



Multidisciplinary learning through collective performance favors decentralization

John Meluso^{a,1} and Laurent Hébert-Dufresne^{a,b}

Edited by Simon Levin, Princeton University, Princeton, NJ; received March 3, 2023; accepted July 17, 2023

Many models of learning in teams assume that team members can share solutions or learn concurrently. However, these assumptions break down in multidisciplinary teams where team members often complete distinct, interrelated pieces of larger tasks. Such contexts make it difficult for individuals to separate the performance effects of their own actions from the actions of interacting neighbors. In this work, we show that individuals can overcome this challenge by learning from network neighbors through mediating artifacts (like collective performance assessments). When neighbors' actions influence collective outcomes, teams with different networks perform relatively similarly to one another. However, varying a team's network can affect performance on tasks that weight individuals' contributions by network properties. Consequently, when individuals innovate (through "exploring" searches), dense networks hurt performance slightly by increasing uncertainty. In contrast, dense networks moderately help performance when individuals refine their work (through "exploiting" searches) by efficiently finding local optima. We also find that decentralization improves team performance across a battery of 34 tasks. Our results offer design principles for multidisciplinary teams within which other forms of learning prove more difficult.

team performance | problem-solving | social networks | collective intelligence | multidisciplinary teams

Multidisciplinary teams play crucial roles in academia, industry, and government (1). Scientists from different fields collaborate to solve complex problems (1). Chefs train as sauciers or pastry chefs, then combine their products into the same meals (2, 3). Engineering teams draw representatives from various backgrounds, each working on parts that integrate with the others (4, 5). But these engines of innovation face an obstacle familiar to many: Disciplinary boundaries make collaborating across fields challenging (6–11). How, then, do members of multidisciplinary teams learn from and with one another as they advance toward collective goals?

Prior work on learning in teams would suggest two possibilities. The first treats individuals' actions as "copyable" (e.g. refs. 12–21), like learning sales strategies from a fellow salesperson or learning the steps to diagnose an illness from a senior physician. Here, people learn through feedback from their own actions (individual learning, Fig. 1*A*) and by copying techniques from those around them (vicarious or social learning, Fig. 1*B*) (22–26). Tasks (27), networks (15), and learning processes (16) all affect these types of learning.

In the second perspective, team members work together concurrently or "transactively" (e.g., refs. 28–38), like software engineers developing a new app together (39) or a medical team treating patients in a busy emergency room (40). These teams learn from performance feedback while they dynamically encode, store, retrieve, and integrate different kinds of knowledge (team learning, Fig. 1*C*) (28, 30). Such teams are often described in terms of their collective intelligence, a validated quality describing a group's ability to perform a wide variety of tasks (29, 36) which likewise depends on tasks (41), interaction patterns (17, 29, 38) and skills (36).

However, these perspectives prove limited within multidisciplinary teams. Disciplinary boundaries often impede individuals from copying each other's actions or working closely together (6–11). Instead, they divide tasks into interdependent subtasks and work in parallel (42–44). Of course, members of these teams still interact with one another, affect each other's outcomes, and the team's collective outcomes. This gives rise to the "performance nonseparability problem" (45, 46): A team's collective performance obfuscates how individuals' actions affect shared outcomes, making it hard to figure out the effect of each change. Learning becomes even more challenging when many changes happen at the same time (they may not work together) or individuals have limited visibility into each other's actions (47, 48). So while extant forms of learning remain

Significance

Like chefs at a fast-moving restaurant or engineers in a multidisciplinary project, team members often complete separate, interrelated subsets of larger tasks with limited insight into the work of others. These contexts make it difficult for individuals to assess the value of their own contribution to the collective work. