carrying the evidence into policy and schools.

5 Proposal: The Map of Ideas

The step proposed here is a non-linear Wikipedia: a shared space where knowledge is read through its connections.

The idea has a long history. Vannevar Bush imagined a machine that would keep the associative trails a reader makes between documents [34], and Ted Nelson proposed links that run both ways [35]. Concept maps draw ideas and the relations between them for teaching [36], Wikidata links facts into one open graph [37], and Kialo (kialo.com) arranges the arguments for and against a single claim as a tree. Wikipedia itself is read one article at a time.

The Map of Ideas (mapideas.org, Figure 1) brings these threads together around the space between ideas. A link carries how one idea reads from the other, in both directions. A criticism or answer is an idea of its own, joined to what it answers, so different views stand side by side and nothing is ranked. Ideas connect across fields, from mathematics and astronomy to religions, cultures and art, and any idea can become the center. Each tradition is described as its own people

would. An AI guide answers from the map and never writes to it. A first version holds more than 300 ideas and open questions across 17 topics.

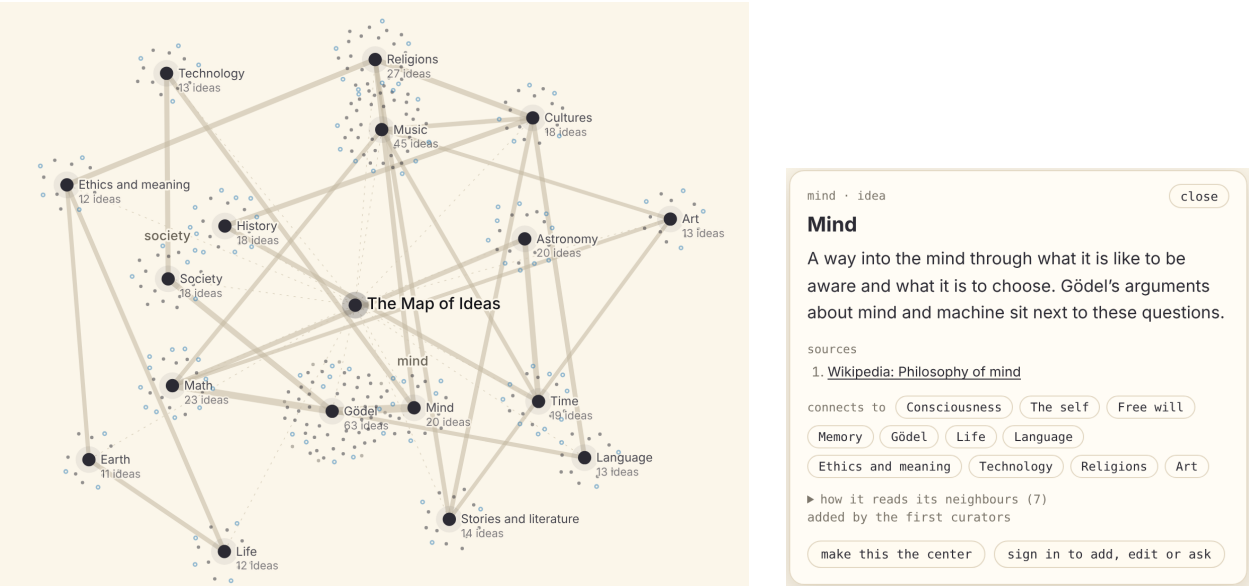


Figure 1: The Map of Ideas. Left, the whole map: topics and the lines between them, thicker where more ideas connect, with each topic's ideas around it. Right, an idea as a reader sees it, with its sources, the ideas it connects to across topics, and how it reads each neighbour.

For people outside a field. A map can also be read by someone who is not an expert. A reader sees the positions around a question and how they relate, and can follow a thread from what they know into what they do not. They can also add their own view, with its own biases and visions, next to the others, where it can be read and answered. For outreach, teaching and public debate, this turns a field into a place to walk through, and opens it to people who would not enter it through its papers.

How it could grow. Wikipedia began as Nupedia, an encyclopedia written only by experts, which grew too slowly; opening it to everyone made it grow [38]. The map can take both lessons. Experts from different fields and traditions seed and look after a small core for each topic, and everyone else adds, answers and connects around it, with new edits on the core waiting for a curator. Contributing should be simple and trusted: easy sign-in, a watchlist of what changed, marks for ideas that need more work. Personal maps can grow from people's own notes, shared by choice, and trails through ideas can be saved and passed on. Debates and interdisciplinary residencies held on the map leave what they find for others. The code and the content stay open, as Wikipedia's do, so that anyone can build on them.

