Epistemic Monocultures and the Effect of AI Personalization

Joseph O'Brien (j3obrien@ucsd.edu)
Department of Philosophy, University of California, San Diego
La Jolla, CA USA

Sina Fazelpour (s.fazel-pour@northeastern.edu)
Department of Philosophy and Religion & Khoury College of Computer Sciences, Northeastern University
Boston, MA USA

Hannah Rubin (hmrfyb@missouri.edu)
Department of Philosophy, University of Missouri
Columbia, MO USA

Kekoa Wong (me@kekoawong.com)
Verily, Google
Boston, MA USA

Abstract

It has been argued that when scientists employ algorithmic tools to assist in problem-solving, epistemic monocultures may emerge in which research tools, topics, findings, etc. are homogenized. As a result, fertile areas of research might be left unexplored, impeding scientific progress. To explore the nature of these epistemic monocultures, we develop an agent-based model where agents have the ability to query an AI system to assist in their search of an epistemic (NK) landscape. In general, we find that AI use negatively affects the community of researchers by reducing heterogeneity, but both the rate of AI queries and how AI is used impact the ultimate success of the community. We then implement a potential solution suggested in the literature, AI personalization, and find somewhat mixed results on its potential for mitigating homogenization in research communities.

Keywords: social networks; epistemic homogenization; AI in science; computational modeling

Introduction

Outcome Homogenization

When a relatively small set of models are employed in research settings, downstream AI models may inherit the biases of the foundation models they are based on (Bommasani et al., 2021). Recent work on this phenomenon of algorithmic homogenization postulates that the use of a relatively small set of foundation models might result in the homogenization of outcomes for users (Kleinberg & Raghavan, 2021). In work further exploring algorithmic homogenization, Bommasani, Creel, Kumar, Jurafsky, and Liang (2022) advance the "component sharing hypothesis," which argues that algorithmic systems built using shared components, whether model architectures or datasets, will lead to homogenization of outcomes for users of these systems. To explore the effects of algorithmic homogenization in the sciences, we develop an NK landscape model where agents are given the ability to query AI and receive suggestions about changes to their research practices that improve fitness.

Messeri and Crockett (2024) develop a taxonomy for how AI might be used in scientific practice. Using their taxonomy, we explore the conceptualization of "AI as Oracle" and "AI as Quant." Under the AI as Oracle conceptualization, AI is used by scientists to digest and communicate published scientific knowledge. Messeri and Crockett write that as the amount of published research expands beyond what any one researcher might reasonably be expected to digest, AI might be employed to assist in sifting through the deluge of articles to determine which might be of use to a given scientist. Messeri and Crockett write that the use of AI as Oracle may result in "illusions of objectivity," where scientists mistakenly believe that the algorithmic tools they use do not have a standpoint. But as Bommasani et al. (2022) suggest, AI models do inherit the biases of their training data and developers. This conflict may result in outcome homogenization, as scientists may come to believe the literature summaries provided by AI are complete and free of bias.

The "AI as Quant" conceptualization involves the use of AI to extract meaning from the ever-increasing complexity of modern datasets. The computational power offered by AI systems might allow scientists to elucidate connections between variables in the data that might not have otherwise been apparent. Messeri and Crockett argue that this use of AI as Quant might lead to the illusion of explanatory depth, whereby scientists using AI may have an inflated perception of their own scientific understanding. When AI systems suggest changes to scientific research practices or offer insights into complex data, scientists might mistakenly believe they have explanations for the causal mechanisms at play. These arguments are supported by recent surveys asking scientists how AI might impact their research field (Van Noorden & Perkel, 2023). Of 1,600 respondents, 66% agreed that AI would assist in processing data, and 58% responded that AI would allow scientists to carry out computations that were previously infeasible. 69% of researchers agreed that the use of AI in science may lead to an overreliance on pattern recognition without the development of scientific understanding.

Using the NK Landscape framework, we show that when agents have access to a non-personalized AI Oracle or Quant,

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

they exist in an epistemic monoculture, since the AI provides the same output regardless of the user. We observe the negative effects of epistemic moncultures by measuring the loss of transient diversity in the community. We then implement a personalized version of both "AI as Oracle" and "AI as Quant" and show that personalization can be effective in mitigating epistemic monocultures, so long as the rate of AI query remains low.

Methods

The NK Landscape

We construct an agent-based model to explore the impact of algorithmic tools on scientific theorizing. We use the NK-landscape framework from Kauffman and Levin (1987), which has previously been used to explore questions surrounding scientific collaboration and diversity of research practices (Lazer & Friedman, 2007; Gomez & Lazer, 2019; Alexander, Himmelreich, & Thompson, 2015; Wu, 2023; Gabriel & O'Connor, 2024). This framework, originally used to explore cell evolution in genetic networks, captures the complexity of scientific practice more clearly than the traditional epistemic landscape framework, due to the ability to capture interdependencies between a set of research practices (Weisberg & Muldoon, 2009; Alexander et al., 2015). The NK landscape is characterized by two parameters where N is the number of dimensions of the landscape, and $1 \leq K \leq N-1$ represents the ruggedness of the landscape. Each agent is randomly initialized on the landscape with a vector of length N representing their location. The bits of this vector take a value of 0 or 1 and represent a given agent's beliefs, tools, and research practices. Agents explore the landscape by randomly choosing a bit of their vector, changing its value, and evaluating the resulting change in fitness. By randomly changing bits of their vector, agents attempt to discover the set of beliefs, tools, etc. that maximize their fitness. The change of fitness that results from changing a bit is dependent on a randomly initialized mask of K other bits, capturing the interdependencies that exist in a set of research practices. For the sake of simplicity, for all reported results, we set $N = 20$ and $K = 5$.[1]

We also initialize a scale-free graph called the Barábasi-Albert graph (Barabási & Albert, 1999). The Barábasi-Albert graph takes two parameters as input, where m is used to construct an initial fully-connected subgraph of m nodes, and n is the total number of nodes in the final output graph.[2] Many researchers have suggested that scientific citation networks do exhibit this scale-free feature ((Price, 1965); (Redner, 2005); (Zhong & Liang, 2024)), and the Barási-Alberts graph is just one of the various approaches to constructing scale-free graphs. When an agent forms a connection with another, they gain the ability to see that agent's solution during social learning.

We take as our starting point the work of (Lazer & Friedman, 2007), who use NK-landscapes to explore social learning for scientific problems. They include a velocity parameter, v, to capture the rate of social learning. Every v rounds, agents engage in social learning by adopting the solution of one of their neighbors which results in the highest increase in fitness.[3] If none of an agent's neighbors have a higher fitness, the agent retains its own solution. They show that increased rates of social learning lead to communities performing better in the short term but worse in the long run, as it is easier for agents to get stuck in local optima. Research communities that frequently look to their neighbors to adopt solutions spend less time exploring the landscape, sacrificing the ability to find solutions with higher fitness levels. These results suggest a trade-off between maintaining a diversity of research practices and rapidly communicating previously discovered successful strategies.

Problem Decomposability

We extend Lazer and Friedman's original model by considering how scientists might interact with algorithmic tools and AI. Lazer and Friedman (2007) discuss how their model might be extended by considering agents who specialize in some subsection of the problem space. Here, we attempt to capture the notion of decomposability formulated by Lazer and Friedman and explore how some scientific problems might be decomposable into subsections, some of which AI systems might be particularly well positioned to explore.

Recent research on self-driving labs (SDLs) lends credence to the claim that there exist scientific problems that are decomposable, where different agents might be well positioned to explore different aspects of the problem. Abolhasani and Kumacheva (2023) write that one common challenge in chemistry and materials science is the need to explore a vast number of variables to find the most effective compositions or manufacturing routes for a given compound. Current strategies rely on the specific domain knowledge of scientists and require changing variables one at a time in a combinatorial fashion, which can be costly and time-consuming. SDLs help solve this problem by integrating machine learning with robotics to explore vast swaths of the chemical space with very little input from scientists. When the SDL locates a promising chemical pairing, it can alert the scientists, who can then perform a more in-depth exploration of the pairing. This provides just one example of how current scientific problems can be decomposed into parts for which AI might prove exceedingly useful.

Study 1: AI as Oracle

To capture the idea of "AI as Oracle", we give agents a subroutine in their decision procedure we call "query" which

[1] Setting $N = 20$ and $K = 5$ is the standard for much of the literature on NK landscapes. See the discussion section for results when K is increased.

[2] For reported results, we set $m = 20$ and $n = 100$. In experiments varying the value of m, we find no significant impacts on long-run community fitness.

[3] In our reported results, we keep velocity set at 25%. In tests varying the velocity parameter, we observe many of the same phenomena as Lazer and Friedman (2007).

allows them to query an AI system in order to receive a suggestion about how they ought to change their set of research practices. We conceive of "AI as Oracle" as being a form of quasi-social learning, so agents in these simulations probabilistically query AI rather than engage in social learning with their neighbors on the network.

The Query Subroutine

We conceptualize the idea of decomposability on the NK landscape by giving the AI the ability to operate on multiple bits of the agent's vector during a single time step. The "AI as Oracle" conceptualization argues that AI might be best used for conducting literature reviews and generating hypotheses, while humans are better suited for iterating on the suggestions provided by AI.

Non-personalization The non-personalized AI randomly selects a solution from the ten best-performing agents and provides the last $N//5$ bits of the solution to the querier. The non-personalized AI can be thought of as an "extended social learning with randomization" procedure. The AI has access to all of the agent's solutions from the previous round. This time lag functions much in the way that current language models do. Language models are not updated instantaneously, and querying for relevant papers may involve a time lag of this sort. Agents also only adopt a solution that makes them better off, protecting against cases where the most successful agent is made worse off by being given a worse performing solution.[4] In this run of simulations, the non-personalized AI can broaden the scope of research an agent has access to beyond their neighbors on the network. While agents might not have direct access to the best-performing agent due to their location on the network, querying the AI provides them with a solution that might not be accessible to them via social learning.

Personalization The personalized query subroutine operates on the last $N//5$ bits of the agent's vector.[5] The AI tests every permutation of these bits and provides the agent with the solution that leads to the greatest increase in fitness, given their current location. Importantly, both velocity and the AI trigger are probabilistic parameters, such that agents do not all engage in social learning or AI queries during the same round. Whereas some NK landscape models have all agents undergo social learning in the same round, we think it is more plausible to suggest that the sorts of scenarios where social learning might occur, whether in workshops or in email exchanges, this hardly ever involves the entire community of scientists. Likewise, we suggest that the use of AI is a largely solitary pursuit, and thus, agents should choose to query the AI each round with some non-zero probability.

[4]See also (Wu, 2023), where agents randomly select a better-performing solution from their neighbors, rather than always choosing the best-performing solution.

[5]In tests on the effects of increasing the number of bits that AI can manipulate, we find only marginal differences in community outcomes.

Results

Here we present the results of our simulations for the "AI as Oracle" conceptualization. For each set of parameters, 1,000 simulations were run to lower the observed variance between simulations. For each community, we measure the average epistemic fitness at each time step and explore how the rate of AI query affects community-level outcomes. Importantly, agents take only one action each round. Every round, there is some probability that an agent will engage in social learning. If an agent engages in social learning, there is a probability that they will choose to query the AI. Finally, if none of these actions are taken, the agent explores the landscape by randomly manipulating a single bit of their solution.

In this run of simulations, we compare both the personalized and non-personalized AI communities to a community of agents that lack AI, and a mixed community where 50% of agents have access to the personalized AI. We observe that increased rates of AI queries have negative impacts on the community. Communities with access to AI perform worse when the rate of AI query increases, but the mixed community is able to mitigate this effect. We observe that the non-personalized community outperforms even the personalized AI community during the early rounds of the simulations, but in the long run, the non-personalized community converges to a suboptimal solution.

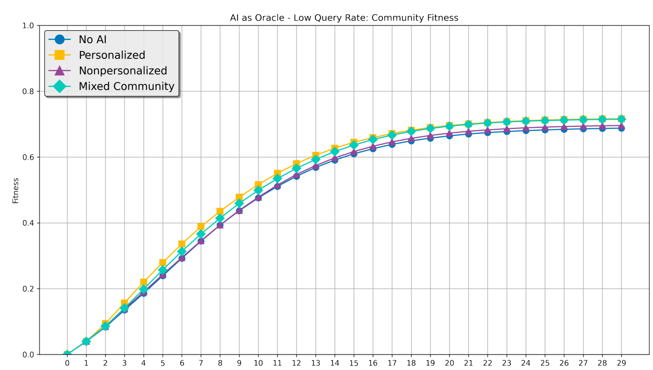

Figure 1: Initial Nodes = 20, Velocity = 0.25, AI Trigger = 0.25, Proportion of agents with AI access = 0.5

When the rate of AI query is relatively low (Figure 1), the community of agents with personalized AI outperforms all other communities. Interestingly, in the early rounds of the simulations, agents with access to a non-personalized AI perform at a level on par with the personalized community but quickly converge to a suboptimal solution in the long run.

As the rate of AI use increases (Figure 2), the personalized AI community also converges to a suboptimal solution. However, the mixed community appears not to suffer from the effects of increased AI usage. When query rates increase, the mixed community outperforms every other community.

In order to explore questions surrounding the formation of epistemic monocultures, we also introduce Hamming distance as a proxy for measuring outcome homogenization and

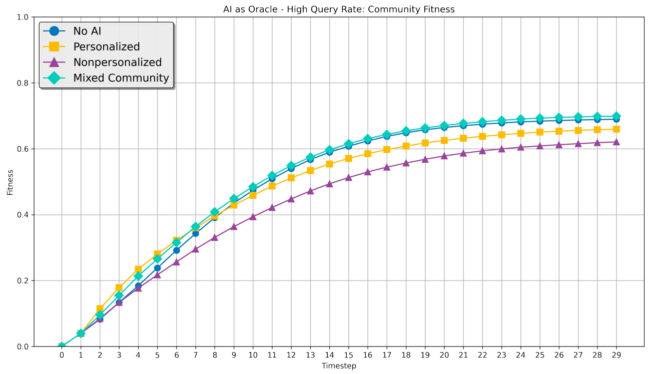

Figure 2: Initial Nodes = 20, Velocity = 0.25, AI Trigger = 0.75, Proportion of agents with AI access = 0.5

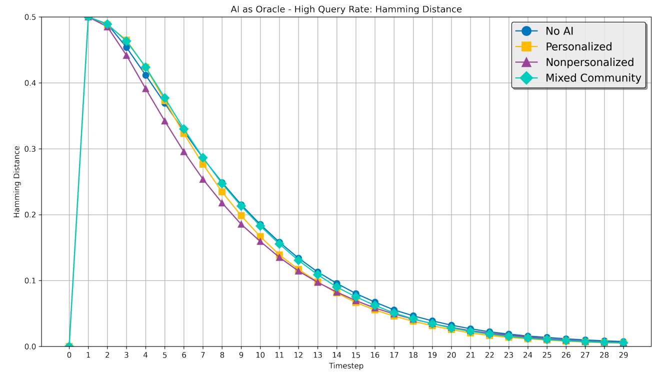

Figure 4: Initial Nodes = 20, Velocity = 0.25, AI Trigger = 0.75, Proportion of agents with AI access = 0.5

the diversity of research practices in a given community. For two one-dimensional vectors, the Hamming distance is simply the proportion of disagreeing bits of the two vectors. We use Hamming distance since Euclidean distance is difficult to measure on high-dimensional landscapes like the NK landscape.

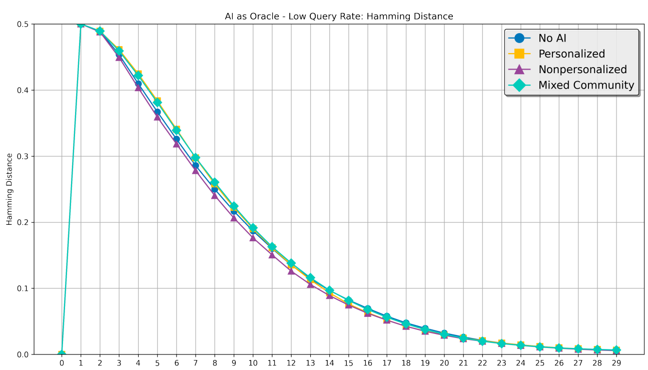

Figure 3: Initial Nodes = 20, Velocity = 0.25, AI Trigger = 0.25, Proportion of agents with AI access = 0.5

Looking at the average Hamming distance between agents across rounds, we find a possible explanation for the non-personalized community's poor performance.[6] The Hamming distance for the non-personalized community decreases more rapidly with the community without AI access, indicating a loss of transient diversity (Figure 3). This result adds evidence to the claim that AI usage might result in scientific communities leaving fruitful areas of the research landscape unexplored by creating epistemic monocultures.

As one might expect, as the rate of AI query continues to increase, we see an even steeper decline in Hamming distance for the community with access to non-personalized AI (Figure 4). Increasing the rate at which agents query the AI appears to exacerbate the effects of homogenization. Interestingly, when the rate of AI query is low, we do not observe the same rapid decrease in Hamming distance for the community with access to personalized AI (Figure 3. This is in line with suggestions from (Bommasani et al., 2022), who write that the personalization of AI systems may be a possible solution for mitigating outcome homogenization and the formation of epistemic monocultures. One reason the personalized community has an average Hamming distance akin to the community without AI is that the personalized AI subroutine iterates over the last $N//5$ bits of an agent's solution, returning the solution with the maximal fitness increase *given an agent's current location.*

Study 2: AI as Quant

In the conceptualization of "AI as Quant", the query subroutine is meant to capture two different notions of expertise in human-AI teams. Human agents have expertise in that they have access to a greater number of interventions. They can make changes to their set of research practices at any location along their vector, while the AI can only suggest changes on a restricted subset of the vector. If we think of the agent's vector as their set of beliefs, tools, and research practices, it makes sense to think that the AI might not have the ability to change all of these. It might be much more difficult for AI to change a given scientist's beliefs rather than suggest a change in their use of tools. The expertise of AI is captured by the ability to change more than one bit of the vector in a single time step. While agents explore the search space by randomly selecting a single bit to flip and evaluating the change in fitness, the AI can deal with increased complexity by suggesting changes for $N//5$ bits simultaneously. This is in accord with Messeri and Crockett's argument that AI might be used to draw insights from complex data.

The Query Subroutine

We conceptualize this idea of modularity on the NK landscape by giving the AI the ability to operate on multiple bits

[6]Note that during time step one, all of the agents are initialized, which explains the jump in Hamming distance from 0 to 0.5 in Figures for Hamming distance. Since we compute the average Hamming distance for the entire community, the maximum average Hamming distance is 0.5.

of the agent's vector during a single time step. The "AI as Quant" conceptualization is meant to argue that AI might be best used for breadth-first search, while humans are better suited to depth-first search. Importantly, whereas "AI as Oracle" can be thought of as an alternative to social learning, we view "AI as Quant" as an alternative to exploring the landscape. Instead of exploring, an agent may choose to query the AI in order to evaluate how changing multiple research practices may impact their level of epistemic significance.

Non-personalization The suggested solution is the last $N//5$ bits of the best-performing agent. This subroutine can be thought of as "partial social learning." While agents might not have direct access to the best-performing agent due to their location on the network, querying the AI gives them a partial solution that they can iterate on in an attempt to increase their own level of epistemic significance.

Personalization The personalized "AI as Quant" operates in much the same way as our "AI as Oracle." In the "AI as Oracle" simulations, we might think of the AI as offering research articles that are directly relevant to the agent's research. In contrast, the "AI as Quant" conceptualization operates on the last $N//5$ bits, suggesting changes they should make to their research practices or sets or tools.

Results

The setup for this run of simulations is largely similar to Study 1. The important difference is that rather than having access to AI during social learning, agents in the "AI as Quant" simulations probabilistically query AI on rounds that they explore the landscape. Every round, there is some probability that an agent will engage in social learning. If they do not engage in social learning, there is a probability that an agent with access to AI will choose to query the AI. Finally, if none of these actions are taken, the agent explores the landscape by randomly manipulating a single bit of their solution.

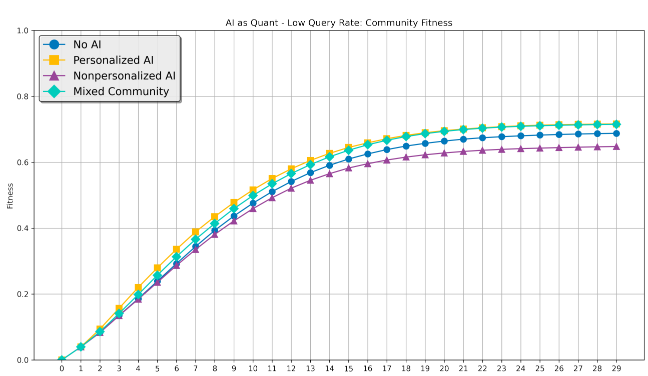

Figure 5: Initial Nodes = 20, Velocity = 0.25, AI Trigger = 0.25, Proportion of agents with AI access = 0.5

The first case we study is one in which the probability of agents querying the AI is relatively low. We observe that when the probability of querying the AI is 0.25 the community with access to the personalized AI performs the best in both the short and long run, but still performs only marginally better than the mixed community where only half of the agents have AI access (Figure 5). This suggests that it may not be necessary for every member of the research community to employ AI in the course of their research. So long as rates of social learning remain under the thresholds discussed by (Lazer & Friedman, 2007), the solutions found by agents with access to AI are able to spread throughout the community, including to those agents who might not have access to personalized AI.

These observed patterns hold across a range of variables, for which the personalized AI community tends to perform best overall, followed closely by the mixed community. Somewhat unsurprisingly, regardless of the parameter combination, the community with access to a non-personalized AI is always the worst-performing. This AI routine suggests that agents change the last $N//5$ bits of their vector to match the best-performing agent. We might gloss this AI subroutine as a sort of "mini social learning" procedure, as agents who query this AI receive a portion of the solution from the best-performing agent in the community. Unfortunately, due to the inter-dependencies between bits of the vector, in later rounds of the simulation, this subroutine tends to make agents worse off, so they choose not to adopt the solution suggested by the AI. Since the AI was queried, the agent also forgoes the chance to explore the landscape, and as such, the community tends toward stagnation when the probability of querying the AI is high.

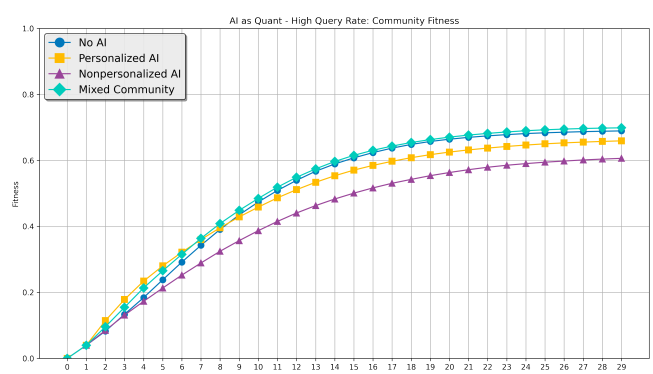

Figure 6: Initial Nodes = 20, Velocity = 0.25, AI Trigger = 0.75, Proportion of agents with AI access = 0.5

We observe a different phenomenon when the probability of querying AI continues to increase (Figure 6). Here, we see that when the probability of triggering the AI query is 0.75, both the mixed community and the community without access to AI perform better than the personalized AI community. This phenomenon is similar to one regarding social learning from (Lazer & Friedman, 2007). They observe that high rates of social learning actually have a negative impact on fitness in the community in the long run. This is similar to what we observe in the case where the probability of querying the AI is high. The personalized AI community performs besst in the short run, but as the community continues to explore, the high rate of AI queries causes the community to prematurely converge to sub-optimal solutions.

When we measure the epistemic benefit to agents when they query the personalized AI, we find that the average gain in fitness in early rounds tends to be quite high, and payoffs tend to decrease in future rounds. In late rounds of the simulation, the increase in fitness an agent gets from querying the personalized AI is often close to zero, resulting in a similar situation to the non-personalized AI community. These agents still choose to query AI, but upon receiving a solution that does not result in an increase in fitness, the agent is also unable to engage in a search of the nearby landscape. When the probability of querying the AI is too high, this happens to the majority of agents in later rounds, and the community reaches a sub-optimal solution.

Another point of interest is that even when the probability of querying the AI is high, the mixed community still performs better than the community with access to the AI. This suggests that having a community where only a portion of agents have access to AI may help mitigate the early stagnation that occurs when a community is comprised entirely of AI agents. This mixed community also outperforms the community of agents without access to AI. This is likely because the mixed community can make use of the large increases in fitness that agents with access to AI get in early rounds, while in later rounds, they can rely on their exploratory members to continue to drive the fitness of the community. These results are similar to a phenomenon observed by (Wu, 2023), who shows that a mixed community of agents, where some of them adopt the best strategy from their neighbors during social learning while others adopt a random better-performing strategy, outperforms both the community entirely comprised of "best" agents and the community entirely comprised of "better" agents. Mixed communities appear to have knock-on benefits that allow them to leverage the strategy of each of the sub-communities that comprise them.

When we observe the Hamming distance for the "AI as Quant" simulations, we observe results similar to our "AI as Oracle" study.[7] One possible reason for the loss of transient diversity and resulting underperformance in the non-personalized community is that, in the first few rounds, when there are relatively few well-performing agents in the community, it may be useful for AI to provide researchers with a portion of the most successful agent's solution. However, in later rounds, the suggested solution from AI frequently makes an agent worse off, so they choose not to accept the provided solution. The choice not to accept the suggestions of AI ensures that some diversity of research practices is maintained, but results in a situation where agents waste time each round querying the AI rather than making progress exploring the landscape. A similar phenomenon appears to hold regardless of the rate of AI query.

Having explored both the "AI as Oracle" and "AI as Quant" conceptualizations from Messeri and Crockett (2024), we find evidence that overreliance on AI might result in the formation of an epistemic monoculture for agents with access to a non-personalized AI. Using Hamming distance, we show that during early rounds, there is a loss of transient diversity. This is an interesting result, as many authors have argued that maintaining transient diversity is important for ensuring scientific progress (Zollman, 2010; Fazelpour & Rubin, 2022; Wu & O'Connor, 2023).

General Discussion

Our results show that AI personalization may be a successful solution for preventing monocultures from forming, but increased rates of queries lead to underperformance even when agents have access to personalized AI. In testing the effect of varying the ruggedness of the landscape on community outcomes (by increasing the value of K to $K = 10$), we find that on an extremely rugged landscape, increasing the probability of querying AI is beneficial for agents with access to a personalized AI. This is likely because on the rugged landscape, there is less correlation between nearby locations, so the increased rate of AI query remains beneficial in later rounds of the simulations.

Interestingly, the mixed community does not fall prey to worries about increased query rates. Even when the probability of querying AI increases, the mixed community still outpaces all other communities. This leads to interesting questions about the impact of AI allocation on the mixed community. In our simulations, AI is randomly allocated to 50% of the agents in the mixed community, but future research should explore this assumption and its effects on the community. One might vary the network structure and explore how the allocation of AI to central or peripheral nodes on the network impacts the community's fitness.

We think it is likely that in the real world, AI may be unevenly allocated, such that it is available to a minority of researchers at large, well-funded institutions. It's likely that these well-funded researchers may be centrally located on the social network, and it would be a worthwhile project to determine whether central or peripheral AI access has an impact on the community-level outcomes. Exploration into these questions may help shed additional light on questions regarding AI's use in scientific practice and the formation of epistemic monocultures.

References

Abolhasani, M., & Kumacheva, E. (2023). The rise of self-driving labs in chemical and materials sciences. *Nature Synthesis*, *2*(6), 483–492.

Alexander, J. M., Himmelreich, J., & Thompson, C. (2015, July). Epistemic landscapes, optimal search, and the di-

[7]Figures are not presented here for the sake of space, but we observe the same phenomena as in the results for "AI as Oracle" Hamming disance.

vision of cognitive labor. *Philosophy of Science*, *82*(3), 424–453. doi: 10.1086/681766

Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *science*, *286*(5439), 509–512.

Bommasani, R., Creel, K. A., Kumar, A., Jurafsky, D., & Liang, P. S. (2022, December). Picking on the same person: Does algorithmic monoculture lead to outcome homogenization? *Advances in Neural Information Processing Systems*, *35*, 3663–3678.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... others (2021). On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.

Fazelpour, S., & Rubin, H. (2022). Diversity and homophily in social networks. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).

Gabriel, N., & O'Connor, C. (2024, April). Can confirmation bias improve group learning? *Philosophy of Science*, *91*(2), 329–350. doi: 10.1017/psa.2023.176

Gomez, C. J., & Lazer, D. M. J. (2019, November). Clustering knowledge and dispersing abilities enhances collective problem solving in a network. *Nature Communications*, *10*(1), 5146. doi: 10.1038/s41467-019-12650-3

Kauffman, S., & Levin, S. (1987, September). Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, *128*(1), 11–45. doi: 10.1016/S0022-5193(87)80029-2

Kleinberg, J., & Raghavan, M. (2021, June). Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences*, *118*(22), e2018340118. doi: 10.1073/pnas.2018340118

Lazer, D., & Friedman, A. (2007, December). The network structure of exploration and exploitation. *Administrative Science Quarterly*, *52*(4), 667–694. doi: 10.2189/asqu.52.4.667

Messeri, L., & Crockett, M. J. (2024, March). Artificial intelligence and illusions of understanding in scientific research. *Nature*, *627*(8002), 49–58. doi: 10.1038/s41586-024-07146-0

Price, D. J. D. S. (1965). Networks of scientific papers: The pattern of bibliographic references indicates the nature of the scientific research front. *Science*, *149*(3683), 510–515.

Redner, S. (2005). Citation statistics from 110 years of physical review. *Physics today*, *58*(6), 49–54.

Van Noorden, R., & Perkel, J. M. (2023). Ai and science: what 1,600 researchers think. *Nature*, *621*(7980), 672–675.

Weisberg, M., & Muldoon, R. (2009, April). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, *76*(2), 225–252. doi: 10.1086/644786

Wu, J. (2023, October). Better than best: Epistemic landscapes and diversity of practice in science. *Philosophy of Science*, 1–11. doi: 10.1017/psa.2023.129

Wu, J., & O'Connor, C. (2023). How should we promote transient diversity in science? *Synthese*, *201*(2), 37.

Zhong, X., & Liang, H. (2024). On the scale-free property of citation networks: An empirical study. In *Companion proceedings of the acm web conference 2024* (pp. 541–544).

Zollman, K. J. (2010). The epistemic benefit of transient diversity. *Erkenntnis*, *72*(1), 17–35.