Title
Cognitive diversity promotes collective creativity: an agent-based simulation

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors
Rocca, Roberta
Tylen, Kristian

Publication Date
2022

Peer reviewed

Cognitive diversity promotes collective creativity: an agent-based simulation

Roberta Rocca (rbrrcc@gmail.com)
Department of Psychology, University of Texas at Austin, Austin (TX), USA, and
The Interacting Minds Centre, Aarhus University, Aarhus, Denmark

Kristian Tylén (kristian@cc.au.dk)
Department of Linguistics, Cognitive Science, and Semiotics, Aarhus University, Aarhus, Denmark, and
The Interacting Minds Centre, Aarhus University, Aarhus, Denmark

Abstract

In an agent-based simulation, we investigate the implications of social interaction and cognitive diversity on creative processes of divergent thinking. Agents performed a verbal association task individually and jointly in pairs. We created pairs of varying cognitive diversity by manipulating properties of the vector spaces defining their semantic memory. We find that cognitive diversity positively stimulates the flexibility of agents' collective cognitive search, giving rise to higher fluency (more solutions) and originality (more 'rare' solutions). While cognitively similar agents tend to exploit local semantic neighborhoods, diversity promotes more explorative search, with longer distances traveled in semantic space. This helps diverse pairs reach more distant areas of semantic space and escape cognitive fixation. However, our model also suggests that too high levels of diversity can have detrimental effects, as overly exploratory behaviors make pairs leave solution saturated areas prematurely and increase the risk of reaching semantic "dead ends".

Keywords: Agent-based modeling, social interaction, cognitive diversity, divergent thinking, creativity

Introduction

Across a number of domains within design, innovation, research and education, ideation processes unfold in collaborative contexts where two or more individuals interact to find novel and useful solutions to a problem. Divergent thinking is considered a central component of creative ideation, and many classical creativity tests attempt to measure the ability of individuals to flexibly and fluently produce as many and as different candidate solutions as possible in response to a prompt (Baer, 2014; Runco, 2010).

The influence of social interaction on creative processes of divergent thinking remains controversial: some studies report benefits from interaction while some do not (Aggarwal & Woolley, 2019; Brophy, 1998; Kohn & Smith, 2011; Mullen et al., 1991). One problem in this regard is that most existing studies only measure the performance of individuals and groups "offline" in terms of the number and quality of resulting solutions, while the underlying cognitive mechanisms of unfolding divergent thinking processes are often not accessible (Said-Metwaly et al., 2017).

The purpose of this study is to address the underlying computational cognitive and social mechanisms of collective divergent thinking, in order to unravel their dynamics and how they relate to performance. In particular, we use agent-based simulation to investigate how the degree of cognitive diversity between interaction partners affects collective search processes.

Cognitive search as information foraging

A prevalent metaphor in the field of problem solving is the idea of a solution space. When presented to a problem, the problem-solver searches for a solution by navigating a mental 'space of possible solutions' analogous to moving through a landscape (Newell & Simon, 1972). Some solutions may appear more immediately accessible and intuitive, that is "closer" in space, while others are located "further away" and might be hard to find.

With analogy to animal foraging behavior, it has been suggested that the mental search for ideas, memories, or solutions unfolds as an 'information foraging' process characterized by a succession of short and long 'jumps' through the solution space (Baronchelli & Radicchi, 2013; Hart et al., 2017; Szary & Dale, 2014). In this context, the short jumps correspond to a situation where a series of closely related solutions, often belonging to the same domain, category or kind, are visited, referred to as a phase of *exploitation*. A long jump, on the other hand, refers to a situation where a solution is found which is relatively distant from the last solution visited and thus is less available by association. Phases consisting of such long jumps are commonly referred to as *exploration* phases, as they often straddle domains, categories or kinds (Hills et al., 2015). In statistical models of foraging behavior, an optimal search pattern is one that presents a particular distribution between short and long jumps: you want to exhaust a local neighborhood before you move on, but you do not want to risk 'getting stuck' as the local neighborhood get sparser. Therefore, if a new and more solution-saturated neighborhood becomes available, you might want to jump there (Baronchelli & Radicchi, 2013).

Creativity and divergent thinking

Related to the concept of information foraging is the idea that human creativity is characterized by a particularly flexible style of cognitive search. A creative solution is often defined as one that is novel and useful, and creative processes can be portrayed as the search for particular novel and useful solutions (Kaufman & Sternberg, 2010). Several classical creativity tests (e.g., the Alternative Uses Test, AUT) thus measure the extent to which the participant is capable of divergent thinking, that is, providing as many different and

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

original solutions in response to a prompt as possible within a set time frame (Gilhooly et al., 2007).