How to know it works. A map shown to everyone is also a shared prior: it could widen what people reach, or become the new common answer. Bridges gives a way to tell. Players link the same words alone, after exploring the map, or after seeing a bridge from an AI, and each group is read against a reference map drawn by other players. The map fosters diversity if the group that used it reaches further and still lands where people go. After answering, players could also see where their own connection sits among everyone's.

6 Conclusion

The choices that shape how AI meets the space of ideas rest today with few, and the directions that could widen it belong to many. A map of ideas, tested against the range real people reach, is one place to start: if it widens what people find, it can grow with them; if it narrows it, the test will show that too.

This proposal is itself shaped by one perspective. The map was started by one person, is written mostly in English, and its first topics follow her interests. That is one more reason to test it against people other than its author, and to open it to the people of each field and tradition, so that the map grows beyond the view it started from.

Data and Code Availability

The Map of Ideas is at `https://mapideas.org`, and the Bridges game at `https://mapideas.org/bridges`; their code, and this paper in the `paper/` folder, are at `https://github.com/nefinia/mapideas`. The Bridges measures, AI answers and analysis are at `https://github.com/nefinia/epistemic-bridges`, and the steering pilot at `https://github.com/nefinia/steering-the-voice`. The Map of Ideas, the game and false epistemic redundancy existed before the sprint started on 13 November; the Bridges study, the steering pilot and this paper were made during it.

Author Contributions and Acknowledgements

S.G.G. conceived the vision, the map, the game and the studies, and wrote this paper with Claude (see LLM Usage Statement). Thanks to everyone who played Bridges and left feedback.

References

[1] L. Hong, S. E. Page. *Groups of diverse problem solvers can outperform groups of high-ability problem solvers.* PNAS, 101(46), 2004.

[2] B. Uzzi, S. Mukherjee, M. Stringer, B. Jones. *Atypical combinations and scientific impact.* Science, 342(6157), 2013.

[3] Q. Hao, F. Xu, Y. Li, J. Evans. *AI Expands Scientists' Impact but Contracts Science's Focus.* Nature (accepted), 2026. arXiv:2412.07727.

[4] L. Messeri, M. J. Crockett. *Artificial intelligence and illusions of understanding in scientific research.* Nature, 627, 2024.

[5] A. R. Doshi, O. P. Hauser. *Generative AI enhances individual creativity but reduces the collective diversity of novel content.* Science Advances, 10(28), 2024.

[6] E. Zhou, D. Lee. *Generative artificial intelligence, human creativity, and art.* PNAS Nexus, 3(3), 2024.

[7] B. R. Anderson, J. H. Shah, M. Kreminski. *Homogenization Effects of Large Language Models on Human Creative Ideation.* Creativity & Cognition, 2024.

[8] M. Jakesch, A. Bhat, D. Buschek, L. Zalmanson, M. Naaman. *Co-Writing with Opinionated Language Models Affects Users' Views.* CHI, 2023.

[9] M. H. Tessler et al. *AI can help humans find common ground in democratic deliberation.* Science, 386(6719), 2024.

[10] J. Kleinberg, M. Raghavan. *Algorithmic monoculture and social welfare.* PNAS, 118(22), 2021.

[11] R. Bommasani, K. A. Creel, A. Kumar, D. Jurafsky, P. Liang. *Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization?* NeurIPS, 2022.

[12] D. Agarwal, M. Naaman, A. Vashistha. *AI Suggestions Homogenize Writing Toward Western Styles and Diminish Cultural Nuances.* CHI, 2025.

[13] N. Kosmyna et al. *Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task.* arXiv:2506.08872, 2025.

[14] R. Kirk, I. Mediratta, C. Nalmpantis, J. Luketina, E. Hambro, E. Grefenstette, R. Raileanu. *Understanding the Effects of RLHF on LLM Generalisation and Diversity.* ICLR, 2024.

[15] E. Wenger, Y. N. Kenett. *Large language models are homogeneously creative.* PNAS Nexus, 5(3):pgag042, 2026.

[16] L. Jiang et al. *Artificial Hivemind: The Open-Ended Homogeneity of Language Models (and Beyond).* NeurIPS, 2025.

[17] A. J. Peterson. *AI and the Problem of Knowledge Collapse.* arXiv:2404.03502, 2024.

[18] D. Wright, S. Masud, J. Moore, S. Yadav, M. Antoniak, P. Ebert Christensen, C. Y. Park, I. Augenstein. *Epistemic Diversity and Knowledge Collapse in Large Language Models.* arXiv:2510.04226, 2025.

[19] Y. Zhang, H. Diddee, S. Holm, H. Liu, X. Liu, V. Samuel, B. Wang, D. Ippolito. *NoveltyBench: Evaluating Language Models for Humanlike Diversity.* arXiv:2504.05228, 2025.

[20] S. Moon et al. *Virtual Personas for Language Models via an Anthology of Backstories.* EMNLP, 2024.

[21] Y. Deng, M. Brucks, O. Toubia. *Examining and Addressing Barriers to Diversity in LLM-Generated Ideas.* 2026.

[22] C. Small, M. Bjorkegren, T. Erkkilä, L. Shaw, C. Megill. *Polis: Scaling Deliberation by Mapping High Dimensional Opinion Spaces.* Recerca, 26(2), 2021.

[23] S. Wojcik et al. *Birdwatch: Crowd Wisdom and Bridging Algorithms can Inform Understanding and Reduce the Spread of Misinformation.* arXiv:2210.15723, 2022.

[24] S. G. Gallego, A. Habana. *False Epistemic Redundancy: Do AI Ensembles Share a Blind Spot?* Apart Research Digital Minds Research Sprint, 2026. https://apartresearch.com/project/false-epistemic-redundancy-do-ai-ensembles-share-a-blind-spot-b50u

[25] S. G. Gallego. *Epistemic Bridges: Do Language Models Reach the Connections People Make?* Apart Research AI × Epistemics Research Sprint, 2026. https://apartresearch.com/sprints/projects/epistemic-bridges-do-language-models-reach-the-connections-people-make-y0xuxrgxw8

[26] T. Eloundou, S. Manning, P. Mishkin, D. Rock. *GPTs are GPTs: Labor market impact potential of LLMs.* Science, 384(6702), 2024.

[27] E. Brynjolfsson, B. Chandar, R. Chen. *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.* Stanford Digital Economy Lab, 2025.

[28] D. Acemoglu, S. Johnson. *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity.* PublicAffairs, 2023.

[29] A. Korinek, J. Vipra. *Concentrating Intelligence: Scaling and Market Structure in Artificial Intelligence.* Economic Policy, 40(121), 2025.

[30] M. Atari, M. J. Xue, P. S. Park, D. E. Blasi, J. Henrich. *Which Humans?* PsyArXiv, 2023.

[31] E. Durmus et al. *Towards Measuring the Representation of Subjective Global Opinions in Language Models.* COLM, 2024.

[32] F. Bianchi et al. *Easily Accessible Text-to-Image Generation Amplifies Demographic Stereotypes at Large Scale.* FAccT, 2023.

[33] P. Joshi, S. Santy, A. Budhiraja, K. Bali, M. Choudhury. *The State and Fate of Linguistic Diversity and Inclusion in the NLP World.* ACL, 2020.

[34] V. Bush. *As We May Think.* The Atlantic Monthly, July 1945.

[35] T. H. Nelson. *Complex information processing: a file structure for the complex, the changing and the indeterminate.* ACM National Conference, 1965.

[36] J. D. Novak, A. J. Cañas. *The Theory Underlying Concept Maps and How to Construct and Use Them.* Technical Report IHMC CmapTools 2006-01 Rev 01-2008, Florida Institute for Human and Machine Cognition, 2008.

[37] D. Vrandečić, M. Krötzsch. *Wikidata: a free collaborative knowledgebase.* Communications of the ACM, 57(10), 2014.

[38] J. M. Reagle. *Good Faith Collaboration: The Culture of Wikipedia.* MIT Press, 2010.

LLM Usage Statement

Claude helped build the Map of Ideas and the Bridges game before the sprint. During it, Claude searched for and checked the literature, drafted and revised this paper with the author, and made the figure of the map; the ideas, the framing and the proposal are the author's, and she directed every revision. All numbers come from the cited studies and from the analysis scripts of the Bridges study and the steering pilot.