

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

From semantic memory to collective creativity: A generative cognitive foundation for social creativity models

Permalink

<https://escholarship.org/uc/item/4v44r3hh>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 48(0)

Authors

Ahmed, Mirza Nayeem

Baten, Raiyan Abdul

Publication Date

2026

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

From semantic memory to collective creativity: A generative cognitive foundation for social creativity models

Mirza Nayeem Ahmed (ahmed.mir@northeastern.edu)¹ & Raiyan Abdul Baten (rbaten@usf.edu)²

¹Khoury College of Computer Sciences, Northeastern University

²Bellini College of Artificial Intelligence, Cybersecurity and Computing, University of South Florida

Abstract

Simulation-based theory development has yielded powerful insights into collective performance by linking social structure to emergent outcomes, yet it has struggled to extend to collective creativity. Creativity is hard to capture purely at the social level, as novel ideas are generated through cognitive mechanisms. To address this gap, we introduce a multi-level socio-cognitive agent-based framework in which agents share a common semantic vocabulary and substrate but differ in semantic network topology. A single generative parameter tunes semantic modularity, yielding emergent individual differences in ideational breadth. When agents exchange ideation traces, two canonical social-creativity phenomena arise without being imposed: lower pre-interaction ideation overlap predicts larger stimulation gains, and shared inspiration sources induce network-level redundancy. The framework enables mechanistic theory-building about cognition and social structure in collective creativity.

Keywords: Collective creativity; agent-based modeling

Introduction

Simulation-based theory development has rich precedence in the study of collective performance in complex social systems. Agent-based computational models have clarified how collective outcomes depend on interaction topology, information flow, and exploration-exploitation trade-offs in settings such as collective wisdom and collective problem solving—from classic opinion aggregation models to rugged landscape search (Lazer & Friedman, 2007; Mahjabin & Baten, 2025). Yet one major class of human performance remains comparatively under-modeled in this tradition: *collective creativity* (Acar et al., 2023). Creative outcomes are hard to capture purely at the social level, as idea generation relies on cognitive processes unfolding within individuals. As a result, creativity resists models that operate solely on social interaction rules or assign agents “creativity scores” that sever collective dynamics from the cognitive mechanisms that generate ideas.

This gap reflects a non-trivial modeling challenge. We argue that to be useful for mechanistic theory-building, a computational model of collective creativity should bridge cognitive and social levels: *First*, a population of agents must share a common semantic grounding so that an ideation ‘prompt’ refers to the same conceptual object across agents. *Second*, agents must differ in creative ability in a way that *emerges* from cognitive mechanisms rather than being imposed as an exogenous trait. *Third*, social creativity phenomena—such as cognitive stimulation from diverse peers and redundancy

from shared sources—should be explainable by the model as a consequence of interaction. Despite extensive empirical work on peer and network effects on creativity and rich cognitive models of individual idea generation, there remains no computational framework that satisfies these constraints jointly.

To address this challenge, we introduce a multi-level, socio-cognitive agent-based model in which agents are endowed with explicit mechanisms for idea generation and embedded in a shared social environment. Strikingly, we find that a *single structural parameter* of semantic memory can serve as a generative “creativity knob,” inducing individual creative ability differences across agents. Once these agents interact socially through minimal ideation-trace exchange, canonical social-creativity effects emerge without being imposed.

By moving beyond purely social abstractions and grounding collective creativity in generative cognitive structure, this work provides a minimal yet expressive foundation for modeling creative social systems. We show how creative outcomes can be constructed, varied, and propagated from first principles in a social setting—enabling mechanistic theory-building and controlled counterfactual analysis of collective creativity.

Background

Computational models of collective performance

Agent-based computational models have been particularly effective when the modeled quantities are clearly specified and the mechanisms of social interaction are explicit. Paradigms such as DeGroot-style opinion dynamics formalize how individuals iteratively update beliefs through social influence (DeGroot, 1974), while NK landscapes model collective problem solving over rugged payoff spaces (Gomez & Lazer, 2019; Kauffman & Weinberger, 1989). These models establish a methodological template: in well-specified settings, simulation can allow inductively examining macro-level impacts of manipulating hypothetical, counterfactual system features (Homer-Dixon et al., 2013; Lazer et al., 2009).

Collective creativity, however, does not trivially fit this template. Early computational attempts at collective creativity models assigned creativity scores to agents with hard-coded redundancy penalties (Baten et al., 2020), or assigned performance outcomes randomly to capture chance effects (Yin et al., 2019). However, such assignments do not have explicit connection to the underlying cognitive processes, limiting the models’ usefulness in theory-building and hypothesis generation about the interactions between socio-cognitive processes.

Creative cognition as semantic search

At the same time, cognitive science offers rich mechanistic accounts of individual creativity (Kenett, 2023; Kenett et al., 2018; Kumar, 2021). A prominent class of models represents semantic memory as a network of concepts and treats idea generation as a stochastic search process over this structure, such as spreading activation or random walks (Abbott et al., 2015; Avery & Jones, 2018; Baronchelli et al., 2013; Collins & Loftus, 1975; Kenett & Austerweil, 2016; Lundin et al., 2023). The classic associative theory frames novelty as the ability to move beyond obvious, local associations to form non-obvious connections among more remote concepts (Brown & Paulus, 2002; Mednick, 1962). In this view, ideation often begins with easily accessible neighbors of a prompt concept, and novelty emerges when retrieval traverses farther regions of semantic memory and recombines distant elements (Runco, 2014).

This perspective has been strengthened by empirical work linking creative ability to properties of semantic organization (Benedek et al., 2017; Campidelli et al., 2026; Chen et al., 2025; Herault et al., 2024). For instance, highly creative individuals enjoy semantic network structures that facilitate access to distant concept regions during retrieval (Kenett et al., 2014; Siew et al., 2019). These accounts, however, are predominantly individual-level: they do not address how cognitively grounded agents interact in a shared semantic space, exchange inspiration, or collectively generate novelty.

Peer and social network effects on creativity

Empirical work across social and network science bodies emphasize that creativity is shaped by exposure to others' ideas, collaboration structure, and the diversity of social inputs (Baten et al., 2021; Sosa, 2011; Uzzi & Spiro, 2005; Van Dijk et al., 2012; Vasudeva et al., 2013). For instance, relationship strength, position and external ties are known to influence creative performance (Perry-Smith, 2006; Porter et al., 2020). Non-redundant contacts and cross-boundary interaction ("structural holes") can provide informational and ideation advantages (Burt, 2004; Todo et al., 2016). At the same time, shared sources and repeated exposure can lead to redundancy in what people consider and produce (Baten et al., 2022; Kelty et al., 2025; Zhou et al., 2009).

Taken together, prior work provides strong socio-cognitive components but lacks an integrative modeling framework. Cognitive models specify how ideas are generated, while social science documents how peer structure shapes creative outputs. What is missing is a generative framework that connects cognitive mechanisms of idea generation to population-level collective ideation dynamics—a gap we address in this paper.

Constructing Creative Agents with Shared Semantic Grounding

We first build a population of cognitively grounded agents who can (i) ideate on the same prompts in a common conceptual space, and (ii) differ systematically in creative ability *without*

having to be assigned creativity as an exogenous trait. Achieving (i) requires two distinct forms of semantic grounding: a *shared vocabulary* so that a prompt denotes the same concept across agents, and a *shared substrate environment* so that agents are comparable variants of the same semantic universe rather than of unrelated/arbitrary semantic universes. Achieving (ii) requires a structural dimension (a 'knob') of semantic memory that is theoretically linked to creative cognition and can be tuned in a controlled, interpretable way.

We represent each agent's semantic memory as an unweighted, undirected graph $G_i = (V, E_i)$, where nodes V denote concepts and edges encode associative relatedness. Shared vocabulary is enforced by holding the node set V fixed across agents. Shared substrate is enforced by generating agents as structured variants around a common backbone or substrate graph G_0 defined on the same V . The agents thus differ in their edge sets E_i , providing a means to inject individual differences in their semantic structures and subsequent creative outcomes. Our choice of G_0 and the approach to introduce idiosyncratic organization in E_i are reasoned next.

Design requirement: A structural knob tied to creative cognition

A key modeling choice in our work is that we do not assign agents creativity scores. Instead, we vary a *structural property* of semantic memory that prior work associates with creative cognition: semantic modularity. Indeed, recent empirical evidence suggests that more creative individuals tend to exhibit semantic networks with weaker modular organization (Kenett et al., 2014; Siew et al., 2019). Intuitively, highly modular semantic networks decompose into densely connected concept-communities with relatively sparse bridges across communities. Such compartmentalization can trap concept retrieval within local neighborhoods during idea generation (diminishing creativity), whereas a weaker modular structure can provide more cross-community connections to support access to remote concept regions (elevating creativity). This makes modularity an attractive candidate for inducing creative heterogeneity without changing how agents search.

Our design requirement is therefore concrete: we need a generative procedure that (a) yields a family of graphs that are comparable variants of a shared semantic universe (i.e., shared vocabulary V and shared substrate G_0), and (b) provides an interpretable, controllable parameter that systematically shifts the graphs' modularity, $Q(G_i)$ without otherwise modifying the agents' cognitive processes.

A shared substrate with a single generative knob

We meet these requirements using a Watts–Strogatz (WS) small-world construction (Watts & Strogatz, 1998) as a *generative scaffold*. Prior empirical evidence suggests that semantic networks show small-world properties (Kenett et al., 2014; Marupaka et al., 2012; Schilling, 2005). In particular, the WS family provides a minimal and interpretable continuum between locally clustered neighborhoods and the injection of




Figure 1: **Modularity Q decreases as the Watts–Strogatz rewiring probability p increases.** We generate a shared G_0 with $|V| = 100$ nodes and average degree $k = 4$. Points show the mean Q across 15 graphs per p ; shaded bands denote 95% bootstrap confidence intervals over graph replicates.

long-range shortcuts—mimicking the structural intuition that modularity is intended to capture.

Concretely, we first generate a shared backbone G_0 over $n = |V|$ nodes arranged in a ring, where each node is connected to its k nearest neighbors. We then instantiate each agent’s semantic network G_i by rewiring each edge of G_0 with probability $p_i \in [0, 1]$, leading to idiosyncratic associative edges E_i . Low p_i preserves local neighborhood structure, yielding strongly compartmentalized graphs; higher p_i injects long-range shortcuts that increasingly blur community boundaries. Importantly, p_i changes what concepts an agent can reach by altering the topology of associations, not by changing the agent’s search strategy, effort, or hyperparameters.

While it is reasonable to expect that injecting shortcuts should reduce community separation, whether p_i yields a *reliable*, *smooth*, and *usable* knob for modularity is not guaranteed a priori. Modularity is a global, partition-dependent statistic: the same set of shortcuts can disrupt community structure in nonlinear ways under finite size and different community-detection choices. We therefore treat smooth controllability as an empirical property of the generator and verify its existence directly. In our analyses, we operationalize modularity $Q(G_i)$ using the Newman–Girvan definition (Newman & Girvan, 2004) evaluated on a standard modularity-maximizing community partition (Clauset et al., 2004). Under this operationalization, p_i behaves as a reliable, usable knob: modularity Q_i nearly monotonically decreases as rewiring probability p_i increases (Spearman $\rho = -0.92$, $p < 10^{-131}$; Kendall $\tau = -0.78$, $p < 10^{-89}$; Figure 1). The decline is nonlinear ($\text{AIC}_{\text{quadratic}} = -171.422 < \text{AIC}_{\text{linear}} = -142.277$).

This yields an interpretable generative knob for semantic modularity: p_i induces systematic variation in semantic modularity, which in turn is theoretically linked to creative performance through its influence on the reach of memory search.

Robustness. This relationship between p and Q is robust to substrate size $|V| \in \{100, 300\}$, degree $k \in \{4, 10\}$, and 3 different seeds.

Idea generation: a fixed semantic search process

Given a semantic network G_i , an agent generates ideas by performing stochastic semantic search from a prompt concept $s \in V$. We implement retrieval as a length- T random walk on G_i , a parsimonious abstraction of spreading activation (Collins & Loftus, 1975; Mednick, 1962). Crucially, all agents share the same search process and hyperparameters. This keeps the model sharply identifiable: any systematic differences in ideation must arise from representational topology rather than agent-specific process variations.

A single walk produces a visited-node set $\mathcal{V}_i(s) \subseteq V$ and an induced trace subgraph $\mathcal{T}_i(s)$ over those visited nodes and edges. We treat $\mathcal{V}_i(s)$ as the agent’s idea trace in conceptual space: the number of distinct concepts visited, $|\mathcal{V}_i(s)|$, is a minimal proxy for ideational breadth (Hofstra et al., 2020; Uzzi et al., 2013), and the trace subgraph $\mathcal{T}_i(s)$ is the object exchanged in *social interaction settings* explored later.

Exchangeable traces for social interaction

A social creativity model requires that agents can respond to the same prompt and exchange what they generated in a commensurable form. Shared vocabulary ensures that a prompt s denotes the same concept across agents. Shared substrate ensures that agents are comparable variants of the same semantic world. Together, these grounding choices mean that ideation traces are *socially exchangeable* objects: $\mathcal{V}_i(s)$ and $\mathcal{T}_i(s)$ live in the same conceptual space for all agents and can therefore be compared (e.g., based on the extent of overlap), incorporated (e.g., augmentation of edges received through social ‘inspiration’), and analyzed at scale.

Individual Creativity Emerges from Semantic Network Structure

Before studying social effects, we verify the framework’s key precondition: creative behavior must emerge from constructed cognitive representations rather than being imposed as an agent attribute. Specifically, we test whether structural variation induced by the generative knob p_i —acting through semantic modularity $Q(G_i)$ —produces systematic differences in ideational breadth under an identical retrieval process.

Measure: ideational breadth under fixed retrieval

We instantiate a population of semantic graphs $\{G_i\}$ as structured variants of a shared substrate G_0 , as described above. Agents differ only in rewiring probability p_i (restricted to the empirically monotonic modularity regime of the WS generator), which induces variation in modularity $Q(G_i)$. For each agent and prompt $s \in V$, we run R independent length- T random walks and record the number of distinct concepts visited,

$$B_i(s) = |\mathcal{V}_i(s)|.$$

We estimate each agent’s expected ideational breadth by averaging over S sampled prompts and R walk replicates,

$$\hat{B}_i = \frac{1}{S} \sum_{k=1}^S \left(\frac{1}{R} \sum_{r=1}^R B_i(s_k) \right),$$




Figure 2: **Semantic modularity constrains exploratory access.** Each point is a semantic graph ($N_G = 500$). Higher modularity $Q(G_i)$ predicts lower expected ideational breadth ($\hat{B}_i$). We use 30 walk replicates per sampled prompt per graph. Line shows the fitted linear trend; shaded band the 95% CI.

and quantify uncertainty by bootstrapping prompts $\{s_k\}_{k=1}^S$ (with replacement) and recomputing $\hat{B}_i$.

Result: modularity predicts exploratory reach

Across $N_G = 500$ independently generated semantic graphs, modularity $Q(G_i)$ is strongly negatively associated with ideational breadth $\hat{B}_i$ under the same walk-based retrieval procedure ($r = -0.90$, 95% CI $[-0.91, -0.88]$, $p < 10^{-179}$; Figure 2). Rank-based estimates corroborate the monotonic relationship (Spearman $\rho = -0.91$, $p < 10^{-192}$; Kendall $\tau = -0.73$, $p < 10^{-131}$). A linear model indicates that higher modularity reduces breadth ($\hat{B} \sim Q$: $\beta = -28.9$, 95% CI $[-30.1, -27.7]$, $p < 10^{-179}$, $R^2 = 0.81$). A robust Theil–Sen estimator shows that the effect is not driven by outliers ($\beta_{TS} = -28.4$, 95% bootstrap CI $[-29.9, -26.9]$). Model comparison suggests curvature beyond a linear trend ($\text{AIC}_{\text{linear}} = 1441$ vs. $\text{AIC}_{\text{quadratic}} = 1369.4$), showing diminishing returns as Q approaches the low-modularity regime.

These results show that the same single knob used to generate the population—rewiring probability p_i acting through modularity—induces systematic individual differences in semantic exploration without changing how agents search. This pattern aligns with the intuition that highly modular networks trap traversal within local clusters, while cross-concept-community bridges facilitate access to concept regions that would otherwise be difficult to reach. Thus, individual creative potential emerges mechanistically from semantic topology, aligning with semantic-network accounts of creativity and providing an internal coherence check for the framework.

Robustness. This relationship between Q and $\hat{B}$ is robust to walk length $T \in \{10, 20, 30, 50\}$, substrate size $|V| \in \{100, 300\}$, degree $k \in \{4, 10\}$, and 3 seeds.

Implications for social modeling

Importantly, the ability to control *individual* creativity differences through a reliable knob p_i , combined with the *shared* semantic grounding that provide socially exchangeable objects for ideation traces ($\mathcal{V}_i(s)$ and $\mathcal{T}_i(s)$; see the previous section), gives us the minimal conditions needed to study *social creativity* phenomena without encoding them as assumptions. Because retrieval dynamics are fixed, subsequent dyadic and network-level effects cannot be attributed to tuned agent strategies or exogenous agent traits. Next, we proceed to examine whether theorized and empirically observed social creativity phenomena—cognitive stimulation from diverse peers and redundancy from shared sources—arise when agents exchange and incorporate ideation traces in the shared conceptual space.

Cognitive Stimulation Emerges from Generative Social Interaction

A central mechanism in social creativity theories is *cognitive stimulation*: exposure to a partner’s ideas can open new associative routes, resulting in ideas that one could not generate on their own—especially when the partner borrows from conceptual neighborhoods that have low overlap with one’s already-accessed regions (Baten et al., 2020; Chan et al., 2016; Dennis & Williams, 2003; Kelty et al., 2025; P. Paulus, 2000; P. B. Paulus & Brown, 2007; Siangliulue et al., 2015). We test whether cognitive stimulation emerges in our framework and whether its magnitude depends systematically on pre-interaction overlap.

Dyadic interaction protocol

We sample *ordered* exposures ($i \rightarrow j$) from the agent population constructed in the previous sections. For each ordered exposure, we sample multiple starting concepts $s \in V$ (prompts). For each ($i \rightarrow j, s$), the source agent i generates an initial ideation trace by running the fixed random-walk process on its semantic network, yielding a visited-node set $\mathcal{V}_i^{(1)}(s)$ and induced trace subgraph $\mathcal{T}_i^{(1)}(s)$. The recipient agent j independently performs the same initial walk from the same s on its own network, yielding $\mathcal{V}_j^{(1)}(s)$ and $\mathcal{T}_j^{(1)}(s)$.

We operationalize *inspiration* as exposure to (and incorporation of) the source’s trace structure. Specifically, the recipient augments its edge set with associative links present in the source trace but absent from the recipient’s current network:

$$E'_j = E_j \cup (E(\mathcal{T}_i^{(1)}(s)) \setminus E_j),$$

yielding an updated network $G'_j = (V, E'_j)$. The recipient then reruns the same random-walk ideation process from the same starting concept s on G'_j , producing a second-round visited-node set $\mathcal{V}_j^{(2)}(s)$.

For each fixed ($i \rightarrow j, s$) exposure, we run multiple independent walk iterations per round and summarize outcomes by averaging across iterations prior to statistical analysis. Each ($i \rightarrow j, s$) thus contributes a single exposure-level mean, removing within-exposure stochasticity from the random walk




Figure 3: **Lower pre-interaction overlap predicts larger stimulation gains.** We plot the fixed-effects partial relationship between pre-interaction overlap and stimulation benefit after residualizing both variables by ordered-pair fixed effects and prompt fixed effects. We partition residualized overlap into 100 quantile bins and plot, for each bin, the mean residualized overlap against the mean residualized gain. The fitted line has a slope equal to the estimated β from Eq. 1.

process; remaining dependence across exposures sharing the same ordered pair or prompt is accounted for in the regression analysis via fixed effects and clustered standard errors.

Measures

We quantify (i) *pre-interaction overlap* between the source and recipient’s initial ideation, and (ii) the recipient’s *stimulation benefit* after incorporating the source trace. Pre-interaction overlap is measured as Jaccard similarity in round one:

$$\text{Overlap}_{i \rightarrow j}(s) = \frac{|\mathcal{V}_i^{(1)}(s) \cap \mathcal{V}_j^{(1)}(s)|}{|\mathcal{V}_i^{(1)}(s) \cup \mathcal{V}_j^{(1)}(s)|}.$$

Stimulation benefit is measured as the number of *new* concepts the recipient accesses in round two relative to its round-one:

$$\Delta_{i \rightarrow j}(s) = |\mathcal{V}_j^{(2)}(s) \setminus \mathcal{V}_j^{(1)}(s)|.$$

For each exposure $(i \rightarrow j, s)$, we compute $\text{Overlap}_{i \rightarrow j}(s)$ and $\Delta_{i \rightarrow j}(s)$ for each walk iteration and then average over iterations, yielding per-exposure means used in statistical tests.

Result: overlap predicts marginal benefit

Because each ordered pair $(i \rightarrow j)$ reappears across many prompts, exposure-level observations are not independent: the same pair shares stable structural properties that could jointly shape both overlap and gain. To isolate the within-pair relationship, we estimate a two-way fixed-effects model with ordered-pair and prompt fixed effects:

$$\Delta_{i \rightarrow j}(s) = \beta \text{Overlap}_{i \rightarrow j}(s) + \alpha_{i \rightarrow j} + \gamma_s + \varepsilon_{i \rightarrow j, s}, \quad (1)$$

where $\alpha_{i \rightarrow j}$ absorbs all pair-specific time-invariant factors (including both agents’ fixed semantic structures and directionality), and γ_s absorbs prompt-specific difficulty or baseline

exploration differences shared across pairs. We report standard errors clustered by ordered pair to account for residual dependence across prompts within the same $(i \rightarrow j)$.

We observe a clear negative relationship between pre-interaction overlap and stimulation benefit: $\hat{\beta} = -5.8$ (95% CI $[-6.7, -4.9]$, $p < 10^{-34}$; 4,998 exposures across 500 ordered pairs; $R^2 = 0.50$; Figure 3). Interpreted on the overlap scale, a 0.10 increase in Jaccard overlap corresponds to 0.58 fewer newly accessed concepts on average, holding the ordered pair and prompt fixed. Thus, exposures in which the source and recipient initially traverse less-overlapping conceptual regions yield larger gains in newly accessed concepts after inspiration, whereas higher-overlap exposures yield smaller gains.

This effect is mechanistic, as it is not produced by any reward term, diversity bonus, or adaptive strategy. Agents run the same fixed search process before and after inspiration, and the update rule only adds novel associative links from a partner’s trace. Cognitive stimulation arises because low-overlap partners contribute bridges into genuinely new semantic regions, whereas high-overlap partners contribute fewer structurally novel links capable of redirecting subsequent exploration.

Robustness. The negative overlap–benefit relation persists under variations in walk length $T \in \{10, 20, 30\}$ and 3 seeds.

Social Network-Level Redundancy Emerges from Shared Sources of Inspiration

Cognitive stimulation is often framed as a dyadic benefit, but collective creativity unfolds in social networks where multiple recipients may draw inspiration from the same source. This creates a structural tension: a highly creative source can help many individuals expand their search, yet the same shared source can also induce convergence among recipients—yielding *redundancy* as a network-level externality (Baten et al., 2020; Keltz et al., 2025). Rather than encoding this tension through explicit costs, we test whether redundancy arises under a minimal model: when recipients incorporate traces in a shared conceptual space, does *sharing an inspiration source* increase recipient–recipient overlap even when recipients do not interact or adapt to each other?

Controlled design: shared vs. independent sources

We compare two conditions that differ only in whether two recipients share an inspiration source. We sample two *sources* to be highly creative under our generative construction (i.e., with bottom 20% modularity) and two *recipients* to be low-creative (i.e., top 20% modularity). All agents ideate on the same prompt $s \in V$ using the same fixed random-walk process. Recipients then incorporate edges derived from the source trace(s) and ideate again, yielding post-inspiration traces.

Independent-source (CONTROL). Recipient a incorporates edges derived from source h_1 ’s trace, while recipient b incorporates edges derived from source h_2 ’s trace.

Shared-source (TRIAD). Both recipients incorporate edges derived from the same source h_1 ’s trace, holding all other aspects constant.




Figure 4: **Shared inspiration increases redundancy in recipients’ semantic exploration.** Bars show mean post-inspiration recipient–recipient overlap across 495,000 matched instances; whiskers denote 95% CI computed with SEs clustered by inspiration source; *** $p < 0.001$.

This construction yields matched TRIAD–CONTROL comparisons for each (h_1, a, b) instance, differing only in whether the inspiration source is shared.

Measures

We quantify redundancy as similarity between the two recipients’ *post-inspiration* visited-node sets. Let $\mathcal{V}_a^{(2)}(s)$ and $\mathcal{V}_b^{(2)}(s)$ denote recipients’ second-round visited-node sets for prompt s . We measure recipient–recipient overlap as:

$$R(s) = \frac{|\mathcal{V}_a^{(2)}(s) \cap \mathcal{V}_b^{(2)}(s)|}{|\mathcal{V}_a^{(2)}(s) \cup \mathcal{V}_b^{(2)}(s)|}.$$

For each matched instance, we define the paired redundancy difference as $\Delta(s) = R_{\text{TRIAD}}(s) - R_{\text{CONTROL}}(s)$.

Result: shared inspiration increases redundancy

Because many instances share the same inspiration source, comparisons are not independent within source; we therefore base inference on a source-clustered estimator:

$$\Delta_h = \alpha + \varepsilon_h,$$

where each observation is a matched instance and standard errors are clustered by source h . We find that across 495,000 matched instances, shared inspiration sources in TRIAD significantly increases redundancy in recipients’ semantic exploration than CONTROL ($t = 55.9$, $p < 10^{-300}$; Figure 4). The paired mean difference is $\Delta = 0.026$ (95% CI [0.025, 0.027]), corresponding to a large effect size (Cohen’s $d_z = 1.4$).

Mechanistically, this redundancy is not imposed by any explicit penalties. Recipients do not adapt to one another. The increased overlap arises solely from shared structural updates: a common source introduces similar new bridges into both recipients’ semantic networks, biasing subsequent walks toward overlapping regions.

Robustness. The redundancy effect persists under variations in walk length $T \in \{10, 20, 30\}$ and 3 different seeds.

Discussion

This paper’s core contribution is to make collective creativity formally tractable under a *generative* view of cognition, instead of modeling creativity as an exogenous agent trait assigned at the social level. We show that a single semantic structural knob can induce creative ability differences across agents. This way, canonical social-creativity effects become well-posed and arise organically. Long-standing claims in social-creativity theory then become mechanistically testable questions: when does exposure expand exploration, when does it induce convergence, and how do these effects depend on the structures of minds and of social networks?

The framework primarily opens a counterfactual and hypothesis-generation space. Because agent differences are induced by a controlled structural dimension of semantic memory, one can systematically vary cognitive mechanisms that shape how far individuals are willing or able to explore. Demographic and developmental shifts (e.g., age-related changes (Campidelli et al., 2026; Ramscar et al., 2017; Wulff et al., 2022)), as well as constructs such as absorptive capacity, cognitive control, and motivation (Miguélez & Moreno, 2015), can be instantiated as parameters that gate traversal and incorporation *without* abandoning representational grounding. Our present model is intentionally minimal, as all agents run the same search process and inspiration is implemented as full incorporation of a partner’s trace. But precisely because the baseline is clean, extensions become interpretable: partial uptake (absorptive capacity), selective attention to specific trace regions, decay and interference over time, or goal-biased traversal can be introduced to test when stimulation saturates, reverses, or becomes fragile.

At the social level, the same cognitive foundation enables principled questions about how interaction structure shapes creative outcomes. Tie formation dynamics (e.g., homophily on demographic or ideation success cues) can be encoded to study how teams self-assemble, how network evolution alters the stimulation–redundancy tradeoff, and when “spreading out” inspiration sources across the network outperforms repeatedly drawing from a small set of highly generative hubs (Bangash et al., 2025; Baten et al., 2026; Page, 2019). Likewise, team size and temporal exposure schedules can be varied to ask when diversity is beneficial versus when shared sources accelerate convergence (Kamal & Baten, 2025; Wu et al., 2019). While we did not model strategic partner choice, incentives, or endogenous rewiring here, these are natural next layers: the framework provides a mechanistic sandbox in which such social processes can be tested against cognitive constraints rather than assuming them as abstract payoffs.

Code and Data Availability

The analysis code and data can be found here: <https://github.com/cssai-research/gen-collective-creativity>.

Acknowledgments

A faculty startup fund from the University of South Florida supported this work.

References

- Abbott, J. T., Austerweil, J. L., & Griffiths, T. L. (2015). Random walks on semantic networks can resemble optimal foraging. *Neural Information Processing Systems Conference*, 122, 558.
- Acar, O. A., Tuncdogan, A., van Knippenberg, D., & Lakhani, K. R. (2023). Collective creativity and innovation: An interdisciplinary review, integration, and research agenda. *Journal of Management*, 01492063231212416.
- Avery, J., & Jones, M. N. (2018). Comparing models of semantic fluency: Do humans forage optimally, or walk randomly? *Proceedings of the 40th Annual Meeting of the Cognitive Science Society*, 40. <https://escholarship.org/uc/item/8tz9m7sc>
- Bangash, A. S., Veera, K., Islam, I. A., & Baten, R. A. (2025). MuseScorer: Idea originality scoring at scale. *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 19947–19965.
- Baronchelli, A., Ferrer-i-Cancho, R., Pastor-Satorras, R., Chater, N., & Christiansen, M. H. (2013). Networks in cognitive science. *Trends in Cognitive Sciences*, 17(7), 348–360.
- Baten, R. A., Aslin, R. N., Ghoshal, G., & Hoque, E. (2021). Cues to gender and racial identity reduce creativity in diverse social networks. *Scientific Reports*, 11(1), 1–10.
- Baten, R. A., Aslin, R. N., Ghoshal, G., & Hoque, E. (2022). Novel idea generation in social networks is optimized by exposure to a “Goldilocks” level of idea-variability. *PNAS Nexus*, 1(5), pgac255.
- Baten, R. A., Bagley, D., Tenesaca, A., Clark, F., Bagrow, J. P., Ghoshal, G., & Hoque, E. (2020). Creativity in temporal social networks: How divergent thinking is impacted by one’s choice of peers. *Journal of the Royal Society Interface*, 17(171), 20200667.
- Baten, R. A., Bangash, A. S., Veera, K., Ghoshal, G., & Hoque, E. (2026). AI-driven peer recommendation systems can enhance creativity in social networks. *npj Artificial Intelligence*.
- Benedek, M., Kenett, Y. N., Umdasch, K., Anaki, D., Faust, M., & Neubauer, A. C. (2017). How semantic memory structure and intelligence contribute to creative thought: A network science approach. *Thinking & Reasoning*, 23(2), 158–183.
- Brown, V. R., & Paulus, P. B. (2002). Making group brainstorming more effective: Recommendations from an associative memory perspective. *Current Directions in Psychological Science*, 11(6), 208–212.
- Burt, R. S. (2004). Structural holes and good ideas. *American Journal of Sociology*, 110(2), 349–399.
- Campidelli, L., Domanti, U., Fusi, G., Kenett, Y. N., & Agnoli, S. (2026). Creativity, the fountain of youth: Association between creativity and semantic memory networks across the lifespan. *Cognition*, 266, 106318.
- Chan, J., Dang, S., & Dow, S. P. (2016). Comparing different sensemaking approaches for large-scale ideation. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2717–2728.
- Chen, Q., Kenett, Y. N., Cui, Z., Takeuchi, H., Fink, A., Benedek, M., Zeitlen, D. C., Zhuang, K., Lloyd-Cox, J., Kawashima, R., et al. (2025). Dynamic switching between brain networks predicts creative ability. *Communications Biology*, 8(1), 54.
- Clauset, A., Newman, M. E., & Moore, C. (2004). Finding community structure in very large networks. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 70(6), 066111.
- Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407.
- DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, 69(345), 118–121.
- Dennis, A., & Williams, M. (2003). Electronic brainstorming. *Group Creativity: Innovation through Collaboration*, 160–178.
- Gomez, C. J., & Lazer, D. M. (2019). Clustering knowledge and dispersing abilities enhances collective problem solving in a network. *Nature Communications*, 10(1), 1–11.
- Herauld, C., Ovando-Tellez, M., Lebuda, I., Kenett, Y. N., Beranger, B., Benedek, M., & Volle, E. (2024). Creative connections: The neural correlates of semantic relatedness are associated with creativity. *Communications Biology*, 7(1), 810.
- Hofstra, B., Kulkarni, V. V., Galvez, S. M.-N., He, B., Jurafsky, D., & McFarland, D. A. (2020). The diversity–innovation paradox in science. *Proceedings of the National Academy of Sciences*, 117(17), 9284–9291.
- Homer-Dixon, T., Maynard, J. L., Mildenerger, M., Milkoreit, M., Mock, S. J., Quilley, S., Schröder, T., & Thagard, P. (2013). A complex systems approach to the study of ideology: Cognitive-affective structures and the dynamics of belief systems. *Journal of Social and Political Psychology*, 1(1).
- Kamal, M. M., & Baten, R. A. (2025). Beginner’s charm: Beginner-heavy teams are associated with high scientific disruption. *arXiv preprint arXiv:2509.10389*.
- Kauffman, S. A., & Weinberger, E. D. (1989). The NK model of rugged fitness landscapes and its application to maturation of the immune response. *Journal of Theoretical Biology*, 141(2), 211–245.
- Kelty, S., Baten, R. A., Proma, A. M., Hoque, E., Bollen, J., & Ghoshal, G. (2025). The innovation trade-off: How following superstars shapes academic novelty. *Nature Humanities and Social Sciences Communications*, 12(1), 1–13.

- Kenett, Y. N. (2023). Creatively searching through semantic memory structure: A short integrative review. *The Routledge International Handbook of Creative Cognition*, 160–179.
- Kenett, Y. N., Anaki, D., & Faust, M. (2014). Investigating the structure of semantic networks in low and high creative persons. *Frontiers in Human Neuroscience*, 8, 407.
- Kenett, Y. N., & Austerweil, J. L. (2016). Examining search processes in low and high creative individuals with random walks. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 38, 313–318. <https://escholarship.org/uc/item/02p112km>
- Kenett, Y. N., Levy, O., Kenett, D. Y., Stanley, H. E., Faust, M., & Havlin, S. (2018). Flexibility of thought in high creative individuals represented by percolation analysis. *Proceedings of the National Academy of Sciences*, 115(5), 867–872.
- Kumar, A. A. (2021). Semantic memory: A review of methods, models, and current challenges. *Psychonomic Bulletin & Review*, 28(1), 40–80.
- Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, 52(4), 667–694.
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., & Van Alstyne, M. (2009). Computational social science. *Science*, 323(5915), 721–723.
- Lundin, N. B., Brown, J. W., Johns, B. T., Jones, M. N., Purcell, J. R., Hetrick, W. P., O'Donnell, B. F., & Todd, P. M. (2023). Neural evidence of switch processes during semantic and phonetic foraging in human memory. *Proceedings of the National Academy of Sciences*, 120(42), e2312462120.
- Mahjabin, M. R., & Baten, R. A. (2025). The wisdom of intellectually humble networks. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 47.
- Marupaka, N., Iyer, L. R., & Minai, A. A. (2012). Connectivity and thought: The influence of semantic network structure in a neurodynamical model of thinking. *Neural Networks*, 32, 147–158.
- Mednick, S. (1962). The associative basis of the creative process. *Psychological Review*, 69(3), 220.
- Miguélez, E., & Moreno, R. (2015). Knowledge flows and the absorptive capacity of regions. *Research Policy*, 44(4), 833–848.
- Newman, M. E., & Girvan, M. (2004). Finding and evaluating community structure in networks. *Physical Review E*, 69(2), 026113.
- Page, S. E. (2019). *The diversity bonus: How great teams pay off in the knowledge economy*. Princeton University Press.
- Paulus, P. (2000). Groups, teams, and creativity: The creative potential of idea-generating groups. *Applied Psychology*, 49(2), 237–262.
- Paulus, P. B., & Brown, V. R. (2007). Toward more creative and innovative group idea generation: A cognitive-social-motivational perspective of brainstorming. *Social and Personality Psychology Compass*, 1(1), 248–265.
- Perry-Smith, J. E. (2006). Social yet creative: The role of social relationships in facilitating individual creativity. *Academy of Management Journal*, 49(1), 85–101.
- Porter, C. M., Keith, M. G., & Woo, S. E. (2020). A meta-analysis of network positions and creative performance: Differentiating creativity conceptualizations and measurement approaches. *Psychology of Aesthetics, Creativity, and the Arts*, 14(1), 50.
- Ramscar, M., Sun, C. C., Hendrix, P., & Baayen, H. (2017). The mismeasurement of mind: Life-span changes in paired-associate-learning scores reflect the “cost” of learning, not cognitive decline. *Psychological Science*, 28(8), 1171–1179.
- Runco, M. A. (2014). *Creativity: Theories and themes: Research, development, and practice*. Elsevier.
- Schilling, M. A. (2005). A “small-world” network model of cognitive insight. *Creativity Research Journal*, 17(2-3), 131–154.
- Siangliulue, P., Arnold, K. C., Gajos, K. Z., & Dow, S. P. (2015). Toward collaborative ideation at scale: Leveraging ideas from others to generate more creative and diverse ideas. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, 937–945.
- Siew, C. S., Wulff, D. U., Beckage, N. M., & Kenett, Y. N. (2019). Cognitive network science: A review of research on cognition through the lens of network representations, processes, and dynamics. *Complexity*, 2019(1), 2108423.
- Sosa, M. E. (2011). Where do creative interactions come from? The role of tie content and social networks. *Organization Science*, 22(1), 1–21.
- Todo, Y., Matous, P., & Inoue, H. (2016). The strength of long ties and the weakness of strong ties: Knowledge diffusion through supply chain networks. *Research Policy*, 45(9), 1890–1906.
- Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical combinations and scientific impact. *Science*, 342(6157), 468–472.
- Uzzi, B., & Spiro, J. (2005). Collaboration and creativity: The small world problem. *American Journal of Sociology*, 111(2), 447–504.
- Van Dijk, H., Van Engen, M. L., & Van Knippenberg, D. (2012). Defying conventional wisdom: A meta-analytical examination of the differences between demographic and job-related diversity relationships with performance. *Organizational Behavior and Human Decision Processes*, 119(1), 38–53.
- Vasudeva, G., Zaheer, A., & Hernandez, E. (2013). The embeddedness of networks: Institutions, structural holes, and innovativeness in the fuel cell industry. *Organization Science*, 24(3), 645–663.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *Nature*, 393(6684), 440–442.

- Wu, L., Wang, D., & Evans, J. A. (2019). Large teams develop and small teams disrupt science and technology. *Nature*, 566(7744), 378.
- Wulff, D. U., Hills, T. T., & Mata, R. (2022). Structural differences in the semantic networks of younger and older adults. *Scientific Reports*, 12(1), 21459.
- Yin, Y., Wang, Y., Evans, J. A., & Wang, D. (2019). Quantifying the dynamics of failure across science, startups and security. *Nature*, 575(7781), 190–194.
- Zhou, J., Shin, S. J., Brass, D. J., Choi, J., & Zhang, Z.-X. (2009). Social networks, personal values, and creativity: Evidence for curvilinear and interaction effects. *Journal of Applied Psychology*, 94(6), 1544.