Performance in divergent thinking is often measured along three dimensions: *fluency*, *flexibility* and *originality*. Fluency is the number of solutions provided in response to the prompt, often within a set time frame. Flexibility is the diversity or distance between solutions, that is, how large a portion of the solution space is sampled. Lastly, originality is related to the frequency of the individual solutions. Some obvious solutions might be presented by most participants. An original solution is thus one that only very few participants will list (while still being an appropriate response to the task) (Reiter-Palmon et al., 2019).

Social interaction and cognitive diversity

Creativity is often, more or less implicitly, portrayed as a property of an individual person or a cognitive process unfolding in an individual mind (Kaufman & Sternberg, 2010). And although creative practices in many contexts unfold in groups of multiple individuals, it is not clear how creative processes are influenced by social interaction (Brophy, 1998).

Across a number of studies of collective problem solving, groups seem to have an advantage compared to individuals working alone (Bahrami et al., 2010; Wahn et al., 2020; Woolley et al., 2015). However, there are also studies suggesting that divergent thinking processes can be inhibited by social interaction (Basden et al., 2000; Kohn & Smith, 2011; Mullen et al., 1991). As group members communicate about their creative ideas, they can come to bias or even disrupt the search process of the fellow group members, affecting the fluency of the group.

It is also unclear how interaction affects flexibility and originality. Recent studies suggest that social interaction can stimulate processes of cognitive flexibility. Groups will often consist of multiple individuals, each with their perspectives, cognitive styles, and prior experiences. This might lead the group to represent more different intuitions and therefore perform broader cognitive search (Tylén et al., 2020; Wahn et al., 2020). Second, the dialogical sharing of ideas might work to cue individual group members to overcome cognitive fixation, and visit solutions neither group members would otherwise think of (Tylén et al., 2014).

A corollary of this reasoning is that groups will benefit from differences between group members (Aggarwal & Woolley, 2010, 2019; Hong & Page, 2004; Sulik et al., 2021). If group members are highly aligned in their cognitive strategies and/or intuitions, they will tend to be attracted to similar parts of the solution space and will thus have a more limited potential to positively complement each other. On the other hand, if group members differ in their cognitive styles, perspectives and strategies, they have a greater potential to combine efforts and exert an influence on each other's cognitive processes (Fjaellingsdal et al., 2021).

Distributional semantics and creativity

Investigating the architecture of human representational space and the particular ways we search this space in contexts of problem solving or creative practices is extremely challenging. Many contemporary theoretical models start from the assumption that human memory has the structure of an associative network comprising links of varying strength between nodes of meaning (Anderson, 1983; Collins & Loftus, 1975). When presented with a prompt (e.g., a word or picture) a node in the network is activated, and depending on the strength of associations to other nodes, the activation can spread to related nodes (Kenett et al., 2017).

Connections and their relative strengths are formed through experience (Flusberg & McClelland, 2014). If two objects or concepts often occur together in experience, their association strength is enhanced. Since individuals have different experiences, their network of associations might differ contingent on an individual's embeddedness in particular environmental, cultural, professional, and socio-demographic contexts (Hoffman, 2018). The association of the word "rat" to other words/concepts might differ between an individual who holds a rat as a dear pet and someone who works in a kitchen and thus considers them a pest.

While it can be hard to map the individual associative networks making up our semantic memories, word embedding models from *Natural Language Processing* (NLP) seem to produce reasonable approximations (Hashimoto et al., 2016). Trained on large text corpora, models learn representations of the meaning of words from their co-occurrence patterns in natural language use. The resulting embeddings represent individual words as vectors of values in large multidimensional spaces (often 300+ dimensions) where semantic similarities between two words are quantified as the cosine distance between two such vectors in semantic space (Jatnika et al., 2019). As these models make it possible to represent human semantic memory in terms of a "searchable" space, they have recently been introduced in creativity research as a means to quantify aspects of cognitive and semantic search (Beaty & Johnson, 2021; Olson et al., 2021; Orwig et al., 2021).

The present study

This study combines word embedding models and agent-based modeling (ABM) to investigate cognitive search under various conditions. ABM provides a privileged window into the underlying causal mechanics of human cognition and interaction. By explicitly modeling agents with theoretically motivated cognitive inclinations, ABM makes it possible to test hypotheses concerning their emergent behavior in social contexts as a function of a range of manipulable parameters.

In this simulation, we investigate the effect of social interaction and cognitive diversity on processes of divergent thinking. We equip individual agents with each their own "semantic memory", that is, a 400-dimensional word2vec model including word vectors for animal names in English. Agents are presented with a simple association task: they have to list as many animal names as possible before reaching cognitive fixation (i.e., a situation where they experience no immediate association to another animal). For each simulation, performance in the task is measured in terms of fluency (the number of animals named in total), flexibility (the average distance between consecutive pairs of animals in a game), and originality (how "rarely" animals are named by agents). The task is carried out in two conditions; i) an individual condition, where agents perform the task alone; and ii) a collaborative condition, where agents are paired two-and-two and perform the task together by taking turns. To investigate the effect of diversity on collective search, cognitive diversity of agents is manipulated by adding controlled levels of noise to agents' semantic memories.

Based on the literature reviewed in previous sections, we hypothesize that pairs of agents will outperform the better individual pair member, but only to the extent that they differ in their cognitive organization. In other words, we predict that diversity will have a positive impact on divergent thinking, with higher-level diversity groups outperforming lower-level diversity groups and individuals with respect to fluency, flexibility and originality of solutions.

Methods and Materials

Word embeddings

Artificial agents were each equipped with a semantic memory constrained to a single domain, in this case that of animals. The semantic memory was constructed by taking a list of 240 animals (sampled from https://a-z-animals.com/), training a skip-gram word2vec model (Mikolov et al., 2013) on a full dump of the English Wikipedia using the Python package Gensim (Rehurek & Sojka, 2011), and extracting the resulting 400-dimensional word embeddings for each animal. Notice that, while helping interpretability, the choice of animals as target domain is entirely agnostic to the mechanism of the agent-based simulation itself. Other domains or even artificially generated data could have been used instead.

Cognitive diversity To simulate human cognitive diversity, we generated a number of agent populations with varying levels of internal semantic diversity.

First, we extracted 20 equally spaced values (henceforth *distance thresholds*) in the range of pairwise distances between animals in the original semantic space. Each of these values v_i was used to generate a corresponding agent population, P_i as follows. First, we extracted the list L containing all pairs of animals in the semantic space whose distance is *lower* than the distance threshold for that population, v_i. To instantiate an agent A_k in the population P_i, we performed the following steps:

1. we randomly shuffled L, generating a new list L_k that differs from the original only in the *order* of pairs;
2. we went through the list, from the first to the last pair. For each pair of animals we encounter, we *swap* the position of those animals in the word2vec space;
3. the resulting word2vec matrix is used as the semantic memory of the new agent.

We repeat this process 100 times for each distance threshold, resulting in 20 populations with 100 agents each. To perform the simulation in the *interactive* condition, we randomly sample 100 pairs of agents from each population.

The extent to which the semantic memory of an agent in a given population P_i will differ, on average, from those of other agents in the *same* population increases as a function of v_i. For lower distance thresholds, swapping pairs in word2vec space only induces "local" diversity, as only close-by animals are swapped. For higher distance thresholds, not only will *more* animal pairs be swapped, but swapping of animals which are *further away* in space will be allowed, which induces diversity not only in local neighborhoods, but also in the global semantic structure of the space. We refer to these increasing levels of within-population diversity as *diversity levels*.

The advantage of inducing diversity by swapping animals in space, compared to the perhaps more cognitively plausible strategy of adding increasing amounts of random noise to the word2vec matrix, is that this allows us systematically manipulate semantic diversity while entirely preserving the topology of the original space. As a consequence, agents will have the exact same levels of performance in the individual condition, making any observed differences in performance metrics in the *interactive* condition uniquely dependent on the effect of interaction and diversity.

As the distribution of pairwise Euclidean distances between agents for each diversity level changes minimally from distance level 12 onwards, we only report results for diversity levels 1-12.

Simulation mechanics

In the simulation, each agent performs a word association task in two conditions, alone (the *individual condition*) or by taking turns with another agent (the *interactive condition*). In both conditions, each simulation is initiated by seeding one out of the 240 animals in the agents' semantic memory.

In the individual condition, the agent responds to the initial seed (say, "dog") by naming the animal which is closest to the seed in semantic space (say, "cat"). The animal named in response to the seed ("cat") now becomes cue for the following turn, to which the agent follows up, again, by naming the animal closest to the cue in its semantic space (say, "mouse"). All response animals (plus the initial seed) are dropped from the agent's semantic memory as they are named, so that each animal can only be named once within one association chain. This procedure goes on iteratively, until at least one of the following criteria is met: a) the agent has named all animals; b) no animal has a distance from the cue which is lower than a set threshold – an analogy to cognitive fixation, where no plausible semantic association comes to mind.

Based on pilot simulations, we selected a threshold which yields individual performance values centered around approximately half the number of animals in the semantic space, thus avoiding floor or ceiling effects and leaving room for observing manipulation-induced variation.

In the interactive condition, the mechanics of the game are the same with the only exception that agents take turns naming animals. That is, in each pair, agents are randomly assigned the roles of *Agent 1* and *Agent 2*. At the beginning of each association chain, Agent 1 is prompted with an initial seed word (say, "dog"), to which it responds with the animal that is closest to the seed in *its own* semantic space (say, "cat"). Agent 2 is then passed "cat" as a cue, to which it follows up by naming the animal that is closest to that in *its own* semantic space (say, "tiger"), and so on iteratively, until a stopping criterion is met. Animals are dropped from both agents' semantic memories as they are named.

In both conditions, for each agent/pair, the association game is simulated 240 times, each time prompting the agent/pair with a different initial seed among the animals defined in their semantic memory.

Performance for each agent/pair in each trial is measured along three metrics: 1) *fluency*, defined by the number of animals named in a trial; 2) *flexibility*, defined as the mean distance between each pair of consecutively named animals. In the interactive condition, for each response given by Agent 1, distances are calculated relative to the semantic spaces of Agent 1, and vice versa. Lastly, for each animal named in a given trial, we compute its *originality* as the proportion of trials (out of all simulations from the individual condition) in which the animal is *not* named. The 3) *originality* of a given trial is the average originality of all animals named in the trial.

To assess the effect of social interaction and diversity, we compute the collective 'benefit' relative to the performance of individual pair members. For fluency and flexibility, this is operationalized as the percentage increase in fluency and flexibility from the individual condition. For originality, this is computed as the percentage increase in originality compared to the most original pair member.

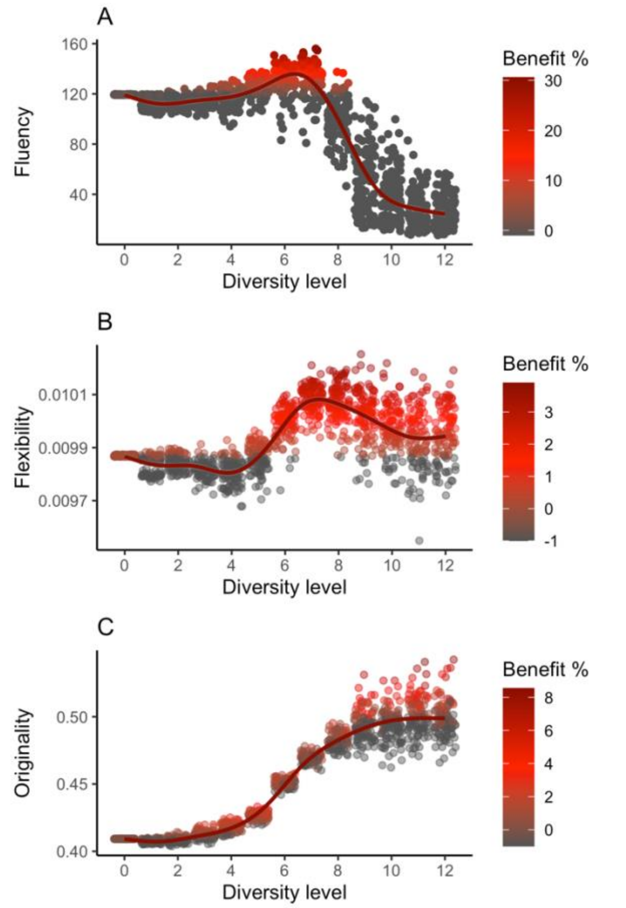

Figure 1: Effect of cognitive diversity on collective divergent thinking. **A**: Fluency (number of animals listed). Benefit is percentage increase in fluency relative to the more fluent pair member. **B**: Flexibility (mean semantic distance traveled in speaker's semantic space). Benefit is percentage increase in flexibility relative to the more flexible pair member. **C**: Originality (inverse frequency of solutions). Benefit is percentage increase in originality relative to the most original pair member.

Results

Fluency

We found moderate positive effects of social interaction on fluency modulated by the level of diversity. The fluency of pairs compared to fluency in the individual condition seems to gradually increase with increasing diversity until diversity level 7, after which it drops drastically (see Figure 1A). Note that this is possibly due to the fact that for the highest diversity levels, a pair member has higher probability of listing animals for which the partner has no above threshold association to other animals which leads the trial to end.

Flexibility

The average distance traveled in semantic space is also affected by cognitive diversity. We observe more exploratory behavior up until diversity level 7, after which flexibility decreases again (see Figure 1B and Table 1). Increased diversity thus gives rise to more exploratory search behaviors, which possibly allows pairs to also reach less dense or more peripheral regions of the semantic space (see Figure 2). This may yield performance advantages by avoiding getting stuck in overexploitation of a local minima, but it can also have negative effects when agents leave an area prematurely before having exploited the local neighborhood. In addition, entering peripheral regions of the semantic space comes with the risk of getting stuck with no above threshold association to the next entry, which might explain why flexibility is decoupled from fluency at the highest levels of diversity.

Originality

Last, but not least, the originality of is also mediated by social interaction and diversity, with more original solutions being moderately correlated with diversity levels (see Figure 1C).

Notice, however, that originality also increase as a function of diversity in the individual condition. The collective benefit is again most notable at intermediate levels of diversity (3-6), after which patterns becomes more heterogeneous.

Table 1: Examples of the first 20 entries from sample association chains produced by individual agents and pairs from the seed "goat", at diversity level 2 and 7, respectively.

Diversity level 2	
Agent 1	***goat***, *sheep, cow, chicken, reindeer, grizzly bear, polar bear, black bear, tapir, caiman, sea snake, coral snake, sawfish, squid, cuttlefish, octopus, shark, whale, fin whale, sperm whale, dolphin, …*
Agent 2	***goat***, *chicken, pig, sheep, cow, reindeer, grizzly bear, polar bear, black bear, tapir, caiman, sea snake, coral snake, sawfish, squid, cuttlefish, octopus, shark, blue whale, fin whale, dolphin, …*
Pair	***goat***, *sheep, pig, duck, goose, pheasant, deer, elk, antelope, bison, black bear, polar bear, grizzly bear, blue whale, fin whale, sperm whale, killer whale, dolphin, whale, shark, cuttlefish, …*
Diversity level 7	
Agent 1	***goat***, *lobster, ladybug, weasel, mongoose, pelican, rattlesnake, red admiral, lynx, mammoth, nuthatch, prairie dog, termite, silkworm, louse, halibut, dolphin, platypus, jaguar, tiger, bear, …*
Agent 2	***goat***, *killer whale, harvestman, rooster, spider, anteater, clown fish, cockroach, alligator, lizard, sperm whale, moose, pheasant, swallowtail, tiger, centipede, crocodile, albatross, saurian, baboon, …*
Pair	***goat***, *lobster, chameleon, blackbird, bear, pekingese, dragonfly, fly, cuckoo, sawfish, dingo, gazelle, camel, antelope, stingray, moose, pheasant, dolphin, killer whale, chicken, crab, …*

Discussion

In an agent-based simulation, we contrasted divergent thinking dynamics in individual agents and pairs of interacting agents using a semantic association task. Besides, we manipulated the relative cognitive diversity of pair members to investigate the impact of diversity on unfolding divergent thinking processes.

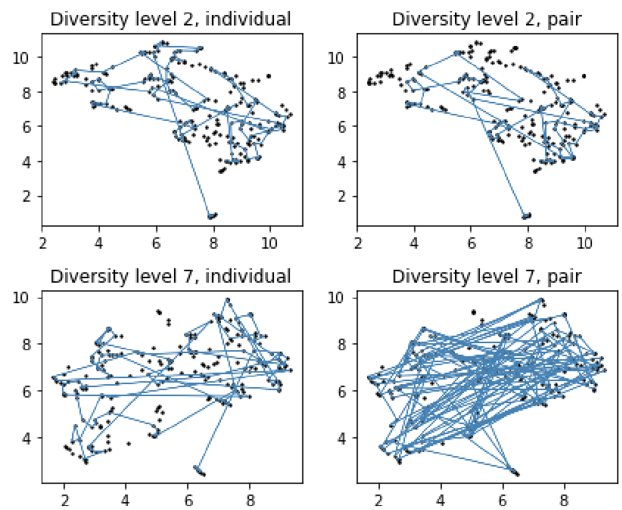

Figure 2: Search trajectories in sample trials from a pair (right panels) and its best individual (left panels) sampled from diversity level 2 and 7. The semantic spaces are a two-dimensional reduction of the semantic memories of the individual (computed using UMAP).

We find that pairs outperform the better performing member in terms of fluency (i.e., proving as many candidate solutions to the prompt as possible), but only at a particular range of diversity. The fluency effects seem to be - at least partly - explained by the flexibility by which agent pairs explore their semantic spaces. At each trial, agents' name the animal closest to the previous animal in their semantic space. When agents are similar in their cognitive organization, their search paths through semantic space will show a high degree of overlap, which prevents them from gaining higher performance than the better pair member. In particular, if an agent only has weak associative connections to a particular area of semantic space, the cognitively similar partner is likely also not to bring the pair to this area.

With increasing levels of diversity, pair members' search strategies come to complement each other. Animals that are barely related in the semantic memory of one pair member, are potentially closely connected in the mind of the other pair member, which allows the pair to escape fixation and together explore more distributed corners of the solution space. This finds expression in the increased flexibility of the pair, here measured as the distance traveled in semantic space. However, importantly, we also observe that the effect of diversity changes with higher levels of diversity. When pair members' semantic spaces are too different (diversity level > 7), the effects on fluency reverses and performance breaks down. This seems to happen when the semantic inclinations of one pair member tend to bring the pair to places for which the partner has no above-threshold associations.

The relative originality of solutions also seems to be moderately affected by cognitive diversity, with diverse pairs listing on average more original solutions relative to the more original pair member. However, with high levels of diversity, the originality effects attenuate and become very varied, possibly due to the fact that many of these trials tend to be rather short.

The current observations provide a privileged window into the complex cognitive dynamics and mechanisms of collaborative creative processes. Importantly, social interaction is not a one-size-fits-all. Rather, the impact of social interaction on the unfolding of collective search is modulated by the relative cognitive diversity of the agents. In particular, diversity seems to affect the flexibility by which agents navigate their solution spaces.

Cognitively similar agents seem to have less potential to affect each other's search patterns, which leads to collective behaviors characterized by more exploitation of local semantic neighborhoods. Their collective performance (fluency) will thus often coincide with - or even be lower than - the better performing pair member, as they experience cognitive fixation at similar areas of the semantic space. In addition, the fact that they are attracted to the same region of semantic space seems to affect the relative originality of their solutions, again with the pair not providing more original solutions than the better performing pair member.

In contrast, with increasing diversity, agents impact each other's search patterns giving rise to emergent patterns of more explorative search (longer jumps in semantic space). This seems to bring the pair to also visit more "distant" regions of semantic space (without experiencing fixation), which, in turn, gives rise to more fluent and original responses.

Importantly, we observe the effect of diversity on creativity is not linear. With the highest levels of diversity, performance breaks down. While a balanced mix of exploration and exploitation seems productive, a too exploratory behavior has the detrimental effect that the pair will often leave a fruitful part of the space (densely inhabited by animals) prematurely and jump to a different part of space. In addition, high diversity seems often to bring a pair member to a corner of the semantic space for which associations are weak, which leads to fixation and ends the trial.

While quadratic effects of diversity following a similar dynamics are previously reported in the literature (e.g. Aggarwal et al., 2015), it is an open question whether the higher levels of diversity operationalized in this simulation give rise to plausible behavior. High levels of diversity are effectively equivalent to scenarios where agents' semantic memories are entirely unrelated and retain little similarity with the original semantic model. When consulting examples of animal associations produced by the pair from diversity level 7 (see Table 1), they are already quite idiosyncratic, moving from "goat" ⇒ "lobster" ⇒ "chameleon" ⇒ "blackbird". We cannot exclude the possibility that the quadratic patterns giving rise to detrimental effects at the highest levels of diversity (level 8-12) could be an artifact of implausibly extreme parameters settings. If so, diversity might be positively linearly related to performance in contexts of divergent thinking. Future experimental studies with human participants could inform a more motivated choice of parameter values for diversity.

In the present implementation of the ABM, agents' associations are based on semantic relations. However, the increasing idiosyncratic associations resulting from shuffling words in the agents' memories are intended to represent individual variability formed through episodic experience (Denervaud et al., 2021). But human agents are likely to deploy more complex heuristics. For instance, similarities in the sound of a word (e.g., "horse" ⇒ "sea horse") or the visual similarity between two otherwise unconnected animals (e.g., "mosquito" ⇒ "narwhal") can motivate a connection despite their weak semantic or episodic association. Future implementations of the ABM should incorporate this complexity to a wider extent.

In the association game performed by agents in this study, interactions are constrained to a strict turn-taking scheme, which can lead to a trial ending when the turn-holding agent does not have a sufficiently strong association to a new animal. A future implementation could relax this constraint and allow the partner to chip in when the turn-holder experiences fixation. This is likely to create greater benefits for pairs with high cognitive diversity (as they often differ in terms of which areas of semantic space are well connected), and might be a better model of real-life collaborative creative contexts where interaction is not subject to strict turn-taking constraints (Holler et al., 2016). Stopping rules could also be relaxed in other ways, for example by allowing agents to access distance values for animals named *before* the last turn. Future studies should test the generalizability of our conclusions across different sets of assumptions, search strategies and heuristics.

We have opted for a conservative measure of collective benefit, where benefit is computed relative to the best individual, but there are other relevant comparisons to make. Previous experimental research on divergent thinking has shown that social interaction can have an inhibiting effect when compared to the offline concatenation of individual contributions (Kohn & Smith, 2011; Mullen et al., 1991; Szary & Dale, 2014). Even if pairs generally perform better than the best individual pair member, they might still constrain each other in ways that prevent them from realizing their full potential, as they will also at times disrupt each other's association chains. The current ABM makes it possible also to test relative collaborative inhibition effects and how they are attenuated by cognitive diversity.

With this study, we investigated the dynamics of collective creative processes, and in particular how aspects of social interaction modulate divergent thinking. We observe that social interaction might not always benefit divergent thinking. If collaborative agents are too similar in their cognitive makeup, they contribute redundant behaviors and thus do not benefit from collaborating. Similarly, if they are too different, their associative inclinations cause too disparate and jumpy search patterns with detrimental effects. However, with moderate levels of cognitive diversity, increased cognitive flexibility emerges which helps agents escape cognitive fixation and leads to advantages in fluency and originality of joint solutions.

References

Aggarwal, I., & Woolley, A. W. (2010). *Do you see what i see? The effect of members' cognitive styles on team processes and performance*. 70th Meeting of the Academy of Management, Montreal, Canada.

Aggarwal, I., & Woolley, A. W. (2019). Team creativity, cognition, and cognitive style diversity. *Management Science*, *65*(4), 1586–1599.

Aggarwal, I., Woolley, A. W., Chabris, C. F., & Malone, T. W. (2015). Cognitive diversity, collective intelligence, and learning in teams. *Proceedings of Collective Intelligence 2015*.

Anderson, J. R. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *22*(3), 261–295.

Baer, J. (2014). *Creativity and divergent thinking: A task-specific approach*. Psychology Press.

Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally interacting minds. *Science*, *329*(5995), 1081–1085. https://doi.org/10.1126/science.1185718

Baronchelli, A., & Radicchi, F. (2013). Lévy flights in human behavior and cognition. *Chaos, Solitons & Fractals*, *56*, 101–105.

Basden, B. H., Basden, D. R., & Henry, S. (2000). Costs and benefits of collaborative remembering. *Applied Cognitive Psychology: The Official Journal of the Society for Applied Research in Memory and Cognition*, *14*(6), 497–507.

Beaty, R. E., & Johnson, D. R. (2021). Automating creativity assessment with SemDis: An open platform for computing semantic distance. *Behavior Research Methods*, *53*(2), 757–780.

Brophy, D. R. (1998). Understanding, measuring, enhancing collective creative problem-solving efforts. *Creativity Research Journal*, *11*(3), 199–229.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, *82*(6), 407.

Denervaud, S., Christensen, A. P., Kenett, Y., & Beaty, R. E. (2021). Education shapes the structure of semantic memory and impacts creative thinking. *Npj Science of Learning*, *6*(1), 1–7.

Fjaellingsdal, T. G., Vesper, C., Fusaroli, R., & Tylén, K. (2021). *Diversity promotes abstraction and cognitive flexibility in collective problem solving*.

Flusberg, S. J., & McClelland, J. L. (2014). Connectionism and the emergence of mind. *The Oxford Handbook of Cognitive Science*, *1*.

Gilhooly, K. J., Fioratou, E., Anthony, S. H., & Wynn, V. (2007). Divergent thinking: Strategies and executive involvement in generating novel uses for familiar objects. *British Journal of Psychology*, *98*(4), 611–625.

Hart, Y., Mayo, A. E., Mayo, R., Rozenkrantz, L., Tendler, A., Alon, U., & Noy, L. (2017). Creative foraging: An experimental paradigm for studying exploration and discovery. *PloS One*, *12*(8).

Hashimoto, T. B., Alvarez-Melis, D., & Jaakkola, T. S. (2016). Word embeddings as metric recovery in semantic spaces. *Transactions of the Association for Computational Linguistics*, *4*, 273–286.

Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., & Cognitive Search Research Group. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, *19*(1), 46–54.

Hoffman, P. (2018). An individual differences approach to semantic cognition: Divergent effects of age on representation, retrieval and selection. *Scientific Reports*, *8*(1), 1–13.

Holler, J., Kendrick, K. H., Casillas, M., & Levinson, S. C. (2016). Turn-taking in human communicative interaction. *Frontiers in Psychology*.

Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *PNAS*, *101*(46), 16385–16389.

Jatnika, D., Bijaksana, M. A., & Suryani, A. A. (2019). Word2vec model analysis for semantic similarities in english words. *Procedia Computer Science*, *157*, 160–167.

Kaufman, J. C., & Sternberg, R. J. (2010). *The Cambridge handbook of creativity*. Cambridge University Press.

Kenett, Y. N., Levi, E., Anaki, D., & Faust, M. (2017). The semantic distance task: Quantifying semantic distance with semantic network path length. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *43*(9), 1470.

Kohn, N. W., & Smith, S. M. (2011). Collaborative fixation: Effects of others' ideas on brainstorming. *Applied Cognitive Psychology*, *25*(3), 359–371.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, *26*.

Mullen, B., Johnson, C., & Salas, E. (1991). Productivity loss in brainstorming groups: A meta-analytic integration. *Basic and Applied Social Psychology*, *12*(1), 3–23.

Newell, A., & Simon, H. A. (1972). *Human problem solving* (Vol. 104). Prentice-Hall Englewood Cliffs, NJ.

Olson, J. A., Nahas, J., Chmoulevitch, D., Cropper, S. J., & Webb, M. E. (2021). Naming unrelated words predicts creativity. *Proceedings of the National Academy of Sciences*, *118*(25).

Orwig, W., Diez, I., Vannini, P., Beaty, R., & Sepulcre, J. (2021). Creative connections: Computational semantic distance captures individual creativity and resting-state functional connectivity. *Journal of Cognitive Neuroscience*, *33*(3), 499–509.

Rehurek, R., & Sojka, P. (2011). Gensim–python framework for vector space modelling. *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic*, *3*(2).

Reiter-Palmon, R., Forthmann, B., & Barbot, B. (2019). Scoring divergent thinking tests: A review and systematic framework. *Psychology of Aesthetics, Creativity, and the Arts*, *13*(2), 144.

Runco, M. A. (2010). Divergent thinking, creativity, and ideation. In J. C. Kaufman & R. J. Sternberg (Eds.), *The Cambridge handbook of creativity* (pp. 413–446). Cambridge University Press

Said-Metwaly, S., Van den Noortgate, W., & Kyndt, E. (2017). Approaches to measuring creativity: A systematic literature review. *Creativity. Theories–Research-Applications*, *4*(2), 238–275.

Sulik, J., Bahrami, B., & Deroy, O. (2021). The Diversity Gap: When diversity matters for knowledge. *PsyArxiv*. https://doi.org/10.31234/osf.io/hx7ze

Szary, J., & Dale, R. (2014). Familiarity modulates the dynamics of collaborative inhibition in a trivia game. *Proceedings of the 36th Annual Meeting of the Cognitive Science Society*.

Tylén, K., Fusaroli, R., Bjørndahl, J. S., Rączaszek-Leonardi, J., Østergaard, S., & Stjernfelt, F. (2014). Diagrammatic Reasoning: Abstraction, interaction, and insight. *Pragmatics and Cognition*, *22*(2), 264–283.

Tylén, K., Fusaroli, R., Smith, P., & Arnoldi, J. (2020). The social route to abstraction: Interaction and diversity enhance rule-formation and transfer in a categorization task. *PsyArXiv*. https://doi.org/10.31234/osf.io/qs253

Wahn, B., Czeszumski, A., Labusch, M., Kingstone, A., & König, P. (2020). Dyadic and triadic search: Benefits, costs, and predictors of group performance. *Attention, Perception, & Psychophysics*, 1–19.

Woolley, A. W., Aggarwal, I., & Malone, T. W. (2015). Collective Intelligence and Group Performance. *Current Directions in Psychological Science*, *24*(6), 420–424. https://doi.org/10.1177/0963721415599543