
Creativity in dynamic networks: How divergent thinking is impacted by one's choice of peers

Raiyan Abdul Baten¹, Daryl Bagley², Ashely Tenesaca², Famous Clark²,
James P. Bagrow³, Gourab Ghoshal⁴, Mohammed Ehsan Hoque^{2,*}

¹Department of Electrical and Computer Engineering, University of Rochester, NY, USA

²Department of Computer Science, University of Rochester, NY, USA

³Department of Mathematics & Statistics, University of Vermont, VT, USA

⁴Department of Physics and Astronomy, University of Rochester, NY, USA

*mehoque@cs.rochester.edu

Abstract

Creativity is viewed as one of the most important skills in the context of future-of-work. In this paper, we explore how the dynamic (self-organizing) nature of real social networks impacts the fostering of creative ideas. We run 6 trials (N=288) of a web-based experiment involving divergent ideation tasks. We find that network connections gradually adapt to individual creative performances, as the participants predominantly seek to follow high-performing peers for creative inspirations. We unearth both opportunities and bottlenecks afforded by such self-organization. While exposure to high-performing peers is associated with better creative performances of the followers, we see a counter-effect that choosing to follow the same peers introduces semantic redundancies in the followers' ideas. We formulate an agent-based simulation model that highlights the implications of these processes. Our findings may help design large-scale interventions to improve the creative aptitude of people interacting in a social network.

Keywords— Network science, computational social science, creativity, dynamic social networks

Recent advances in robotics, AI and machine learning are increasingly focused on mimicking or even surpassing human capabilities. This particular innovation, however, has serious implications on our future workforce [1, 2, 3]. Approximately 51% of the tasks done in the US economy can be automated [4], and for each robot on the factory floor, some six jobs are lost [5]. The need for manual labor in predictable and repetitive work is declining [6], while the demand is soaring for expertise in creative tasks, problem-solving, and other social-cognitive avenues of soft-skills [4, 7, 8]. Many of the critical and challenging tasks of the human civilization requires humans to collaborate with others [9, 10], and perform creatively at both individual and collective levels. Thus, enhancing the creative abilities of collaborating humans has become one of the aspirational challenges today. This motivation for creativity-at-scale leads to the exploration of social networks of creative collaborators. For instance, the development of an innovative product such as an aircraft or a computer operating system is only made possible by an interacting network of creative problem solvers, who benefit from each other's expertise [11]. Discussions in an academic network of researchers, faculties and students can stimulate ideas for novel explorations in the members of the network. A graphic designer can benefit from inspirations found in online networks like Reddit, Behance or Twitter, among other instances of social network driven creative accomplishments. Adopting a social network lens helps understand the mechanisms, bottlenecks and opportunities for maximizing creative outcomes in a larger scale.

Researchers have examined the effects of various network attributes on creativity [12, 13]. For instance, it has been reported that relationship strength, position and external ties are some of the factors that influence creative performance [14]. However, a key element missing from most prior literature is the dynamic characteristic of real social networks. Human interactions are structured in social networks, where people have control over who they interact with. Given an objective, they can choose to make or break ties to update the connectivity patterns around them, often in response to behavior, performance, prestige, age, gender, popularity, self-similarity and other cues of the social partners [15, 16, 17, 18, 19, 20]. This dynamic characteristic affords opportunities in human populations that static networks cannot: for example, dynamic networks promote cooperation among humans [21], improve their collective intelligence [22, 23], and even help build up people’s speaking skills [24].

When it comes to creative performance, the dynamic nature of social networks has received rather little research attention. Perry-Smith *et al.* proposed a spiraling model, capturing the cyclical relationship between creativity and network position, where one fosters the other [12]. The argument being, if someone is creative at something, it might draw more attention to the creative person, resulting in an increased centrality and visibility. Conversely, a central person is able to inspire creative thoughts in others and can also get inspired by others more readily than a peripheral person, thus helping in further creative ideation. However, this chain of argument has not been directly tested. Despite some efforts in examining other temporal effects [25, 26, 27], it remains largely unclear how dynamic creative networks evolve with time, what laws they follow, and what implications such evolving has on the creative ideation process and outcomes.

This motivates the desire to understand how creative performances are exhibited and impacted in dynamic social networks. Consequently, in this paper, we first explore how connectivity patterns adapt to individual performance cues in a creativity-centric dynamic network (RQ1). Second, we test how such adaptations impact individual and collective ideation performances in a dynamic network, against the controls of static (fixed connections) and solo (unconnected) network conditions (RQ2).

We run six trials of a web-based experiment, where participants in a dynamic social network performed creative idea generation tasks in 5 consecutive rounds. The participants chose after each round which of the peers’ ideas they wanted to be shown as stimuli (see Experimental setup for details). Following Perry-Smith *et al.*’s argumentation [12], we anticipated in RQ1 that people looking for creative inspiration from others will use some form of success cues to determine who among the peers are more creative and, therefore, more promising to be advantageous to form ties with (“follow”). Our results show that in the dynamic networks, the participants were indeed drawn towards following the most creative ideators. In particular, we find that the statistical rarity and novelty ratings of one’s ideas to be robust predictors of his/her popularity in the dynamic networks.

If more people choose to form ties with the highly creative ideators in the network, what implications will it have on the individual and collective creative performances? The associative theory of creative cognition suggests that when people are primed properly, e.g., by exposure to ideas from others [28, 29], their long-term memory circuitry can be stimulated. This can enable retrieving concepts that are remotely stored in the memory [30, 31, 32, 33, 34, 35]. Combining various aspects of such remote concepts can help in synthesizing novel ideas [36, 37, 38]. Based on this, we anticipated that following highly creative peers in a dynamic social network can allow for positive stimulation of novel ideas in people. For instance, following a person who generates rare or unique ideas will increase the chances that the follower comes across ideas that have little overlap with his/her own. This can allow for the stimulation of further ideas through novel association of concepts, resulting in ideas that would not have occurred to the follower otherwise [39, 40, 41, 42]. However, we also anticipated a counter-effect that if many people follow the same highly creative ideators in a dynamic network, the followers’ inspiration sets will become overlapping, which might introduce redundancy in the stimulated ideas [43, 44]. Our results show that following highly creative ideators is indeed associated with one’s better creative performance. However, different participants who chose to follow the same alters (same stimuli) show an increasing semantic similarity in their independently stimulated ideas with time. These results suggest that self-organizing in a dynamic network (i.e., preferentially forming ties with the most creative peers) might lead to conflicting opportunities and constraints. In the end, we do not find any significant difference in the individual and collective creative performances between the dynamic and static conditions. We formulate a simulation model that captures these empirically-derived intuitions, and helps us assess the generality of the reported processes and insights.




Figure 1: (A) The bipartite network structure used in the static condition, and as the initial condition in the dynamic condition. The ideas of the 6 alters were pre-recorded and later shown to 18 egos each in static and dynamic conditions. Each ego was connected to 2 alters. (B) The study protocol for each of the 5 rounds. In turn 1, the participants generated alternate use ideas on their own on a given prompt object. In turn 2, egos in the dynamic and static conditions received social exposure from their alters and could add further ideas to their lists. The solo condition did not have any stimuli. Egos in the dynamic condition could update their set of 2 alters at the end of each round, which the participants in other conditions could not.

Experimental setup. It is challenging to identify or generate a dataset in the wild that allows for traceable links between ideas and their stimuli, and also provides temporal evolution information of a dynamic network. We therefore resort to an artificial social network created in the virtual laboratory.

In this experiment, we are interested in divergent creativity, which deals with a person’s ability to come up with or explore many possible solutions to a given problem [45]. We use a customized version of Guilford’s Alternate Uses Test [46], which is a widely-adopted approach for quantifying divergent creative performances. The study had 5 rounds. In each round, the participants were instructed to consider an everyday object (e.g., a brick), whose common use was stated (e.g., a brick is used for building). The participants needed to come up with alternative uses for the object: uses that are different than the given use, yet are appropriate and feasible. We choose 5 common objects from Form B of Guilford’s test as the ideation objects in the 5 rounds.

We recruited 288 participants from Amazon Mechanical Turk, who are diverse in their age, racial, ethnic and gender distributions, and live in the United States (see SI for details). We placed them randomly in one of three network conditions: (1) Dynamic, (2) Static, and (3) Solo. The static and solo conditions act as controls against which we assess the performances in the dynamic condition. For the dynamic and static conditions, we adopted a bipartite network structure [47], as shown in Figure 1A. There are two types of nodes in the network, *alters* and *egos*. First, we pre-recorded the ideas of 6 *alters*, who generated ideas on their own. Then, we used these ideas as the stimuli for 36 *egos*—18 of them placed in a dynamic network condition, and the other 18 in static. This bipartite design helped us ensure a uniform stimuli-set for all the egos in the static and dynamic conditions. We repeated the process for 6 independent trials, each with its unique set of alters. Under the solo condition, 36 additional participants generated ideas in isolation.

Initially, the egos in the dynamic and static conditions were connected to 2 alters each using the network structure shown in Figure 1A. In each round, the egos in both of these conditions first generated ideas on their own for 3 minutes (‘turn 1’). They were then shown the ideas of the 2 alters they were connected to, and given an additional 3 minutes to list further ideas (‘turn 2’). The egos were told not to resubmit any of the alters’ exact ideas, and that only non-redundant ideas would contribute to their performance. They were also told that there will be a short test at the end of the study, where they will need to recall the ideas shown to them. This was to ensure that the participants paid attention to the stimuli ideas, which has been shown to positively impact ideation performances [37, 33, 41, 48]. Then, the egos rated the ideas of all of the 6 alters on novelty (5-point ratings, 1: not novel, 5: highly novel) [49, 50]. Finally, in the dynamic condition, the egos were allowed to update their links at the end of the round, by optionally following/unfollowing alters to have an updated list of 2 alters each. Except for the alters’ username and ideas, no other information about the alters was shown to the egos. The egos were required to submit the rationale behind their choices of updating/not updating links in each round. This was in place to make the dynamic egos accountable for their choices, which has been shown to raise epistemic motivation and improve

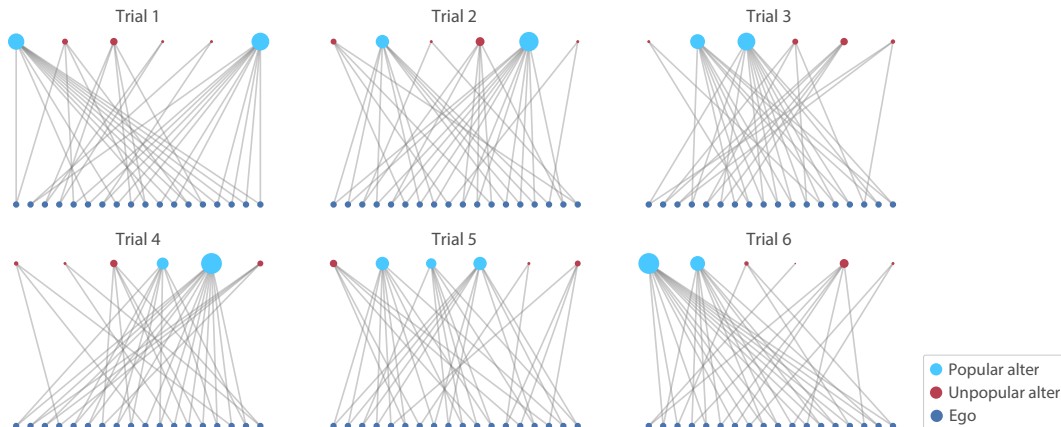


Figure 2: Final (evolved) dynamic networks in the 6 trials. The diameters of the circles in the upper-rows (alters) are proportional to the alters’ number of followers at the end of the 5th round. Some of the alters finished with more than the initial number of 6 ego-followers (‘popular alters’, marked in cyan), while some finished with ≤ 6 ego-followers (‘unpopular alters’, marked in maroon).

systematic information processing [50, 51]. The egos in the static condition could not update their links, marking the only difference between the two study conditions of static and dynamic.

The participants in the solo condition were given 6 minutes to list their ideas without any external stimuli. Detailed descriptions and examples guided the participants throughout the study. Everyone was paid \$10 after completing the requirements, and the top 5 egos/solo participants (in each group of 18) with the most number of non-redundant ideas were awarded \$5 bonuses. Figure 1B summarizes the protocol.

Measures. We operationalize creative performance using three metrics: (1) Number of non-redundant ideas, (2) Average novelty ratings and (3) Creativity quotient. We briefly introduce the metrics below, and refer to the Methods section for details. The non-redundant idea count is a measure of statistical rarity of the ideas. If an idea is given by at most a threshold number of participants in a given participant pool, it is considered non-redundant. The non-redundant idea count is then incremented for all of the participants who submitted that idea [52, 53].

Each idea of the alters was rated on novelty by 36 egos in the static and dynamic conditions in the corresponding trial. Additionally, each idea of the egos and solo participants was rated by at least 4 raters, who were hired separately from Amazon Mechanical Turk. We take the mean of the ratings received by each idea as that idea’s novelty indicator, and consider the average novelty rating received by an individual as his/her creativity metric.

Our third metric is the creativity quotient, Q . This metric combines fluency (quantity of ideas) and flexibility (the ability to generate a wide variety of ideas) of the submitted ideas using information theoretic measures [54, 55, 56]. In all three metrics, a higher value indicates a better creative performance.

Results

Link update patterns in the network evolution. In response to RQ1, we first explore the evolution of network connectivity patterns and the associated cues. The dynamic networks evolved as the egos updated their lists of 2 alters across the rounds. The final evolved networks of the 6 trials are shown in Figure 2, where the diameters of the upper-row circles (alters) are drawn in proportion to the alters’ number of followers at the end of the 5th round. All of the alters started with 6 ego-followers, but after the network evolution, some of the alters lost followers, and some of them gained. Figure 3 shows the number of connection updates per ego at the end of each round. As can be seen, the number of connection changes had a downward trend across rounds ($p < 0.001$ for the negative slope), often showing a sharp drop at the end of the second round. Out of a maximum of 2 possible updates, an average of 0.97 connections were updated per ego after the first rounds of the trials,




Figure 3: Number of connections updated per ego at the end of each round. Across rounds, the number of connection updates per ego has a downward trend ($p < 0.001$ for the negative slope).




Figure 4: (A) Cumulative non-redundant idea count comparisons between popular and unpopular alters across rounds. The total number of non-redundant ideas in all 5 rounds is significantly higher for the popular alters than unpopular alters ($p < 0.001$, see footnote 1). (B) The popular alters have significantly higher average novelty ratings than the unpopular alters across all 5 rounds ($p < 0.001$, footnote 2). (C) The total creativity quotient in all 5 rounds is significantly higher for the popular alters than the unpopular alters ($p < 0.001$, footnote 3). *** $p < 0.001$.

while 0.3 connections were updated per ego after the final rounds. This suggests that as the egos received information about the alters' performances through the rounds, they made up their minds on whom to follow, and readjusted later if necessary.

We denote alters who finished with > 6 and ≤ 6 followers as 'popular' and 'unpopular' alters respectively, for ease of reference. The rationale being, if the egos don't update their links at all, the alters will still have the initially assigned 6 followers, so we take > 6 followers as the threshold for defining popularity. The total non-redundant idea counts of the popular alters in all 5 rounds were significantly higher than the counts of the unpopular alters ($p < 0.001$ in the aggregated data over all trials¹, Figure 4A; $p < 0.05$ in each of the 6 trials, SI Figure S1). Similarly, the average novelty ratings of the ideas of the popular alters were significantly higher than those of the unpopular alters ($p < 0.001$ in the aggregated data², Figure 4B; $p < 0.05$ in 5 out of 6 trials, SI Figure S2). When it comes to the creativity quotient, the total Q scores of the popular alters in all 5 rounds were again significantly higher than the unpopular ones ($p < 0.001$ in the aggregated data³, Figure 4C). However, in this case, the differences were significant in 3 of the trials ($p < 0.05$) and insignificant in the other 3 (SI Figure S3).

¹Popular (p) vs unpopular (u) alters: 2-tailed test, $m_p = 23.8$, $m_u = 14.4$, $t(34) = 7.291$, $p < 0.001$, 95% C.I. for $m_p - m_u = [6.9, 12]$

²2-tailed test, $m_p = 3.1$, $m_u = 2.6$, $t(34) = 5.7$, $p < 0.001$, 95% C.I. for $m_p - m_u = [0.3, 0.6]$

³2-tailed test, $m_p = 57.8$, $m_u = 36.7$, $t(34) = 5.81$, $p < 0.001$, 95% C.I. for $m_p - m_u = [13.9, 28.2]$

Table 1: Regression results of predicting the alters’ relative popularity from their relative creativity markers. β = standardized regression coefficient. $**p < 0.01$. $***p < 0.001$.

Predictor	Model 1: \bar{r}' only		Model 2: \bar{r}', u'		Model 3: all predictors	
	β	t	β	t	β	t
\bar{r}'	0.1851***	9.762	0.1278***	4.972	0.1026**	3.435
u'	—	—	0.0767**	2.983	0.0704**	2.763
Q'	—	—	—	—	0.0399	1.562
R^2	0.737		0.793		0.808	
Adjusted R^2	0.729		0.780		0.790	

Using multivariate linear regression analysis, we explore how the creative performances of the alters correspond to their final number of followers. As the dependent variable y_i , we take the fraction of egos connected to an alter i at the end of the 5th round. The three independent variables are: (1) the relative number of non-redundant ideas, $u'_i = \frac{u_i}{\sum_i u_i}$, (2) relative average novelty ratings, $\bar{r}'_i = \frac{\bar{r}_i}{\sum_i \bar{r}_i}$, and (3) relative creativity quotient, $Q'_i = \frac{Q_i}{\sum_i Q_i}$. Here, u_i , \bar{r}_i and Q_i are the total number of non-redundant ideas, average novelty ratings and total creativity quotients of alter i in all 5 rounds together. We take the relative performance of the alters with respect to other alters in a given trial, since the egos could only choose from a fixed pool of alters. While all of the three independent variables correlate strongly with the dependent variable (Pearson’s $\rho = 0.80, 0.86, 0.75$ respectively, $p < 0.001$ in each), multivariate regression takes care of information overlap therein and allows us to explore the relative contributions of the three independent variables. Mathematically,

$$y_i = \beta_0 + \beta_1 u'_i + \beta_2 \bar{r}'_i + \beta_3 Q'_i \quad (1)$$

The results are shown in Table 1. We first test the three independent variables separately using univariate regression, and find \bar{r}' to lead to the best adjusted $R^2_{\bar{r}'} = 0.729$ (shown as Model 1 in the table). Adding u' makes adjusted $R^2_{\bar{r}', u'} = 0.780$ (Model 2), while all the features together give adjusted $R^2_{\bar{r}', u', Q'} = 0.790$ (Model 3). In Model 3, it can be seen that the β values are significant for \bar{r}' and u' , but not for Q' . These p -values are computed against a null hypothesis of no association between the dependent and independent variables, showing that the associations reported here are systematic for the first two predictors (\bar{r}' and u'). Thus, the three independent variables together explain 79% of the variation in the dependent variable, where the first two independent variables contribute the most. This indicates that the egos’ decisions of following alters were more strongly captured by the novelty of the alters’ ideas (\bar{r}') and moderately by the statistical rarity of the ideas (u'). The relative creativity quotient metric, which combines fluency (quantity of ideas) and flexibility (quantity of different categories represented by the ideas), had much of its information overlapped with the other two metrics.

The key take-away here is that the egos in a dynamic network are drawn to form ties with the consistently high-performing alters, as we anticipated. In typical creativity studies, participants are shown inspiration stimuli randomly or based on intrinsic qualities of the stimuli [29, 28]. In our experiment, a key contrast is that the networks are allowed to dynamically self organize—in other words, one can choose for oneself who to take inspirations from. Thus, the implications of such adaptations on the ideation process, as we explore below (RQ2), become direct manifestations of the dynamic nature of the social networks.

Exposure to high-performing alters is associated with better creative performance of the egos.

We argued previously that forming ties with high-performing alters should increase the chances that an ego comes across ideas that have little overlap with his/her own. This lack of overlap, in turn, can increase the chances of stimulating new ideas in the ego by facilitating novel associations between remote concepts. To test this, we take advantage of the fact that in turn 1, the egos generated ideas on their own prior to any social exposure, which allows us to test the overlap their ideas have with their alters’ ideas (measured by idea-set overlap, Jaccard index, see Methods). In turn 2, the egos could see the alters’ ideas, which allows us to explore whether the creativity measures of stimulated ideas have any association with how good the respective alters were.

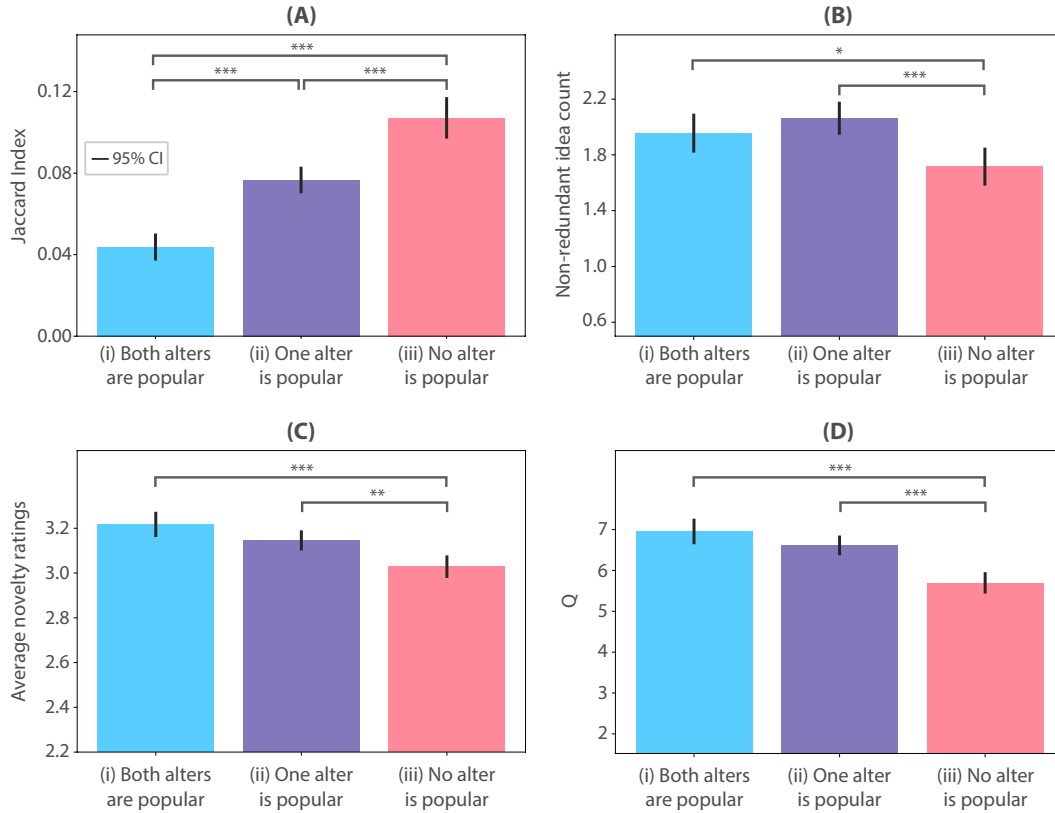


Figure 5: (A) Average overlap (measured with Jaccard Index) between the egos’ turn-1 ideas and their alters’ ideas, aggregated over all rounds and all trials. Comparisons are made among three cases of egos: those with (i) both, (ii) only one and (iii) no alter(s) who are round-wise popular. The overlaps are significantly higher from (i) to (ii), and from (ii) to (iii) ($p < 0.001$ in both cases, see footnotes 4, 5). This shows that following more high-performing alters increases the ego’s exposure to ideas that are different from his/her own. Panels (B), (C) and (D) show the creative performances of these egos in turn-2. In all three metrics, we find that egos in condition (iii) perform significantly worse than both (i) and (ii), but there is no significant difference between (i) and (ii) (footnote 6). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, all p -values Bonferroni-corrected.

We first find out the round-wise popular alters. If an alter has > 6 followers at the end of a round, we take that as a marker that his/her performance was superior in that round, and label the alter as a popular alter for that round. Then, for each round, we split the egos (dynamic and static conditions together) into three groups, where (i) both, (ii) only one, and (iii) none of the followers of an ego are round-wise popular alters. For each ego, we take the average overlap between the ego’s turn-1 ideas and his/her two alters’ ideas. Figure 5A shows that group (i) had significantly less idea-overlap than group (ii) ($p < 0.001^4$), and group (ii) had significantly less overlap than group (iii) ($p < 0.001^5$). This shows that following more high-performing alters systematically decreased the overlap between an ego’s turn-1 ideas and the alters’ ideas, as was anticipated. SI Figure S4 shows that group (i) consistently had less overlap than groups (ii) and (iii) in each of the 5 rounds.

We then compare the creativity metrics of the turn-2 ideas among the same three groups. As shown in Figures 5B, 5C and 5D, both groups (i) and (ii) significantly outperformed the turn-2 ideas

⁴Group (i) vs (ii): 2-tailed test, $m_1 = 0.04$, $m_2 = 0.08$, $t(747) = -6.493$, Bonferroni-corrected $p < 0.001$, 95% C.I. for $m_1 - m_2 = [-0.043, -0.023]$

⁵Group (ii) vs (iii): 2-tailed test, $m_2 = 0.08$, $m_3 = 0.11$, $t(805) = -5.194$, Bonferroni-corrected $p < 0.001$, 95% C.I. for $m_2 - m_3 = [-0.042, -0.019]$

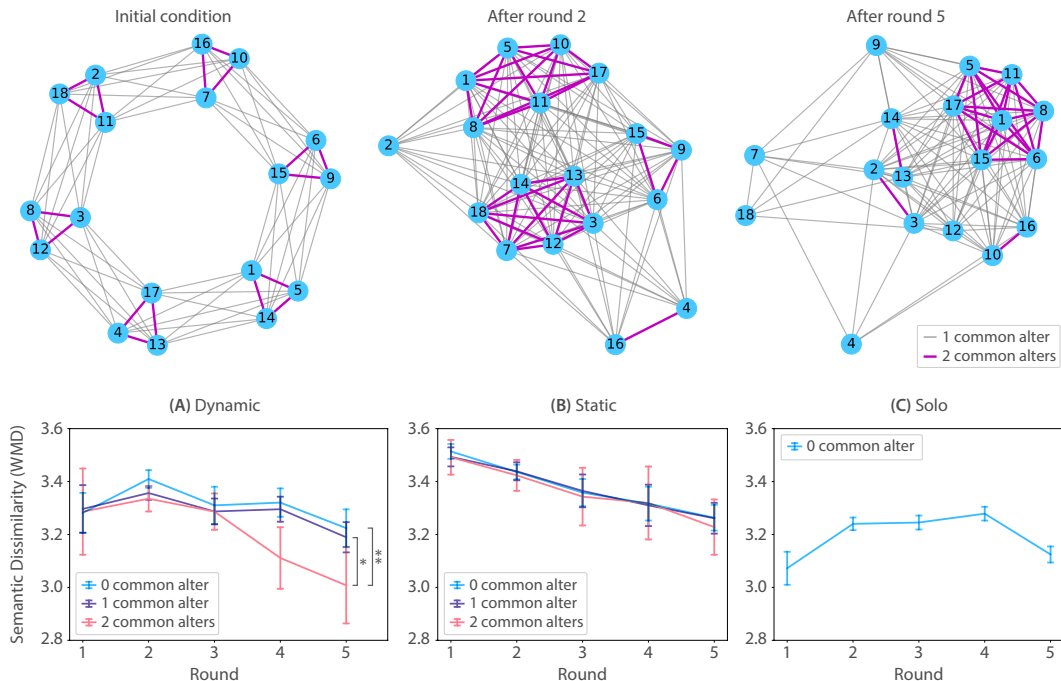


Figure 6: (*Top row*) One-mode projections of the bipartite networks on the ego-nodes. Three such projections from the dynamic condition of trial 4 are shown: the initial projected network, and the projections after rounds 2 and 5. In the static condition, the initial condition remained fixed. (*Bottom row*) Semantic dissimilarity (measured by Word Mover’s Distance, WMD) between the idea-sets of node-pairs are shown for (A) dynamic, (B) static and (C) solo conditions. Node-pairs that share two common alters in the dynamic condition show significantly less dissimilarity by the fifth round than the 0 and 1 common alter cases (footnote 7). The whiskers denote 95% C.I. $*p < 0.05$, $**p < 0.01$, all p -values Bonferroni-corrected.

of group (iii) in all of the three metrics⁶. However, there was no significant difference in any of the performance metrics between groups (i) and (ii) (2-tailed test, $p > 0.05$). This implies that following at least one high-performing alter is associated with better creative performance of the egos. To summarize, these results suggest evidence for better stimulation of ideas when egos are exposed to high-performing alters’ ideas.

Following the same alters introduces semantic similarities in the egos’ ideas. We further motivated a counter-argument that if many egos follow the same high-performing alters, their stimuli set will become overlapping. This might make the egos’ turn-2 ideas similar, despite them ideating independently. To test this, we explore whether the semantic (dis)similarities of pairs of egos have any association with the number of common alters they have. We estimate the semantic nature of submitted ideas using neural word embeddings (Word2Vec [57]) and compare the dissimilarity of the embeddings using Word Mover’s Distance [58] (see Methods).

We first take one-mode projections [47] of the round-wise bipartite networks on the ego nodes. In the projected graphs, two ego-nodes are connected with an edge if they have common alters (top row of Figure 6). For each node-pair, we compute the semantic dissimilarity between the two nodes’ turn-2 idea-sets. We then compare the average dissimilarities among node-pairs with 2, 1 and 0

⁶2-tailed t-tests; Non redundant idea count: for groups (i) vs (iii): $m_1 = 1.96$, $m_3 = 1.72$, $t(602) = 2.403$, Bonferroni-corrected $p < 0.05$, 95% C.I. for $m_1 - m_3 = [0.04, 0.44]$; for (ii) vs (iii): $m_2 = 2.06$, $m_3 = 1.72$, $t(805) = 3.751$, Bonferroni-corrected $p < 0.001$, 95% C.I. for $m_2 - m_3 = [0.17, 0.53]$. Average ratings: for (i) vs (iii): $m_1 = 3.22$, $m_3 = 3.03$, $t(601) = 4.98$, Bonferroni-corrected $p < 0.001$, 95% C.I. for $m_1 - m_3 = [0.11, 0.26]$; for (ii) vs (iii): $m_2 = 3.15$, $m_3 = 3.03$, $t(805) = 3.345$, Bonferroni-corrected $p < 0.003$, 95% C.I. for $m_2 - m_3 = [0.05, 0.18]$. Creativity quotient: for (i) vs (iii): $m_1 = 6.95$, $m_3 = 5.7$, $t(602) = 6.12$, Bonferroni-corrected $p < 0.001$, 95% C.I. for $m_1 - m_3 = [0.86, 1.66]$; for (ii) vs (iii): $m_2 = 6.61$, $m_3 = 5.7$, $t(805) = 4.984$, Bonferroni-corrected $p < 0.001$, 95% C.I. for $m_2 - m_3 = [0.56, 1.28]$




Figure 7: Individual-level comparisons of (A) cumulative non-redundant idea counts, (B) average novelty ratings and (C) cumulative creativity quotients among all the participants of the three study conditions. The static and dynamic conditions significantly outperformed the solo case in the total number of non-redundant ideas (see footnote 8). All other pairwise comparisons show insignificant results. The whiskers in (A) and (C) denote 95% CI. *Bonferroni-corrected $p < 0.05$.

common alters (corresponding to the purple, gray and missing edges respectively in the projected graphs). We find that as the rounds progressed and the dynamic-egos rewired their connections, the turn-2 ideas of node-pairs with 2 common alters gradually became less dissimilar in the dynamic condition ($p < 0.05$ for the negative slope, panel A in the bottom row of Figure 6). Node-pairs with 0 and 1 common alters did not show any such decreasing trend. At the end of the 5th round, the node-pairs with 2 common alters were significantly less dissimilar than 0 and 1 shared alter cases⁷. In the static condition, the alters were the same as the dynamic condition, but the network ties remained fixed. All of the three comparison cases of 0, 1 and 2 common alter node-pairs showed a steadily decreasing dissimilarity ($p < 0.001$ for the negative slope in all three cases), but there was no difference among the three comparison cases (2-tailed test, $p > 0.05$, panel B in the bottom row of Figure 6). For the solo case, there was no stimuli at all, and the semantic dissimilarity did not have any systematic trend ($p = 0.68$ for the slope, panel C in Figure 6).

This shows that as the rounds progressed and the dynamic networks evolved, the ideas of egos who followed the exact same alters increasingly grew similar. It is important to note that this effect is different from groupthink [59], where the desire for harmony or conformity results in consensus among group members. In our case, the egos acted independently without the knowledge of other egos' responses, yet became increasingly similar in association with their choices of the same alters.

Individual creative performance comparisons among various study conditions. We proceed to analyze the individual creative performances in various study conditions. Figure 7A compares the number of non-redundant ideas of the egos and solo participants, aggregated over all the trials. The participants in both the dynamic and static conditions significantly outperformed the solo participants in the total number of non-redundant ideas generated ($p < 0.03$ in both cases⁸). However, there was no significant difference between the participants in the dynamic and static conditions (2-tailed test, $p > 0.05$). Looking at the trials individually, we find that only in trial 3 the static condition significantly outperformed the dynamic condition in the total number of non-redundant ideas generated (2-tailed t-test, $p < 0.05$, SI Figure S5). However, this result was not replicated in any of the other trials ($p > 0.05$, SI Figure S5).

In the case of average novelty ratings, the egos in the dynamic condition showed significantly higher average ratings than the ones in the static condition ($p < 0.05$), but after Bonferroni correction, the difference was no longer significant⁹ (Figure 7B). The other condition-pair comparisons (solo vs dynamic and solo vs static) did not show any significant difference (2-tailed test, $p > 0.05$, same

⁷2-tailed test, 2 vs 0 common alter(s): $m_2 = 3.01$, $m_0 = 3.22$, $t(452) = -2.962$, Bonferroni-corrected $p < 0.01$, 95% C.I. for $m_2 - m_0 = [-0.36, -0.07]$; 2 vs 1 common alter(s): $m_2 = 3.01$, $m_1 = 3.19$, $t(632) = -2.788$, Bonferroni-corrected $p < 0.02$, 95% C.I. for $m_2 - m_1 = [-0.31, -0.05]$

⁸Aggregated over all trials, 2-tailed test; dynamic (d) vs solo (c): $m_d = 6.33$, $m_c = 4.44$, $t(142) = 2.7$, Bonferroni-corrected $p < 0.03$, 95% C.I. for $m_d - m_c = [0.52, 3.26]$; static (s) vs solo (c): $m_s = 6.77$, $m_c = 4.44$, $t(142) = 2.898$, Bonferroni-corrected $p < 0.02$, 95% C.I. for $m_s - m_c = [0.75, 3.9]$

⁹Aggregated over all trials, 2-tailed test; dynamic (d) vs static (s): $m_d = 3.09$, $m_s = 3.03$, $t(214) = 2.042$, $p < 0.04$, Bonferroni-corrected $p > 0.05$, 95% C.I. for $m_d - m_s = [0.004, 0.109]$

figure). In the individual trials, the egos in the dynamic group significantly outperformed the static egos in trials 4 and 6 ($p < 0.05$ in 2-tailed tests), but there was no significant difference in the other four trials (2-tailed test, $p > 0.05$, SI Figure S6).

The creativity quotient metric did not show any significant difference between any of the condition-pairs in any of the individual trials and in the aggregated data (2-tailed tests, $p > 0.05$ for each condition pair, see Figure 7C for aggregated data, SI Figure S7 for individual trial data).

Collective creative performance comparisons among various study conditions. Next, we compare the creative outcomes at a collective level. We refer to the sets of 18 egos in the static and dynamic conditions in each trial as the ‘collective’ or ‘group’ entity. For the solo condition, the 36 participants are randomly split into 2 similar groups of 18 participants. Thus, we compare among 6 dynamic-groups (from the 6 trials), 6 static-groups and 2 solo-groups. The Methods section provides details on the group-level creativity metric computations. As detailed in SI Figures S8, S9 and S10, there was no significant difference between any condition-pair in any of the three metrics (2-tailed tests with Bonferroni correction, $p > 0.05$ in each case).

Thus, at individual and collective levels, we observe no significant benefit in dynamic networks compared to their static counterparts. This is in contrast to another important human performance avenue—collective intelligence—where dynamic link adaptations have been shown to have individual and collective performance benefits over the static condition [22]. However, in typical collective intelligence tasks, people can imitate their peers’ answers to get closer to the known ‘correct’ response. In our study, the task was fundamentally different as it encouraged open ended ideation, and none of the three creativity metrics we used showed any systematic benefit of the dynamic condition.

Simulation model for the observed processes. We present a simulation model for the observed processes, as elaborated in the SI Index. We begin by generating idea-sets for alters such that some of the alters have more non-redundant ideas than others (capturing popular and unpopular alters, respectively). Starting from the same initial network structure as the empirical setup, the egos gradually rewire their connectivity patterns to increasingly follow the popular alters. We model the consequent network-driven and cognitive effects on the ideation process as follows. (A) As rewiring takes place, more egos connect to the popular alters. This is a network-driven process, which makes the stimuli sets overlapping for egos who have either or both alters in common. (B) Given the stimuli set, the egos can generate stimulated ideas, which is a cognitive process. We model the cognitive process to capture the intuition that rare stimuli ideas are associated with better stimulation of new ideas. (C) There can be redundancies in the stimulated ideas of independently ideating egos, if they share the same alters. This is again a network-driven process, as the redundancy is initiated/facilitated by the egos’ similar choices of peers.

The simulation results show that as the egos increasingly follow popular alters (process A), they independently get exposed to more rare ideas, which is associated with their generation of new stimulated ideas (process B). However, as the inter-ego redundancy increases (process C), those stimulated ideas lose statistical rarity. Our simulated model captures these processes and insights that we observed in the empirical data. The results highlight the implications of the two extreme cases of inter-ego redundancy: if there is no inter-ego redundancy whatsoever, then all the stimulated ideas from a given stimulus are different from each other. In that case, process C becomes irrelevant, and the cognitive processes become key to the creative outcomes of the agents in the network. On the other hand, at full inter-ego redundancy, all the stimulated ideas from a single stimulus become exactly the same, in which case none of the stimulated ideas remain statistically rare anymore. These insights are robust to various cognitive stimulation functions that we used to test the generality of the obtained results.

Discussion

Social cues and heuristics are used by humans from their early childhood, lasts throughout the lifespan as they navigate through societal interactions, and contributes to their immense success as a species. In this study, we first explored how the connectivity patterns in a creativity-centric dynamic network adapted to people’s performance cues. From 6 independent trials, we found evidence that the egos’ following/unfollowing patterns are governed significantly by the novelty (measured by average novelty ratings) and statistical rarity (measured by non-redundant idea counts) of the alters’

ideas. The three performance metrics used in the study lead to adjusted $R^2 = 0.79$ in predicting the relative popularity of the alters, suggesting that the cues reported here explain a reasonably high 79% of the variation in the independent variable. Perry-Smith *et al.*'s spiral model suggests that highly creative people will enjoy increased visibility in a dynamic network [12]. Our results validate that idea and explain the relevant cues governing such tie updates. These tie formations are different from preferential attachment [47], since the egos were blind to the number of followers of the alters, and therefore, made their choices without the knowledge of which alters were already popular. The use of a bipartite network structure helped us keep the egos' stimuli set uniform and thus understand the dynamic link update patterns in a cleaner manner. However, this setting comes with the trade off of having only unidirectional edges between alters and egos, as the alters' ideas were pre-recorded. This asymmetry prohibited us from testing the full spiral model, where a back-and-forth influence mechanism is proposed among the interacting nodes.

While the alters were passive actors in the study, our results have implications for them as well. As an example, one can consider the social media influencers, who act as third-party endorsers and shape audience perceptions through blogs, tweets, and other social media channels [60]. The rise of such micro-celebrity has triggered a lot of research interest (e.g., see [61]). In particular, corporate brands are interested in understanding how such influence works, towards leveraging the marketing capital therein [62]. We see in our results that the alters need to generate not only statistically rare ideas, but also ideas that are of high quality and novelty to win more attention over other alters (i.e., competitors). This has implications for the influencers regarding how to stay relevant and ahead of others, in parallel to the Red Queen hypothesis [63].

We confirmed that following high-performing alters is associated with better creative performance of the egos. Indeed, as the egos followed more high-performing alters, the overlap reduced between the ego's own ideas and the alters' ideas—which can partly explain the positive stimulation of ideas that we observed. In the dynamic networks, the egos showed a pattern of flocking behind the high-performing alters, thereby improving their own chances of positive stimulation of novel ideas. However, there was a catch. We saw that ego-pairs who followed the same alters (i.e., exposed themselves to the same stimuli set) in the dynamic condition had an increasing semantic similarity with time, and at the end of the fifth round, had significantly higher semantic similarity than ego-pairs who did not have both of their alters in common. This shows one way dynamic networks can hurt original idea generation from its members: choosing to follow the same stimuli might inadvertently and increasingly make people's ideas similar to each other, despite everyone ideating independently. It is important to note how these processes are driven by the egos' own choices of alters, which captures the dynamic nature of real social networks. In the end, we found no systematic difference in creative outcomes between the dynamic and static conditions, in any of the three metrics. Our simulation model captures the interplay between the network-driven and cognitive processes, and corroborates the generality of the findings.

These insights can lead to research directions that aim to make social networks more creatively competent, both individually and as a whole. As an example, consider a case where many academicians follow the same popular domain experts on social media, seeking inspirations for novel research. Indeed, high quality stimuli can be expected to help the followers generate high quality ideas themselves. But at the same time, there can be redundancies in the stimulated ideas of different followers, hurting the chances of out-of-the-box ideas to be generated. Research on social network intervention strategies will then need to strike a sweet spot to help the followers get the best out of their networks. One way to achieve this could be to recommend people new followers who are diverse and different from one's current set of followers. This might decrease the chances for redundancy that stems from everyone dynamically choosing to follow the same popular stars. The same arguments can apply for offline social networks, such as in large-scale creative projects, where explorations on intelligent intervention can protect against network-driven redundancy being introduced among independent problem-solvers. As for an ego, one can further take away implications that seeking out high-quality creative inspirations can be worthwhile, but flocking behind popular people might not always be optimal for one to stand out.

Drawing from relevant literature, our study settings were designed to reduce performance bottlenecks and increase cognitive stimulation of the egos. The key bottlenecks known to affect creative performance in brainstorming sessions [64, 41, 42] were not present here: there was no evaluation apprehension from peers as the egos were asynchronously exposed to the alters' ideas, no social loafing as the individuals were responsible for their own performance, and no production

blocking as the egos could think on their own in turn 1 of each round. Furthermore, the quiz at the end and the justifications for connection update decisions recorded each round were in place to increase cognitive stimulation and epistemic motivation, as per recommendations in literature [37, 33, 41, 48, 50, 51]. As expected, the dynamic and static conditions, which had access to stimuli from alters, enjoyed a significantly higher count of non-redundant ideas compared to the solo condition, which did not have any stimuli.

Our study is not without limitations. The participants generated alternative uses of common objects in our study—which, despite being widely adopted in the creativity research domain, hardly resembles real-life creative challenges. The study settings prohibited us from exploring the effects of bidirectional creative influence. Also, the study lasted for 5 rounds, which can be prohibitively short to demonstrate the full effects of dynamicity in creative networks. Longitudinal studies with practical creative challenges might generate further insights on the research questions.

Methods

Metrics

The idea of ‘divergent thinking’ leads an individual to come up with numerous and varied responses to a given prompt/situation. Individual differences can then be assessed through measures of fluency (quantity of ideas generated), originality (quantity of unusual ideas), flexibility (the number of different categories implied by the ideas/the ability to generate a wide variety of ideas), elaboration (development of ideas), and other parameters [41, 31, 65]. We choose the three metrics described below, based on previous literature.

Non-redundant idea counts

We operationalized creative performance as the number of non-redundant ideas generated, both at individual and group-levels [52, 53]. First, we filtered out inappropriate answers that did not meet the requirements of being feasible and different from the given use. Then, all the ideas in a given round were organized so that the same ideas are binned or collected together. We followed the coding rules described by Bouchard and Hare [66] and the rules specified in the scoring key of Guilford’s Alternate Uses test, Form B¹⁰.

Once all the ideas were binned, we computed the non-redundant idea counts by looking at the statistical rarity of the ideas submitted by the participants. Namely, an idea was determined to be non-redundant if it was given by at most a threshold number of participants in a given pool of ideas. For the alters, the threshold was set to 1, and the pool was set to be the round-wise idea-set of the 6 alters in the given trial. When comparing trial-wise dynamic and static egos, the threshold was heuristically set to 2, with the pool being the round-wise idea-set of the 36 egos in the trial. In other words, if 3 or more egos in a trial submitted the same idea, it was no longer considered non-redundant. When comparing the data of the solo, static and dynamic conditions aggregated over all trials, the threshold was once again heuristically set to 2, and the pool was set to all the ideas generated by all these participants. Finally, the total number of non-redundant ideas generated by an individual was taken as his/her creativity marker.

2 research assistants independently binned similar ideas together from the first 3 trials of the study. They were shown the anonymized ideas in a random order. Based on their coding, the non-redundant idea counts of the participants were computed, which had a high agreement (Krippendorff’s $\alpha = 0.85$; Spearman’s $\rho = 0.92$, $p < 0.001$, 95% C.I. = [0.885, 0.941]). Then, the first research assistant carried out the coding procedure for the rest of the dataset, which were used in the analyses.

We also computed non-redundant ideas at a group (collective) level of the 18 egos in the static and dynamic conditions in each trial. The cardinality of the set of non-redundant ideas was taken as the performance marker at the group level, computed using the same thresholds and idea pools described above. In other words, if an idea is given by 2 egos in the same group, it is considered a non-redundant idea and the count is increased for both the egos. However, at the group level, that non-redundant idea is counted only once.

¹⁰Guilford’s Alternate Uses Test is Copyright © 1960 by Sheridan Supply Co., all rights reserved in all media, and is published by Mind Garden, Inc, www.mindgarden.com.

Creativity Quotient

For illustrative purposes, consider two imaginary participant’s alternate use ideas for a tire of a truck. Person A submits “boat dock bumper”, “cushion around a mailbox” and “go-kart track boundaries”, all of which are variants of shock absorption. On the other hand, person B submits “As the Os in an autOmOtive repair shOp sign”, “Donate a bunch to Donald Trump to create a wall out of them” and “Use it as a goal hoop in Quidditch”, which are much more spread out and dissimilar in terms of the categories they belong to. Intuitively, person B should be rewarded more. However, it is a matter of subjective judgment to determine the acceptable degree of overlap among different ideas and evaluate them accordingly.

One way to resolve this is to analyze the ideas computationally. The Creativity Quotient metric accounts for both the quantity of ideas generated and the quantity of distinct categories those ideas fall into [55]. Bossomaier *et al.* proposed a semantic network approach for determining the creativity quotient [54]. The idea being, if the ideas of a participant (in a single round, in our case) are very ‘similar’, perhaps they are mostly subtle variations of a small number of categories. Conversely, if they are very dissimilar, perhaps the participant touched many categories—which marks a better creative performance.

The computation of the creativity quotient, Q , relies on an information theoretic measure of semantic similarity derived from WordNet [67]. Concepts appear as syn-sets (set of synonyms) in WordNet, and the nouns come with an ‘is a’ relationship. We begin by mapping the ideas in the dataset to their corresponding concepts or syn-sets. First, we pre-process the ideas by removing stopwords and punctuations, and running a spell-checker algorithm to fix misspelled words. We then split each idea into the set of concepts it is made up of, preferably by replacing verbs and adjectives with related nouns, whenever possible. Then, we find the information content of each of those concepts. Seco *et al.* argued that since the taxonomic structure of WordNet is organized in a meaningful way, concepts with many hyponyms should convey less information than the ones with a small number of hyponyms [56]. In other words, infrequent concepts (such as leaf nodes) should hold more information than the nodes abstracting them. We can therefore quantify the Information Content, I , of a concept c as,

$$I(c) = \frac{\log\left(\frac{h(c)+1}{w}\right)}{\log\left(\frac{1}{w}\right)} = 1 - \frac{\log(h(c) + 1)}{\log(w)} \quad (2)$$

where $h(c)$ is the number of hyponyms of concept c and w is the total number of concepts in the taxonomy. The denominator, in essence, normalizes the metric with respect to the most informative concept, to have an I value between $[0, 1]$.

Using the information contents of the concepts, we then proceed to determine how similar a given pool of ideas are in a certain round. The pool can be the ideas of an individual participant or the idea-set of an entire group. We compute the semantic similarity between every pair of concepts in the pool, c_1 and c_2 , using the following formulation [68]:

$$sim(c_1, c_2) = 1 - \left(\frac{I(c_1) + I(c_2) - 2 \times sim_{MSCA}(c_1, c_2)}{2} \right) \quad (3)$$

Here, the semantic similarity, $sim(c_1, c_2)$, is a function of the amount of information two concepts have in common, $sim_{MSCA}(c_1, c_2)$. This, in turn, is given by the information content of the Most Specific Common Abstraction (MSCA) that subsumes both the concepts in question. Namely,

$$sim_{MSCA}(c_1, c_2) = \max_{c' \in S(c_1, c_2)} I(c') \quad (4)$$

where $S(c_1, c_2)$ is the set of concepts subsuming c_1 and c_2 .

Given the pairwise mutualities of an idea pool, we compute the multi-information, I_m , as the shared information across that response-set. This is computed by first obtaining the maximum spanning tree from the network of concept similarity values between concept pairs, and then summing over the edge weights in the max spanning tree. Finally, the creativity quotient, Q , is obtained by

$$Q = N - I_m \quad (5)$$

where N is the total number of concepts in the idea pool. We compute the $Q_i^{(p)}$ values for each participant i in each round p . The overall creativity quotient for the participant is then computed by

taking the sum over the 5 rounds,

$$Q_i = \sum_{p=1}^5 Q_i^{(p)} \quad (6)$$

Similarly, the overall creativity quotient for a group, Q_g , is computed by,

$$Q_g = \sum_{p=1}^5 Q_g^{(p)}. \quad (7)$$

Ratings

Each of the 36 participants in the static and dynamic groups rated their respective alters' ideas in each round. This led to a total of 38,880 ego-generated novelty ratings that we use in our analyses. We take the mean rating received by each alter from each ego, and compute the consequent intra-class correlation coefficient among the ego-raters in that trial. The mean intra-class correlation coefficient from all 6 trials was $ICC(3, 36) = 0.945$.

For the egos' ideas, we collected ratings from Amazon Mechanical Turk workers ('raters'), who were different individuals from the study participants. To make the rating comparisons fair, each rater focused on only one round, and rated all the 36 participants' ideas on that round in a given trial (static and dynamic conditions). They were first given 3 minutes to generate ideas on that round's prompt themselves, so that they could appreciate the difficulty-levels associated with the tasks. Then the ideas were shown to them in a random order, which they rated on novelty (5-point scale, 1: not novel, 5: highly novel). The raters were guided with detailed instructions and examples. Each idea was rated by at least 4 raters. The ideas of the 36 participants in the solo condition were similarly rated. 141 raters were hired, who generated a total of 40,320 ratings in the dataset. Computed the same way as the ego-raters above, we find a positive mean intra-class correlation coefficient among the raters, $ICC(3, 4) = 0.317$.

For an idea j , the mean of all its ratings is taken as the idea-rating r_j . For a participant i , the average novelty rating is computed as $\bar{r}_i = mean(r_j)$, for all ideas j submitted by i . This average novelty rating is used as the quality indicator of the participant in our analyses. For the collective level, the mean rating of all the participants is taken as the creative performance marker.

Measure of idea overlap

To measure the overlap between idea-sets A and B , we use the Jaccard Index:

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (8)$$

If $A = B = \emptyset$, we take $J(A, B) = 1$. We resort to the manual annotations done in non-redundant idea count computations to decide which ideas should be considered the same entry in the idea-sets.

Measure of semantic dissimilarity

To measure how semantically dissimilar two sets of ideas are, we use the Word Mover's Distance metric [58]. First, we remove all of the stop-words and punctuation from the ideas, and find their Word2Vec [57] embeddings. The dissimilarity or distance between two idea-sets is then computed by taking the minimum Euclidean distance that the embedded words of one idea-set need to travel to reach the embedded words of another idea-set.

Code and Data Availability

The data analysis and simulation code is available at http://bit.ly/dynamic_creativity19 (temporary link for review). The full data is not publicly released to ensure compliance with the copyright requirements of some of the study materials.

References

- [1] Frank, M. R. *et al.* Toward understanding the impact of artificial intelligence on labor. *Proceedings of the National Academy of Sciences* **116**, 6531–6539 (2019).
- [2] Mitchell, T. & Brynjolfsson, E. Track how technology is transforming work. *Nature* **544**, 290–292 (2017).
- [3] Brynjolfsson, E. & Mitchell, T. What can machine learning do? Workforce implications. *Science* **358**, 1530–1534 (2017).
- [4] Manyika, J. *et al.* A future that works: Automation, employment, and productivity. *McKinsey Global Institute* (2017).
- [5] Acemoglu, D. & Restrepo, P. Robots and jobs: Evidence from US labor markets. *Journal of Political Economy* (2019). <https://doi.org/10.1086/705716>.
- [6] Rotman, D. Making AI into jobs. *MIT Technology Review* **121**, 10–17 (2018).
- [7] Alabdulkareem, A. *et al.* Unpacking the polarization of workplace skills. *Science Advances* **4** (2018).
- [8] Baten, R. A., Clark, F. & Hoque, M. E. Upskilling together: How peer-interaction influences speaking-skills development online. In *8th International Conference on Affective Computing & Intelligent Interaction (ACII)* (2019).
- [9] Wu, L., Wang, D. & Evans, J. A. Large teams develop and small teams disrupt science and technology. *Nature* **566**, 378 (2019).
- [10] Milojević, S. Principles of scientific research team formation and evolution. *Proceedings of the National Academy of Sciences* **111**, 3984–3989 (2014).
- [11] Kazanjian, R. K., Drazin, R. & Glynn, M. A. Creativity and technological learning: The roles of organization architecture and crisis in large-scale projects. *Journal of Engineering and Technology Management* **17**, 273 – 298 (2000).
- [12] Perry-Smith, J. E. & Shalley, C. E. The social side of creativity: A static and dynamic social network perspective. *Academy of Management Review* **28**, 89–106 (2003).
- [13] Kijkuit, B. & Van Den Ende, J. The organizational life of an idea: Integrating social network, creativity and decision-making perspectives. *Journal of Management Studies* **44**, 863–882 (2007).
- [14] Perry-Smith, J. E. Social yet creative: The role of social relationships in facilitating individual creativity. *Academy of Management Journal* **49**, 85–101 (2006).
- [15] Henrich, J. *The Secret of Our Success: How Culture is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter* (Princeton University Press, 2015).
- [16] Henrich, J., Chudek, M. & Boyd, R. The big man mechanism: How prestige fosters cooperation and creates prosocial leaders. *Philosophical Transactions of the Royal Society of London B: Biological Sciences* **370** (2015).
- [17] Welles, B. F. & Contractor, N. Individual motivations and network effects: A multilevel analysis of the structure of online social relationships. *The ANNALS of the American Academy of Political and Social Science* **659**, 180–190 (2015).
- [18] Bollen, J. & Gonçalves, B. Network happiness: How online social interactions relate to our well being. In *Complex Spreading Phenomena in Social Systems*, 257–268 (Springer, 2018).
- [19] Herrmann, E., Call, J., Hernández-Lloreda, M. V., Hare, B. & Tomasello, M. Humans have evolved specialized skills of social cognition: The cultural intelligence hypothesis. *Science* **317**, 1360–1366 (2007).
- [20] Boyd, R., Richerson, P. J. & Henrich, J. The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences* **108**, 10918–10925 (2011).
- [21] Rand, D. G., Arbesman, S. & Christakis, N. A. Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences* **108**, 19193–19198 (2011).

- [22] Almaatouq, A. *et al.* The wisdom of the network: How adaptive networks promote collective intelligence (2018). arxiv.org/abs/1805.04766.
- [23] Bernstein, E., Shore, J. & Lazer, D. How intermittent breaks in interaction improve collective intelligence. *Proceedings of the National Academy of Sciences* **115**, 8734–8739 (2018).
- [24] Shafipour, R. *et al.* Buildup of speaking skills in an online learning community: A network-analytic exploration. *Palgrave Communications* **4**, 63 (2018).
- [25] Jiang, H., Zhang, Q.-P. & Zhou, Y. Dynamic creative interaction networks and team creativity evolution: A longitudinal study. *The Journal of Creative Behavior* **52**, 168–196 (2018).
- [26] Chen, M.-H., Chang, Y.-C. & Hung, S.-C. Social capital and creativity in R&D project teams. *R&D Management* **38**, 21–34 (2008).
- [27] Leenders, R. T. A., Van Engelen, J. M. & Kratzer, J. Virtuality, communication, and new product team creativity: A social network perspective. *Journal of Engineering and Technology Management* **20**, 69–92 (2003).
- [28] Siangliulue, P., Arnold, K. C., Gajos, K. Z. & Dow, S. P. Toward collaborative ideation at scale: Leveraging ideas from others to generate more creative and diverse ideas. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 937–945 (ACM, 2015).
- [29] Chan, J., Dang, S. & Dow, S. P. Comparing different sensemaking approaches for large-scale ideation. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2717–2728 (ACM, 2016).
- [30] Nijstad, B. A., Diehl, M. & Stroebe, W. Cognitive stimulation and interference in idea generating groups. *Group Creativity: Innovation through Collaboration* 137–159 (2003).
- [31] Runco, M. A. *Creativity: Theories and Themes: Research, Development, and Practice* (Elsevier, 2014).
- [32] Carson, S. *Your Creative Brain: Seven Steps to Maximize Imagination, Productivity, and Innovation in Your Life* (John Wiley & Sons, 2010).
- [33] Dugosh, K. L. & Paulus, P. B. Cognitive and social comparison processes in brainstorming. *Journal of Experimental Social Psychology* **41**, 313–320 (2005).
- [34] Mednick, S. The associative basis of the creative process. *Psychological Review* **69**, 220 (1962).
- [35] Collins, A. M. & Loftus, E. F. A spreading-activation theory of semantic processing. *Psychological Review* **82**, 407 (1975).
- [36] Paulus, P. B. & Brown, V. R. Toward more creative and innovative group idea generation: A cognitive-social-motivational perspective of brainstorming. *Social and Personality Psychology Compass* **1**, 248–265 (2007).
- [37] Nijstad, B. A. & Stroebe, W. How the group affects the mind: A cognitive model of idea generation in groups. *Personality and Social Psychology Review* **10**, 186–213 (2006).
- [38] Brown, V. R. & Paulus, P. B. Making group brainstorming more effective: Recommendations from an associative memory perspective. *Current Directions in Psychological Science* **11**, 208–212 (2002).
- [39] Dahl, D. W. & Moreau, P. The influence and value of analogical thinking during new product ideation. *Journal of Marketing Research* **39**, 47–60 (2002).
- [40] Doboli, A., Umbarkar, A., Subramanian, V. & Doboli, S. Two experimental studies on creative concept combinations in modular design of electronic embedded systems. *Design Studies* **35**, 80–109 (2014).
- [41] Paulus, P. Groups, teams, and creativity: The creative potential of idea-generating groups. *Applied Psychology* **49**, 237–262 (2000).
- [42] Dennis, A. & Williams, M. Electronic brainstorming. *Group Creativity: Innovation through Collaboration* 160–178 (2003).
- [43] Burt, R. S. Structural holes and good ideas. *American Journal of Sociology* **110**, 349–399 (2004).

- [44] Zhou, J., Shin, S. J., Brass, D. J., Choi, J. & Zhang, Z.-X. Social networks, personal values, and creativity: Evidence for curvilinear and interaction effects. *Journal of Applied Psychology* **94**, 1544 (2009).
- [45] Kozbelt, A., Beghetto, R. A. & Runco, M. A. Theories of creativity. *The Cambridge Handbook of Creativity* **2**, 20–47 (2010).
- [46] Guildford, J., Christensen, P., Merrifield, P. & Wilson, R. Alternate uses: Manual of instructions and interpretation. *Orange, CA: Sheridan Psychological Services* (1978).
- [47] Newman, M. *Networks* (Oxford University Press, 2018).
- [48] Brown, V., Tumeo, M., Larey, T. S. & Paulus, P. B. Modeling cognitive interactions during group brainstorming. *Small Group Research* **29**, 495–526 (1998).
- [49] Coursey, L. E., Williams, B. C., Kenworthy, J. B., Paulus, P. B. & Doholi, S. Divergent and convergent group creativity in an asynchronous online environment. *The Journal of Creative Behavior* (2018).
- [50] Bechtoldt, M. N., De Dreu, C. K., Nijstad, B. A. & Choi, H.-S. Motivated information processing, social tuning, and group creativity. *Journal of Personality and Social Psychology* **99**, 622 (2010).
- [51] Scholten, L., Van Knippenberg, D., Nijstad, B. A. & De Dreu, C. K. Motivated information processing and group decision-making: Effects of process accountability on information processing and decision quality. *Journal of Experimental Social Psychology* **43**, 539–552 (2007).
- [52] Oppezzo, M. & Schwartz, D. L. Give your ideas some legs: The positive effect of walking on creative thinking. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **40**, 1142 (2014).
- [53] Abdullah, S., Czerwinski, M., Mark, G. & Johns, P. Shining (blue) light on creative ability. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 793–804 (ACM, 2016).
- [54] Bossomaier, T., Harré, M., Knittel, A. & Snyder, A. A semantic network approach to the Creativity Quotient (CQ). *Creativity Research Journal* **21**, 64–71 (2009).
- [55] Snyder, A., Mitchell, J., Bossomaier, T. & Pallier, G. The Creativity Quotient: An objective scoring of ideational fluency. *Creativity Research Journal* **16**, 415–419 (2004).
- [56] Seco, N., Veale, T. & Hayes, J. An intrinsic information content metric for semantic similarity in WordNet. In *Proceedings of the 16th European Conference on Artificial Intelligence, ECAI*, vol. 16, 1089 (2004).
- [57] Mikolov, T., Chen, K., Corrado, G. & Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [58] Kusner, M., Sun, Y., Kolkin, N. & Weinberger, K. From word embeddings to document distances. In *International Conference on Machine Learning*, 957–966 (2015).
- [59] Nemeth, C. & Nemeth-Brown, B. Better than individuals. *Group Creativity: Innovation through Collaboration* **4**, 63–84 (2003).
- [60] Freberg, K., Graham, K., McGaughey, K. & Freberg, L. A. Who are the social media influencers? A study of public perceptions of personality. *Public Relations Review* **37**, 90–92 (2011).
- [61] Khamis, S., Ang, L. & Welling, R. Self-branding, ‘micro-celebrity’ and the rise of social media influencers. *Celebrity Studies* **8**, 191–208 (2017).
- [62] Booth, N. & Matic, J. A. Mapping and leveraging influencers in social media to shape corporate brand perceptions. *Corporate Communications: An International Journal* **16**, 184–191 (2011).
- [63] Derfus, P. J., Maggitti, P. G., Grimm, C. M. & Smith, K. G. The red queen effect: Competitive actions and firm performance. *Academy of Management Journal* **51**, 61–80 (2008).
- [64] Paulus, P. B., Putman, V. L., Dugosh, K. L., Dzindolet, M. T. & Coskun, H. Social and cognitive influences in group brainstorming: Predicting production gains and losses. *European Review of Social Psychology* **12**, 299–325 (2002).
- [65] Sawyer, R. K. *Explaining Creativity: The Science of Human Innovation* (Oxford University Press, 2011).

- [66] Bouchard Jr, T. J. & Hare, M. Size, performance, and potential in brainstorming groups. *Journal of Applied Psychology* **54**, 51 (1970).
- [67] Miller, G. A. Wordnet: A lexical database for English. *Communications of the ACM* **38**, 39–41 (1995).
- [68] Jiang, J. J. & Conrath, D. W. Semantic similarity based on corpus statistics and lexical taxonomy. In *Proceedings of the International Conference on Research in Computational Linguistics* (1998).