INTERFACE

royalsocietypublishing.org/journal/rsif

Research

Cite this article: Baten RA, Bagley D, Tenesaca A, Clark F, Bagrow JP, Ghoshal G, Hoque E. 2020 Creativity in temporal social networks: how divergent thinking is impacted by one's choice of peers. *J. R. Soc. Interface* **17**: 20200667.
http://dx.doi.org/10.1098/rsif.2020.0667

Received: 17 August 2020
Accepted: 22 September 2020

Subject Category:
Life Sciences–Physics interface

Subject Areas:
biophysics

Keywords:
network science, creativity, dynamic social networks

Author for correspondence:
Ehsan Hoque
e-mail: mehoque@cs.rochester.edu

Electronic supplementary material is available online at https://doi.org/10.6084/m9.figshare.c.5144380.

THE ROYAL SOCIETY PUBLISHING

Creativity in temporal social networks: how divergent thinking is impacted by one's choice of peers

Raiyan Abdul Baten[1], Daryl Bagley[2], Ashely Tenesaca[2], Famous Clark[2], James P. Bagrow[4], Gourab Ghoshal[3] and Ehsan Hoque[2]

[1]Department of Electrical and Computer Engineering, [2]Department of Computer Science, and [3]Department of Physics and Astronomy, University of Rochester, Rochester, NY, USA
[4]Department of Mathematics & Statistics, University of Vermont, VT, USA

RAB, 0000-0002-9203-4528; DB, 0000-0003-2926-7203; JPB, 0000-0002-4614-0792; GG, 0000-0001-7593-5626; EH, 0000-0003-4781-4733

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 social networks impacts the fostering of creative ideas. We run six 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 similarities in the followers' ideas. We formulate an agent-based simulation model to capture these intuitions in a tractable manner, and experiment with corner cases of various simulation parameters to assess the generality of the findings. Our findings may help design large-scale interventions to improve the creative aptitude of people interacting in a social network.

1. Introduction

Recent advances in robotics, AI and machine learning are increasingly focused on mimicking or even surpassing human capabilities. These innovations, however, have serious implications on our future workforce [1]. Approximately 51% of the tasks done in the US economy can be automated [2], and for each robot on the factory floor, some six jobs are lost [3]. The need for manual labour in predictable and repetitive work is declining, while the demand is soaring for expertise in creative tasks, problem-solving and other social-cognitive avenues of soft-skills [2,4,5]. Again, many of the critical and challenging tasks of the human civilization require humans to collaborate with others [6], where they need to 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 [7]. Discussions in an academic network of researchers, faculty members and students can stimulate ideas for novel explorations. A graphic designer can find creative inspiration from peer-interactions in online networks like Reddit, Behance or Twitter. Adopting a social network lens helps us better understand the mechanisms, bottlenecks and opportunities for maximizing creative outcomes at-scale.

Researchers have examined the effects of various network attributes on creativity [8]. For instance, relationship strength, position and external ties of people

are known to influence creative performance [9]. 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 [10], often in response to the behaviour, performance, prestige, age, gender, popularity, self-similarity and other cues of the social partners [11–13]. This dynamic characteristic affords opportunities in human populations that static networks cannot: for example, dynamic networks promote cooperation [14,15], collective intelligence [16], and speaking skills [17], among others.

When it comes to creativity, 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 [8]. 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. Again, a central person can inspire creative thoughts in others and 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 such temporal effects from network and evolutionary game-theoretic perspectives [18–20], it remains largely unclear how dynamic creative networks evolve with time, what laws they follow, and what implications such evolving has on creative ideation.

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. Second, we test how such adaptations impact creative ideation performances in a dynamic network, against the controls of static (fixed connections) and solo (unconnected) networks.

We run six trials of a web-based experiment, where participants in a dynamic social network performed idea generation tasks for five rounds. The participants chose after each round which of the peers' ideas they wanted to be shown as stimuli (see Experimental set-up). Following [8], we anticipated that people looking for creative inspiration from others will use success cues to determine who among the peers are more creative and, therefore, more promising to be advantageous to form ties with (follow). We find that in the dynamic networks, the participants were indeed drawn towards following the most creative ideators. The statistical rarity and novelty ratings of one's ideas appear to be robust predictors of his/her popularity in the networks.

If people preferentially form ties with the highly creative ideators in the network, what implications will it have on the creative outcomes? The associative theory of creative cognition suggests that priming people properly, e.g. by exposing them to others' ideas [21,22], can stimulate their long-term memory circuitry. This can enable retrieving remotely stored concepts from memory [23]. Combining such remote concepts can help in synthesizing novel ideas [24,25]. Based on this, we anticipated that following highly creative peers can enable better creative performance in people. For instance, following a person who generates rare ideas will increase the chances that the follower comes across ideas that have little overlap with his/her own. This can stimulate further ideas through novel association of concepts, resulting in ideas that would not have occurred to the follower otherwise [26,27]. 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 [28,29]. Our results show that following highly creative ideators is indeed associated with one's better creative performance. However, participants who chose to follow the same people (same stimuli) show an increasing semantic similarity in their independently stimulated ideas with time. These results suggest that self-organizing in a dynamic network can lead to conflicting opportunities and constraints. We formulate a simulation model that captures these empirically derived intuitions, and helps assess the generality of the processes and insights.

1.1. Experimental set-up

It is challenging to identify a dataset in the wild that offers traceable links between ideas and their stimuli in a temporal manner. We therefore resort to an artificial social network created in the virtual laboratory.

We are interested in divergent creativity, which signifies a person's ability to come up with numerous and varied responses to a given prompt/situation [30]. We use a customized version of Guilford's Alternate Uses Test [31], a widely adopted approach for quantifying divergent creativity. In each of five rounds, the participants were instructed to consider an everyday object (e.g. a brick), whose common use was stated (used for building). The participants needed to submit alternative uses for the object which are different than the given use, yet are appropriate and feasible. We choose the first five objects from the Form b of Guilford's test as the ideation objects respectively in the five rounds.

We recruited 288 participants from Amazon Mechanical Turk, who live in USA and are diverse in their age, racial, ethnic and gender distributions (see electronic supplementary material). 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 (figure 1*a*). There are two types of nodes in the network, *alters* and *egos*. First, we pre-recorded the ideas of six *alters*, who generated ideas independently. 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 ensure uniform stimuli for all the egos in the static and dynamic conditions. We repeated the process for six independent trials, each with its unique alters. Under the solo condition, 36 participants generated ideas in isolation.

Initially, the dynamic and static egos were connected to two alters each using the network structure in figure 1*a*. In each round, the egos first generated ideas independently for 3 min (turn-1). They were then shown the ideas of the two alters they were connected to, and given an additional 3 min to list further ideas (turn-2). The egos were instructed 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

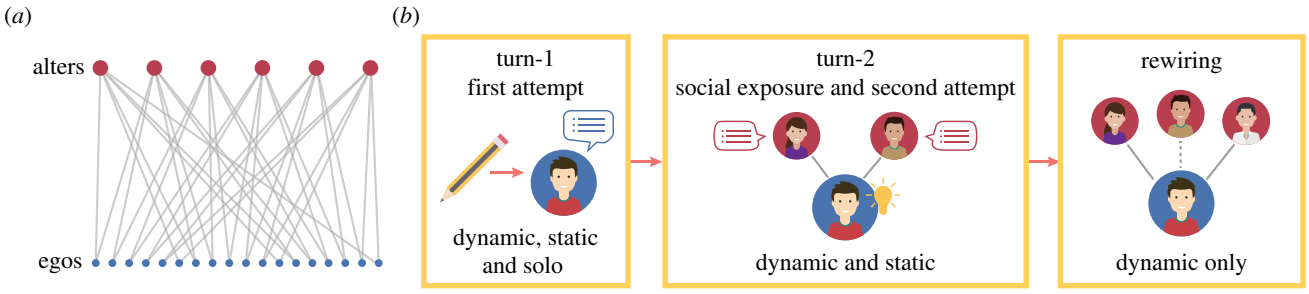

Figure 1. (*a*) The bipartite network structure used in the static condition, and as the initial configuration in the dynamic condition. (*b*) Study protocol for each round. In turn-1, the participants generated ideas independently. In turn-2, the dynamic and static egos received social exposure and could list further ideas. Only the dynamic egos could update which alters to follow at the end of each round.

them. This was to ensure that the participants paid attention to the stimuli, which is known to positively impact ideation [32]. Then, the egos rated the ideas of all of the six alters on novelty (five-point ratings, 1: not novel, 5: highly novel). Finally, in the dynamic condition, the egos could optionally follow/unfollow alters at the end of each round to have an updated list of two followees 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 is known to raise epistemic motivation and improve systematic information processing [33]. The egos in the static condition could not update their links, marking the only difference between the static and dynamic conditions.

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

Measures. We quantify creativity using three metrics: (1) non-redundant idea counts, (2) average novelty ratings and (3) creativity quotient (see Methods for details). If an idea is given by at most a threshold number of participants in a given participant pool, it is considered non-redundant [34]. The non-redundant idea count is thus a measure of the statistical rarity of the ideas.

Each idea in the dataset was rated on novelty by multiple people. We take the mean rating received by an idea as its novelty indicator, and consider the average novelty rating of a given set of ideas as a creativity metric.

The creativity quotient metric, Q, uses information theoretic measures to capture the fluency (quantity of ideas) and flexibility (the ability to generate a wide variety of ideas) of a set of ideas [35–37]. In all three metrics, higher values indicate better creative performances.

2. Results

Link update patterns in the network evolution. We first explore how connectivity patterns adapt to individual performance cues in creativity-centric dynamic networks. The networks evolved as the egos updated their lists of two alters across the rounds. All of the alters started with six ego-followers, but after five rounds of network evolution, some of the alters lost followers, and some of them gained. The final evolved networks of the six trials are shown in figure 2. The number of connection updates per ego at the end of each round had a downward trend ($p < 0.001$ for the negative slope, figure 3). Out of a maximum of two possible updates, an average of 0.97 connections were updated per ego after the first round, while 0.3 connections were updated per ego after the fifth round. This suggests that as the egos received information about the alters' performances, they made up their minds on whom to follow, and readjusted later if necessary.

We denote alters who finished with greater than six and less than or equal to six followers as 'popular' and 'unpopular' alters respectively. If the egos did not update their links at all, the alters would still have the initially assigned six followers, so we take greater than six followers as a marker of popularity. The total non-redundant idea counts of the popular alters in all five rounds were significantly higher than those of unpopular alters (popular (p) versus unpopular (u) alters: two-tailed test, $n_p = 13$, $n_u = 23$, $t(34) = 7.291$, $p < 0.001$, figure 4*a*; $p < 0.05$ in all six trials, electronic supplementary material, figure S1). Similarly, the popular alters outperformed the unpopular alters significantly in the average novelty ratings (two-tailed test, $t(34) = 5.7$, $p < 0.001$, figure 4*b*; $p < 0.05$ in five out of six trials, electronic supplementary material, figure S2). The total creativity quotient Q of the popular alters in all five rounds were again significantly higher than those of unpopular alters (two-tailed test, $t(34) = 5.81$, $p < 0.001$, figure 4*c*; $p < 0.05$ in three out of six trials, electronic supplementary material, figure S3). Electronic supplementary material, table S1 provides further details.

Using multivariate linear regression, we explore how the alters' creative performances correspond to their final popularity. The rational being, by the end of the fifth round, the egos had full knowledge of the qualities of the alters' ideas. So, we can expect the egos' final choices to be captured by the overall performances of the alters. As the dependent variable y_i, we take the fraction of egos connected to an alter i at the end of the fifth round. The independent variables are: (1) the relative number of non-redundant ideas, $u_i' = u_i/\sum_i u_i$, (2) relative average novelty ratings, $\bar{r}_i' = \bar{r}_i/\sum_i \bar{r}_i$ and (3) relative creativity quotient, $Q_i' = Q_i/\sum_i Q_i$. 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 five rounds. 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 the independent variables correlate strongly with the dependent variable (Pearson's $\rho = 0.80$, 0.86, 0.75, respectively, $p < 0.001$ in each),

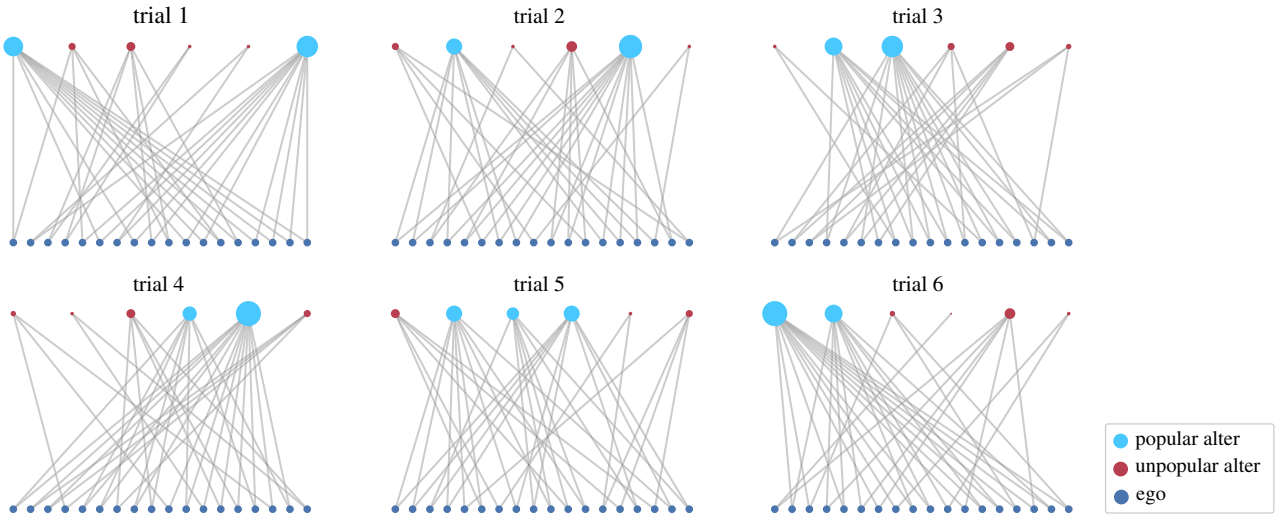

Figure 2. Final/evolved dynamic networks. The diameters of the circles in the upper-rows are proportional to the alters' follower counts at the end of the fifth round.

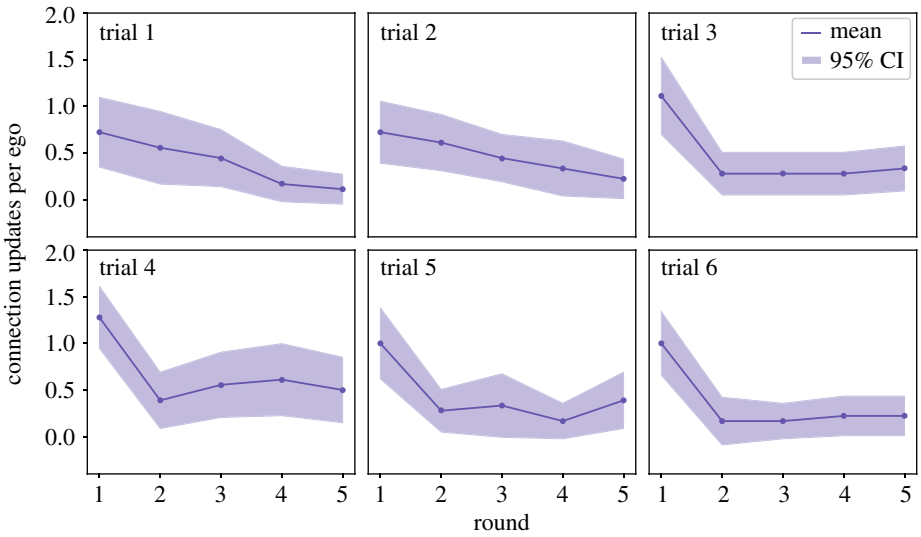

Figure 3. Number of connections updated per ego at the end of each round.

multivariate regression helps us explore the relative contributions of the three independent variables. Mathematically,

$$y_i = \beta_0 + \beta_1 u'_i + \beta_2 \bar{r}'_i + \beta_3 Q'_i. \tag{2.1}$$

We train on 70% data, compute R^2 and adjusted-R^2 on the remaining 30% data, and repeat the process 10 times (each time with a new 70/30 split) to compute a confidence interval around the mean R^2 and mean adjusted-R^2 values. Table 1 summarizes the results. We first test the independent variables separately using univariate regression, and find $\bar{r}'$ to give the best adjusted-$R^2_{\bar{r}'} = 0.69$ (Model 1). Adding u' makes the adjusted-$R^2_{\bar{r}',u'} = 0.72$ (Model 2), while all three features give an adjusted-$R^2_{\bar{r}',u',Q'} = 0.70$ (Model 3). In Model 3, the coefficients are significant for $\bar{r}'$ and u', but not for Q'. This shows that the associations reported here are systematic for the first two predictors ($\bar{r}'$ and u'), which together explain 72% of the variation in the dependent variable. Thus, the egos' connectivity dynamics are strongly captured by the novelty of the alters' ideas ($\bar{r}'$) and moderately by the statistical rarity of those ideas (u'). The relative creativity quotient metric has much of its information overlapped with the other two metrics.

As most of the connection updates happened after the first round, we explore how the alters' first round performances can predict their popularity at the end of the first and fifth rounds. We find that the features $\bar{r}$ and Q from the first round together give the best adjusted-$R^2 = 0.67$ in predicting the alters' popularity after the first round (95% CI = [0.59, 0.76]). Similarly, for predicting the popularity after the fifth round only from the first round's performances, the feature $\bar{r}$ alone gives the best adjusted-$R^2 = 0.60$ (95% CI = [0.52, 0.67]). This drop in the explainability of the dependent variable can be due to the fact that when we try to predict the final popularity of the alters from the first round's data, we miss out on the performance information from the middle rounds, which can have an effect on the ego's final choices. When we incorporate all five rounds' data, we are able to predict the final popularities with adjusted-R^2 as high as 0.72 (Model 2).

The key take-away is that the egos in a dynamic network predominantly form ties with the consistently high-performing alters, as anticipated. In typical creativity studies, participants are shown stimuli randomly or based on intrinsic qualities [21,22]. Here, a key contrast is that the networks are allowed

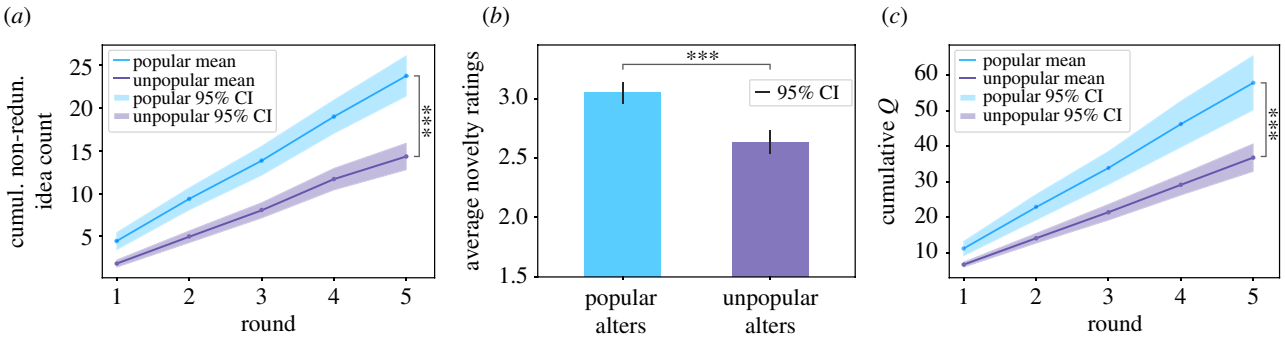

Figure 4. The popular alters significantly outperformed the unpopular alters in all three metrics. (*a*) The total number of non-redundant ideas in all five rounds (two-tailed test, $t(34) = 7.291$, $p < 0.001$), (*b*) average novelty ratings in all five rounds ($t(34) = 5.7$, $p < 0.001$) and (*c*) the total creativity quotient in all five rounds ($t(34) = 5.81$, $p < 0.001$). ***$p < 0.001$.

Table 1. Regression results of predicting the alters' relative popularity from their relative creativity markers. β = standardized regression coefficient.

predictor	Model 1: $\bar{r}'$ only		Model 2: $\bar{r}'$, u'		Model 3: all predictors	
	β	t (s.e.)	β	t (s.e.)	β	t (s.e.)
$\bar{r}'$	0.1851***	9.762 (0.019)	0.1278***	4.972 (0.026)	0.1026**	3.435 (0.030)
u'	—	—	0.0767**	2.983 (0.026)	0.0704**	2.763 (0.025)
Q'	—	—	—	—	0.0399	1.562 (0.026)
R^2	0.72, 95% CI = [0.69, 0.75]		0.78, 95% CI = [0.74, 0.81]		0.79, 95% CI = [0.75, 0.83]	
adjusted-R^2	0.69, 95% CI = [0.66, 0.72]		0.72, 95% CI = [0.67, 0.77]		0.70, 95% CI = [0.65, 0.75]	

**$p < 0.01$. ***$p < 0.001$. $N = 36$ for each model.

to dynamically self-organize, as one can choose for oneself who to take inspirations from. Thus, the implications of such adaptations on the ideation process, as explored below, become direct manifestations of the dynamic nature of the networks.

Exposure to high-performing alters is associated with better creative performances 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 by facilitating novel associations between remote concepts. To test this, we take advantage of the fact that in turn-1, the egos generated ideas independently prior to any social exposure, which allows us to test the overlap their ideas have with their alters' ideas (measured by Jaccard index, see Methods). In turn-2, the egos could see the stimuli, allowing us to explore whether the creative qualities of the stimulated ideas have any association with how creative the respective alters were.

We first find the round-wise popular alters by identifying those with greater than six followers at the end of each round. Then, for each round, we split the egos into three 'groups', where (i) both, (ii) only one and (iii) none of the followees of an ego are round-wise popular. Thus, we analyse each ego's data across five rounds, where the ego belongs to one of these three 'groups' in each round. The group sample sizes are $n_i = 273$, $n_{ii} = 476$ and $n_{iii} = 331$, respectively. We assess the fixed effects of the group and round factors (three and five levels, respectively) against the random effects of the egos to control for repeated measures. We employ the aligned rank transform (ART) procedure, which is a linear mixed model-based non-parametric test [38], available in the ARTool package in R [39].

We analyse the average overlap between each ego's turn-1 ideas and his/her two alters' ideas. We find a significant main effect of the group factor (ART procedure, $F(2, 669.84) = 66.53$, $p < 0.001$). *Post hoc* pairwise comparisons among the group factor levels in the fitted ART model show that group (i) had significantly less idea-overlap than group (ii) ($t(887) = -7.52$, $p < 0.001$), and group (ii) had significantly less overlap than group (iii) ($t(574) = -5.56$, $p < 0.001$). All p-values are Bonferroni-corrected. Thus, following more high-performing alters systematically decreased the overlap between an ego's turn-1 ideas and the alters' ideas, as anticipated (figure 5*a*; electronic supplementary material, tables S2 and S3). Group (i) consistently had less overlap than groups (ii) and (iii) in each round (electronic supplementary material, figure S4).

We then explore the creativity measures of the turn-2 ideas across the three levels of the group factor. We fit three separate models for the three metrics. All of the three creativity metrics show significant main effects for both the group and round factors, but no significant interaction between the factors (main effects of the group factor: (1) non-redundant idea counts, $F(2, 825.86) = 3.7$, $p = 0.025$, (2) average novelty ratings, $F(2, 535.47) = 11.85$, $p < 0.001$, (3) creativity quotient, $F(2, 1036.36) = 6.66$, $p = 0.001$).

We proceed to conduct post-hoc analysis among the three group levels as before. We find that group (iii), i.e. egos who followed unpopular alters only, showed the

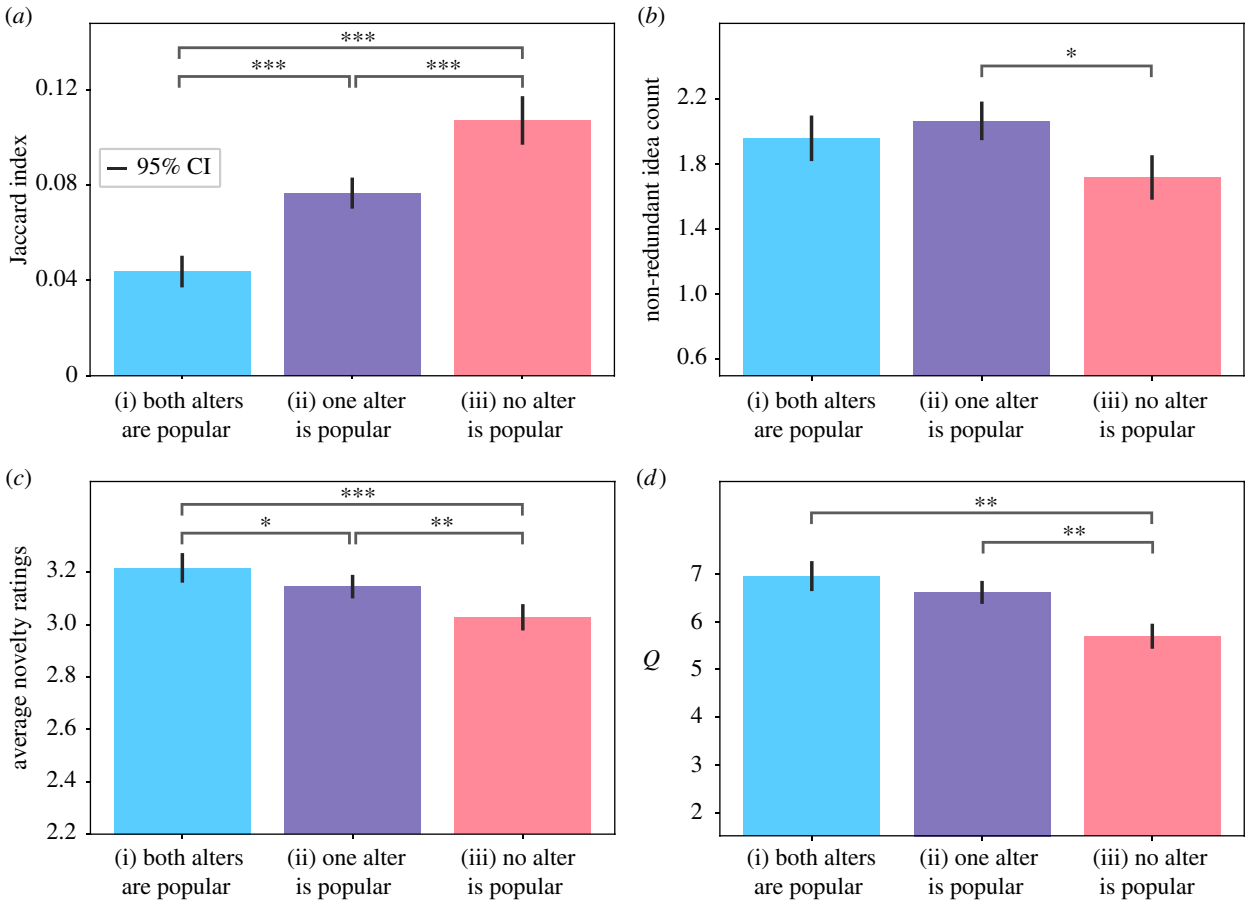

Figure 5. (*a*) Average overlap between the egos' turn-1 ideas and their alters' ideas. We compare among egos with (i) both, (ii) only one and (iii) no round-wise popular alters. Following more high-performing alters decreases the egos' exposure to overlapping ideas. (*b–d*) The creative performances of the same groups of egos in turn-2. In all three metrics, egos in group (iii) perform the worst. *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$; all p-values corrected for multiple comparisons and repeated measures.

worst performance in all three metrics. In particular, group (ii) significantly outperformed group (iii) in non-redundant idea counts ($t(727) = 2.69$, $p = 0.02$), but the other two pairwise comparisons (group (i) versus group (ii) and group (i) versus group (iii)) were insignificant in this metric. In case of the average novelty ratings, group (i) significantly outperformed group (ii) ($t(721) = 1.97$, $p = 0.049$), and group (ii) significantly outperformed group (iii) ($t(458) = 3.4$, $p = 0.002$). As for creativity quotient, both groups (i) and (ii) outperformed group (iii) (respectively, $t(1022) = 3.06$, $p = 0.005$ and $t(1016) = 3.5$, $p = 0.002$). There was no significant difference between groups (i) and (ii). See figure 5*b–d* and electronic supplementary material, tables S4–S5. These results imply that following at least one high-performing alter is associated with better creative performance of the egos. Our results thus suggest evidence for better stimulation of ideas when egos are exposed to high-quality ideas.

Following the same alters introduces semantic similarities in the egos' ideas. We further motivated a counter-argument that if multiple egos follow the same alters, their stimuli sets will become overlapping. This can make the egos' stimulated ideas similar, despite ideating independently. To test this, we explore whether the semantic (dis)similarities between each node-pair's turn-2 ideas have any association with the number of common alters they have. We estimate the semantic nature of ideas using neural word embeddings (Word2Vec [40]) and compare the dissimilarity of the embeddings using Word Mover's Distance [41] (see Methods).

We first take one-mode projections 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). We compute the semantic dissimilarity between each node-pair's turn-2 idea-sets. We then compare the dissimilarities among node-pairs with 2, 1 and 0 common alters (corresponding to the purple, grey and missing edges, respectively, in the projected graphs).

As the dynamic-egos rewired their connections across rounds, the turn-2 ideas of node-pairs with 2 common alters gradually became less dissimilar ($p < 0.05$ for the negative slope, figure 6*a*). Node-pairs with 0 and 1 common alters did not show any such decreasing trend. At the end of the fifth round, the node-pairs with two common alters were significantly less dissimilar than the 0 and 1 common alter cases (sample size of node-pairs: $n_2 = 170$, $n_1 = 464$, $n_0 = 284$; two-tailed test; 2 versus 0 common alter(s): $t(452) = -2.962$, $p < 0.01$; 2 versus 1 common alter(s): $t(632) = -2.788$, $p < 0.02$; Bonferroni-corrected p-values; see electronic supplementary material, table S6). In the static condition, the alters were the same, but the network remained fixed. All of the three comparison cases of 0, 1 and 2 common alter node-pairs showed steadily decreasing dissimilarity ($p < 0.001$ for the negative slope in all three cases), but there was no difference among the three comparison cases ($p > 0.05$, figure 6*b*). In the solo condition, there was no stimuli, and the semantic dissimilarity did not have any systematic trend ($p = 0.68$ for the slope, figure 6*c*).

This shows that as the rounds progressed, the ideas of egos who followed the exact same alters increasingly grew similar. Importantly, this effect is different from groupthink [42], where the desire for harmony or conformity results in consensus

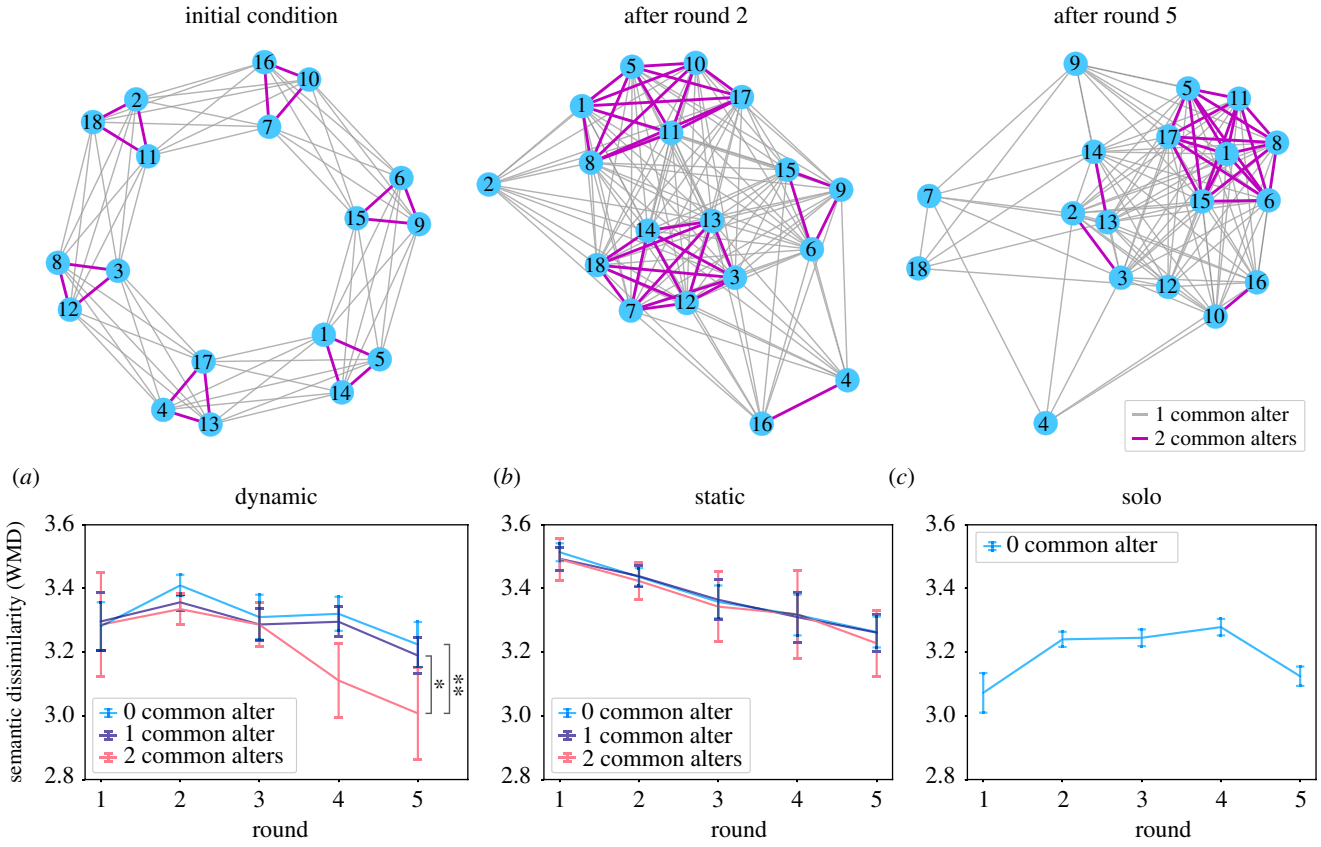

Figure 6. (Top row) Example one-mode projections of the bipartite networks on the ego-nodes from trial 4. In the static condition, the initial condition remained fixed. (Bottom row) Semantic dissimilarity (Word Mover's Distance, WMD) between the idea-sets of node-pairs are shown for (*a*) dynamic, (*b*) 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 (two-tailed tests, 2 versus 0 common alter(s): $t(452) = -2.962$, $p < 0.01$; 2 versus 1 common alter(s): $t(632) = -2.788$, $p < 0.02$). Whiskers denote 95% CI. $^{*}p < 0.05$, $^{**}p < 0.01$ (Bonferroni-corrected).

among group members. In our case, the egos acted independently without the knowledge of other egos' ideas, yet became increasingly similar to the co-followers of the same alter.

Individual creative performance comparisons among various study conditions. We analyse the individual creative performances in various study conditions. The participants in both the dynamic and static conditions significantly outperformed the solo participants in the total number of non-redundant ideas (two-tailed test; dynamic $n_d = 108$, static $n_s = 108$, solo $n_c = 36$; dynamic (d) versus solo (c): $t(142) = 2.7$, $p < 0.03$; static (s) versus solo (c): $t(142) = 2.898$, $p < 0.02$; all p-values Bonferroni-corrected; figure 7*a*; electronic supplementary material, table S7). However, there was no significant difference between the dynamic and static conditions (two-tailed test, $p > 0.05$).

The dynamic egos showed significantly higher average novelty ratings than the static egos (two-tailed test, $p < 0.05$), but after Bonferroni correction, the difference was no longer significant (figure 7*b*). The other condition-pair comparisons (solo versus dynamic and solo versus static) did not show any significant difference ($p > 0.05$). The dynamic egos significantly outperformed the static egos in two trials out of six (two-tailed test, $p < 0.05$).

The creativity quotient metric did not show any significant difference between any of the condition-pairs ($p > 0.05$ for each condition pair, figure 7*c*). Electronic supplementary material, figures S5–S7 summarize trialwise results.

Thus, at an individual level, we observe no consistent and systematic 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 performance benefits over the static condition [16]. However, in typical collective intelligence tasks, people can imitate their peers' answers to get closer to the 'correct' responses. In our study, the task encouraged *open-ended* ideation, and none of the three creativity metrics captured any systematic benefit of the dynamic condition.

Simulation model for the observed processes. We formulate an agent-based simulation model for the observed processes. The purpose is twofold: *first*, to capture the network dynamics and the associated micro- and macro-level stimulation effects in a simple and tractable manner. This helps to provide clarity to the interplay between the network and cognitive processes in the system. *Second*, to explore corner case outcomes of various simulation variables, e.g. the rewiring probability, inter-ego similarity, cognitive stimulation function and network size. These result in insights that enrich the empirical findings. The simulation model, analysis and results are elaborated in the electronic supplementary material. Below, we summarize the key ideas.

We first generate idea-sets for alter-agents such that some alters have more non-redundant ideas than others (capturing popular and unpopular alters). Starting from the same initial network structure as the empirical set-up, the ego-agents gradually rewire their connections to increasingly follow the popular alters. We use a rewiring probability parameter P_r, where $P_r = 0$ (no rewiring) results in the initial network structure, and $P_r = 1$ (full rewiring) results in all of the egos following only the top performing alters. Thus, P_r captures the network's temporal evolution.

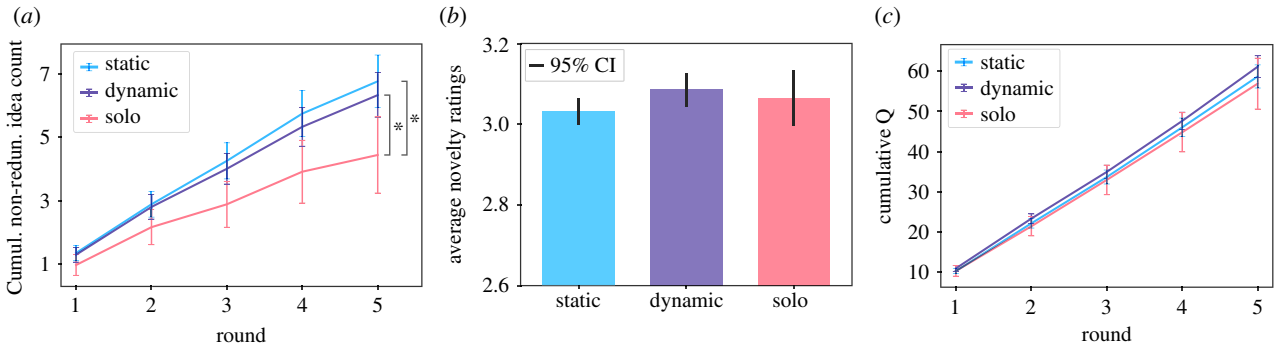

Figure 7. Individual-level comparisons of (*a*) cumulative non-redundant idea counts, (*b*) average novelty ratings and (*c*) cumulative creativity quotients among various study conditions. The static and dynamic conditions significantly outperformed the solo condition in the total number of non-redundant ideas (two-tailed test; dynamic versus solo: $t(142) = 2.7$, $p < 0.03$; static versus solo: $t(142) = 2.898$, $p < 0.02$). Whiskers denote 95% CI. *Bonferroni-corrected $p < 0.05$.

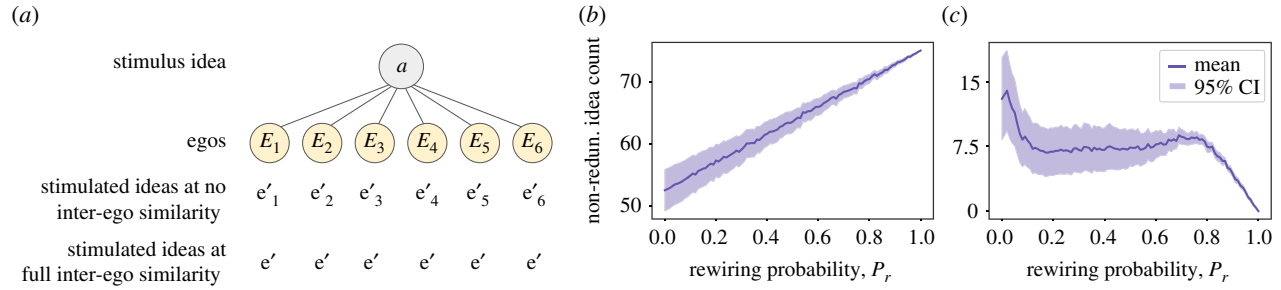

Figure 8. (*a*) Illustration of inter-ego similarity. Egos E_1 to E_6 view the same stimulus *a*, and independently generate a single idea each. At no inter-ego similarity, each of the stimulated ideas is mutually different. At full inter-ego similarity, all of the stimulated ideas become the same. (*b*,*c*) The total non-redundant idea counts in a network with $m = 18$ alters and $n = 54$ egos. As the rewiring probability P_r increases, more egos follow the same highly popular alters in the network. In the absence of inter-ego similarity (*b*), an increase in P_r leads to a better generation of non-redundant ideas. At full inter-ego similarity (*c*), $P_r = 1$ makes all the stimulated ideas mutually redundant.

We simulate the consequent network-driven and cognitive effects on the egos' ideation using three key processes: (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 follow the same alters. (B) Given the stimuli set, the egos can generate stimulated ideas, which is a *cognitive* process. Our empirical results show a positive association between the rarity of the stimuli and the generation of novel ideas. To capture this, we consider three abstractions of the cognitive process, where the statistical rarity of a stimulus idea affects the number of stimulated novel ideas in linear, sub-linear and super-linear manners. (C) There can be similarities in the stimulated ideas of independently ideating egos who share the same alters. This is again a *network-driven* process, as the similarity is initiated/facilitated by the egos' similar choices of peers. We consider two extreme cases: no inter-ego similarity and full similarity, where all the ideas inspired by the same stimulus are mutually different and exactly the same, respectively. Importantly, if any of these three key processes is taken away, one cannot fully capture the insights from our empirical results.

Using the model, we explore various corner cases of network rewiring, inter-ego similarity, cognitive stimulation and network size variables. We find that if there is no inter-ego similarity whatsoever, process C loses relevance, and the cognitive stimulation mechanisms in process B become the key to the creative outcomes of the ego-agents. As P_r increases, more novel ideas are then generated due to better stimulation, and the network's creative outcomes peak at $P_r = 1$. On the other hand, at full inter-ego similarity, process B loses relevance, as all the stimulated ideas from a given stimulus become exactly the same. At $P_r = 1$, all of the egos have the same stimuli. As a result, none of the stimulated ideas are rare anymore and the net creative output of the network drops to zero. While our empirical results of inter-ego similarity lie in between these extreme cases, the simulation nonetheless clarifies the opportunities and constraints in a systematic manner (figure 8).

These insights are robust to the various cognitive stimulation functions we considered in process B. We find that as long as the cognitive stimulation function captures a positive stimulation effect of a high quality/rare stimulus, the same outcomes are observed irrespective of whether the effect is linear, sub-linear or super-linear. We further experimented with various network sizes and found the effects to be robust in networks 3, 10 and 100 times larger than our experimental ones. A detailed theoretical treatment of the model and the associated processes, for example beginning with a mean-field analysis, remain part of our future work.

3. Discussion

Social cues and heuristics are ubiquitously used by humans for navigating through societal interactions, and contribute to their immense success as a species. Here, we first explored how the connectivity patterns in a creativity-centric dynamic network adapt to people's performance cues. From six independent trials, we found that the egos' following/unfollowing patterns are governed significantly by the novelty (average novelty ratings) and statistical rarity (non-redundant idea counts) of the alters' ideas. These two metrics lead to an adjusted-$R^2 = 0.72$ in predicting the relative popularity of the alters. Perry-Smith *et al.*'s spiral model suggests that highly creative people will enjoy increased visibility in a dynamic network [8]. Our results

validate that idea and explain the relevant cues governing such tie formations. These dynamics are different from preferential attachment, since the egos were blind to the existing follower counts/popularity of the alters. The use of a bipartite network helped us keep the egos' stimuli-sets uniform and understand the link update patterns cleanly. However, pre-recording the alters' ideas made the network unidirectional, prohibiting us from testing the full spiral model.

Although the alters were passive actors, our results have implications for them as well. Consider social media influencers, who act as third-party endorsers and shape audience perception through blogs, tweets and other social media channels. The rise of such micro-celebrity has inspired a lot of research (e.g. [43]), and corporate brands are interested in leveraging the marketing potential therein. We find that the alters need to generate not only statistically rare ideas but also ideas of high quality and novelty to win more attention than others. This has implications for the influencers regarding how to stay relevant and ahead of others, in parallel to the Red Queen hypothesis [44].

We confirmed that following high-performing alters is associated with better creative performance of the egos. 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. In the dynamic networks, the egos showed a pattern of flocking behind the high-performing alters, thereby improving their own chances of generating novel ideas. However, there was a counter-effect that ego-pairs who followed the same alters in the dynamic condition had an increasing semantic similarity with time. At the end of the fifth round, they had significantly higher semantic similarity than ego-pairs who did not have both of their alters in common. This shows one way network dynamicity can hurt original idea generation: choosing the same stimuli may inadvertently and increasingly make people's ideas similar to each other, despite independent ideation. Importantly, these processes are driven by the egos' own choices of alters, capturing the dynamic nature of real social networks. Our simulation model captures these interplay between the network-driven and cognitive processes.

These insights can lead to research efforts towards making social networks more creatively competent. Consider academicians who follow the same popular domain-experts on social media, seeking inspiration for novel research. Indeed, high quality stimuli can help the followers generate high-quality ideas themselves. But at the same time, there can be similarities in the stimulated ideas of the followers. Social network intervention strategies will then need to strike a sweet spot to help the followers get the best out of their networks, e.g. by recommending diversified sets of followees to people. If the high-preforming alters act as gravitational force attracting and maintaining the egos' connections, the semantic similarities among the followers' contents could potentially signal when the network needs intervention for re-stimulating ideation. The same arguments can apply to offline social networks as well, e.g. in large-scale creative teams. This can be particularly pertinent for groups known to gravitate towards strong personalities. As for an ego, there are implications that seeking out high-quality 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 performance in brainstorming sessions [26,27] were not present here: there was no evaluation apprehension from peers (alters) due to asynchronous exposure, 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. 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 [26,32,33]. As expected, the dynamic and static conditions, which had stimuli from alters, enjoyed a significantly higher count of non-redundant ideas compared to the solo condition, which had no stimulus.

To make the findings robust, we repeated our experiment for six independent trials. We reported robust out-of-sample R^2 and adjusted-R^2 results. The simulation model further helped us assess the generality of the insights under various corner cases. However, our study is not without limitations. The steep expenses associated with conducting the experiment prohibited us from collecting a larger dataset. We conducted power analyses to ensure that the reported results have sufficient statistical power (see electronic supplementary material). The participants generated alternative uses of common objects, which hardly resembles real-life creative challenges. The study settings prohibited us from exploring the effects of bidirectional creative influence. Also, the study lasted for five rounds, which can be prohibitively short to demonstrate the full effects of network dynamicity. Longitudinal studies with practical creative challenges can generate further insights on the research questions.

4. Methods

4.1. Creativity metrics

We operationalized creativity using the three following metrics, based on the previous literature.

4.1.1. Non-redundant idea counts

To compute the non-redundant idea counts [34], we first rejected submissions that were infeasible and not different from the given use. Then, the round-wise idea-sets were organized/binned to collect the same ideas together. For binning, we followed the coding rules specified in [45] and in the scoring key of Guilford's Alternate Uses test.[1]

Then, we computed the statistical rarity of the ideas. 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 six alters in the 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 other words, if greater than or equal to three 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.

Two research assistants independently binned similar ideas together from the first three trials. They were shown randomly ordered anonymized ideas. The non-redundant idea counts of the participants computed based on their coding had a high agreement (Krippendorff's $\alpha = 0.85$; Spearman's $\rho = 0.92$, $p < 0.001$, 95% CI = [0.885, 0.941]). Then, the first research assistant coded the rest of the dataset, which was used in the analyses.

4.1.2. Creativity quotient

Creativity quotient, Q, accounts for both the quantity of ideas generated and the quantity of distinct categories those ideas fall into [35,36]. If the ideas of a participant are very similar, they are likely subtle variations of a small number of categories. Conversely, if they are very dissimilar, they likely touched many categories—marking better creativity.

The computation of Q uses an information theoretic measure of semantic similarity derived from WordNet [46]. Concepts appear as syn-sets (synonym sets) in WordNet, and the nouns come with an 'is a' relationship. We first remove stopwords and punctuations from the ideas, and run a spell-checker. We then split each idea into its constituting set of concepts, and replace verbs and adjectives with related nouns whenever possible. Then, we find the information content of each of those concepts. 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 [37]. Thus, infrequent concepts (e.g. leaf nodes) should hold more information than the abstracting nodes. We therefore quantify the information content, I, of a concept c as,

$$I(c) = \frac{\log\left((h(c)+1)/w\right)}{\log\left(1/w\right)} = 1 - \frac{\log\left(h(c)+1\right)}{\log\left(w\right)}, \tag{4.1}$$

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

We then determine how similar a given pool of ideas (from a participant in a given round) are. We compute the semantic similarities between every pair of concepts, c_1 and c_2, in the pool [47] using

$$\mathrm{sim}(c_1, c_2) = 1 - \left(\frac{I(c_1) + I(c_2) - 2 \times \mathrm{sim}_{\mathrm{MSCA}}(c_1, c_2)}{2}\right). \tag{4.2}$$

Here, the semantic similarity, $\mathrm{sim}(c_1, c_2)$, is a function of the amount of information the two concepts have in common, $\mathrm{sim}_{\mathrm{MSCA}}(c_1, c_2)$. This, in turn, is given by the information content of the most specific common abstraction (MSCA) that subsumes both the concepts

$$\mathrm{sim}_{\mathrm{MSCA}}(c_1, c_2) = \max_{c' \in S(c_1, c_2)} I(c'), \tag{4.3}$$

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

We then compute the multi-information, I_m, as the shared information across the 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, Q is obtained by,

$$Q = N - I_m, \tag{4.4}$$

where N is the total number of concepts in the idea pool.

4.1.3. Ratings

Each idea of the alters was rated on novelty by 36 egos in the trial, generating 38 880 ratings in total. The mean rating received by idea j is taken as the idea-level rating r_j. For participant i in round p, we compute the average novelty rating $\bar{r}_i^{(p)} = \mathrm{mean}(r_j)$ for ideas j by i in p. 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 six trials was $ICC(3, 36) = 0.945$.

Additionally, each idea of the egos and solo participants was rated by at least four raters, hired separately from Amazon Mechanical Turk, resulting in 40 320 ratings from 141 raters. For fairness, each rater focused on one round in a given trial, and rated all the 36 egos'/solo participants' ideas in that round. The raters were first given 3 min to generate ideas on that round's prompt themselves, to familiarize them with the task. Then they rated randomly presented ideas, guided by instructions and examples. We find a positive mean intra-class correlation coefficient among the raters, $ICC(3, 4) = 0.317$.

4.2. Measure of idea overlap

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

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}. \tag{4.5}$$

If $A = B = \emptyset$, we take $J(A, B) = 1$. We use the binning from non-redundant idea count computation to identify same entries in the idea-sets.

4.3. Measure of semantic dissimilarity

To measure the semantic dissimilarity of two idea-sets, we use the Word Mover's Distance [41]. We first remove stop-words and punctuation. The dissimilarity of two idea-sets is computed by the minimum Euclidean distance that the Word2Vec [40] embedded words of one idea-set need to travel to reach the embedded words of another idea-set.

Ethics. This study was approved by the Institutional Review Board at the University of Rochester, New York, USA.

Data accessibility. See https://github.com/ROC-HCI/Creativity-in-temporal-social-networks for data and code. Due to copyright protection of the creativity test, we provide processed data of the participants' ideas rather than the raw data.

Authors' contributions. R.A.B. designed the study, collected, annotated and analysed the data, and authored the manuscript. D.B. analysed the data. A.T. designed the study and annotated the data. F.C. designed the study and built the data collection platform. J.P.B. contributed to the data analysis and preparation of the manuscript. G.G. oversaw the study design, interpretation of results and preparation of the manuscript. E.H. oversaw the study design, interpretation of results and preparation of the manuscript.

Competing interests. We declare we have no competing interests.

Funding. This work was supported by NSF grant nos. IIS-1750380 and IIS-1447634, and a Google Faculty Research Award.

Acknowledgements. We appreciate the contributions of Shagun Bose, Ethan Otto and Xiaoning Guo to the project.

Endnote

[1]Copyright @ 1960 by Sheridan Supply Co., all rights reserved, published by Mind Garden, Inc., www.mindgarden.com.

Reference

1. Frank MR *et al.* 2019 Toward understanding the impact of artificial intelligence on labor. *Proc. Natl Acad. Sci. USA* **116**, 6531–6539. (doi:10.1073/pnas.1900949116)

2. Manyika J, Chui M, Miremadi M, Bughin J, George K, Willmott P, Dewhurst M. 2017 *A future that works: automation, employment, and productivity*. San Francisco, CA: McKinsey Global Institute. See https://www.mckinsey.com/featured-insights/digital-disruption/harnessing-automation-for-a-future-that-works.

3. Rotman D. 2018 Making AI into jobs. *MIT Technol. Rev.* **121**, 10–17.
4. Alabdulkareem A, Frank MR, Sun L, AlShebli B, Hidalgo C, Rahwan I. 2018 Unpacking the polarization of workplace skills. *Sci. Adv.* **4**, eaao6030. (doi:10.1126/sciadv.aao6030)
5. Baten RA, Clark F, Hoque ME. 2019 Upskilling together: how peer-interaction influences speaking-skills development online. In *8th Int. Conf. on Affective Computing and Intelligent Interaction (ACII)*, pp. 662–668. New York, NY: IEEE.
6. Wu L, Wang D, Evans JA. 2019 Large teams develop and small teams disrupt science and technology. *Nature* **566**, 378–382. (doi:10.1038/s41586-019-0941-9)
7. Kazanjian RK, Drazin R, Glynn MA. 2000 Creativity and technological learning: the roles of organization architecture and crisis in large-scale projects. *J. Eng. Tech. Manage.* **17**, 273–298. (doi:10.1016/S0923-4748(00)00026-6)
8. Perry-Smith JE, Shalley CE. 2003 The social side of creativity: a static and dynamic social network perspective. *Acad. Manage. Rev.* **28**, 89–106. (doi:10.5465/amr.2003.8925236)
9. Perry-Smith JE. 2006 Social yet creative: the role of social relationships in facilitating individual creativity. *Acad. Manage. J.* **49**, 85–101. (doi:10.5465/amj.2006.20785503)
10. Perc M, Szolnoki A. 2010 Coevolutionary games—a mini review. *BioSystems* **99**, 109–125. (doi:10.1016/j.biosystems.2009.10.003)
11. Henrich J. 2015 *The secret of our success: how culture is driving human evolution, domesticating our species, and making us smarter*. Princeton, NJ: Princeton University Press.
12. Herrmann E, Call J, Hernández-Lloreda MV, Hare B, Tomasello M. 2007 Humans have evolved specialized skills of social cognition: the cultural intelligence hypothesis. *Science* **317**, 1360–1366. (doi:10.1126/science.1146282)
13. Boyd R, Richerson PJ, Henrich J. 2011 The cultural niche: why social learning is essential for human adaptation. *Proc. Natl Acad. Sci. USA* **108**(Suppl. 2), 10 918–10 925. (doi:10.1073/pnas.1100290108)
14. Rand DG, Arbesman S, Christakis NA. 2011 Dynamic social networks promote cooperation in experiments with humans. *Proc. Natl Acad. Sci. USA* **108**, 19 193–19 198. (doi:10.1073/pnas.1108243108)
15. Szolnoki A, Perc M, Danku Z. 2008 Making new connections towards cooperation in the Prisoner's Dilemma game. *EPL (Europhysics Letters)* **84**, 50007. (doi:10.1209/0295-5075/84/50007)
16. Almaatouq A, Noriega-Campero A, Alotaibi A, Krafft P, Moussaid M, Pentland A. 2020 Adaptive social networks promote the wisdom of crowds. *Proc. Natl Acad. Sci. USA* **117**, 11 379–11 386. (doi:10.1073/pnas.1917687117)
17. Shafipour R, Baten RA, Hasan MK, Ghoshal G, Mateos G, Hoque ME. 2018 Buildup of speaking skills in an online learning community: a network-analytic exploration. *Palgrave Commun.* **4**, 63. (doi:10.1057/s41599-018-0116-6)
18. Jiang H, Zhang QP, Zhou Y. 2018 Dynamic creative interaction networks and team creativity evolution: a longitudinal study. *J. Creative Behav.* **52**, 168–196. (doi:10.1002/jocb.141)
19. Armano G, Javarone MA. 2017 The beneficial role of mobility for the emergence of innovation. *Sci. Rep.* **7**, 1–8. (doi:10.1038/s41598-017-01955-2)
20. Javarone MA. 2017 Solving optimization problems by the public goods game. *Eur. Phys. J. B* **90**, 171. (doi:10.1140/epjb/e2017-80346-6)
21. Siangliulue P, Arnold KC, Gajos KZ, Dow SP. 2015 Toward collaborative ideation at scale: leveraging ideas from others to generate more creative and diverse ideas. In *Proc. of the 18th ACM Conf. on Computer Supported Cooperative Work & Social Computing*, pp. 937–945. New York, NY: ACM.
22. Chan J, Dang S, Dow SP. 2016 Comparing different sensemaking approaches for large-scale ideation. In *Proc. of the 2016 CHI Conf. on Human Factors in Computing Systems*, pp. 2717–2728. New York, NY: ACM.
23. Nijstad BA, Diehl M, Stroebe W. 2003 Cognitive stimulation and interference in idea generating groups. In *Group creativity: innovation through collaboration* (eds PB Paulus, BA Nijstad), pp. 137–159. New York, NY: Oxford University Press.
24. Paulus PB, Brown VR. 2007 Toward more creative and innovative group idea generation: a cognitive-social-motivational perspective of brainstorming. *Social Pers. Psychol. Compass* **1**, 248–265. (doi:10.1111/j.1751-9004.2007.00006.x)
25. Brown VR, Paulus PB. 2002 Making group brainstorming more effective: recommendations from an associative memory perspective. *Curr. Dir. Psychol. Sci.* **11**, 208–212. (doi:10.1111/1467-8721.00202)
26. Paulus P. 2000 Groups, teams, and creativity: the creative potential of idea-generating groups. *Appl. Psychol.* **49**, 237–262. (doi:10.1111/1464-0597.00013)
27. Dennis A, Williams M. 2003 Electronic brainstorming. In *Group creativity: innovation through collaboration* (eds PB Paulus, BA Nijstad), pp. 160–178. New York, NY: Oxford University Press.
28. Burt RS. 2004 Structural holes and good ideas. *AJS* **110**, 349–399. (doi:10.1086/421787)
29. Zhou J, Shin SJ, Brass DJ, Choi J, Zhang ZX. 2009 Social networks, personal values, and creativity: evidence for curvilinear and interaction effects. *J. Appl. Psychol.* **94**, 1544. (doi:10.1037/a0016285)
30. Kozbelt A, Beghetto RA, Runco MA. 2010 Theories of creativity. In *The Cambridge handbook of creativity* (eds J Kaufman, R Sternberg), pp. 20–47. Cambridge Handbooks in Psychology. Cambridge: Cambridge University Press. (doi:10.1017/CBO9780511763205.004)
31. Guildford J, Christensen P, Merrifield P, Wilson R. 1978 *Alternate uses: manual of instructions and interpretation*. Orange, CA: Sheridan Psychological Services.
32. Nijstad BA, Stroebe W. 2006 How the group affects the mind: a cognitive model of idea generation in groups. *Pers. Soc. Psychol. Rev.* **10**, 186–213. (doi:10.1207/s15327957pspr1003_1)
33. Bechtoldt MN, De Dreu CK, Nijstad BA, Choi HS. 2010 Motivated information processing, social tuning, and group creativity. *J. Pers. Soc. Psychol.* **99**, 622. (doi:10.1037/a0019386)
34. Oppezzo M, Schwartz DL. 2014 Give your ideas some legs: the positive effect of walking on creative thinking. *J. Exp. Psychol.: Learn. Memory Cogn.* **40**, 1142. (doi:10.1037/a0036577)
35. Bossomaier T, Harré M, Knittel A, Snyder A. 2009 A semantic network approach to the creativity quotient (CQ). *Creativity Res. J.* **21**, 64–71. (doi:10.1080/10400410802633517)
36. Snyder A, Mitchell J, Bossomaier T, Pallier G. 2004 The creativity quotient: an objective scoring of ideational fluency. *Creativity Res. J.* **16**, 415–419. (doi:10.1080/10400410409534552)
37. Seco N, Veale T, Hayes J. 2004 An intrinsic information content metric for semantic similarity in WordNet. In *Proc. of the 16th European Conf. on Artificial Intelligence*, *ECAI 2004. Valencia, Spain, 22–27 August,* p. 1089. Amsterdam, The Netherlands: IOS Press.
38. Wobbrock JO, Findlater L, Gergle D, Higgins JJ. 2011 The aligned rank transform for nonparametric factorial analyses using only ANOVA procedures. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems, Vancouver, BC, May*, pp. 143–146. New York, NY: ACM.
39. Kay M, Wobbrock JO. 2019 ARTool: aligned rank transform for nonparametric factorial ANOVAs. R package version 0.10.6.9000. See https://github.com/mjskay/ARTool.
40. Mikolov T, Chen K, Corrado G, Dean J. 2013 Efficient estimation of word representations in vector space. (https://arxiv.org/abs/1301.3781)
41. Kusner M, Sun Y, Kolkin N, Weinberger K. 2015 From word embeddings to document distances. In *Proc. of the 32nd Int. Conf. on Machine Learning, Lille, France. PMLR* **37**, 957–966.
42. Nemeth C, Nemeth-Brown B. 2003 Better than individuals. In *Group creativity: innovation through collaboration* (eds PB Paulus, BA Nijstad), vol. 4, pp. 63–84. New York, NY: Oxford University Press.
43. Khamis S, Ang L, Welling R. 2017 Self-branding, 'micro-celebrity' and the rise of social media influencers. *Celebrity Stud.* **8**, 191–208. (doi:10.1080/19392397.2016.1218292)
44. Derfus PJ, Maggitti PG, Grimm CM, Smith KG. 2008 The Red Queen effect: competitive actions and firm performance. *Acad. Manage. J.* **51**, 61–80. (doi:10.5465/amj.2008.30708624)
45. Bouchard Jr TJ, Hare M. 1970 Size, performance, and potential in brainstorming groups. *J. Appl. Psychol.* **54**(1p1), 51. (doi:10.1037/h0028621)
46. Miller GA. 1995 WordNet: a lexical database for English. *Commun. ACM* **38**, 39–41. (doi:10.1145/219717.219748)
47. Jiang JJ, Conrath DW. 1998 Semantic similarity based on corpus statistics and lexical taxonomy. In *Proc. of the Int. Conf. on Research in Computational Linguistics*, *Taipei, Taiwan, August*. See https://www.aclweb.org/anthology/O97-1002.pdf.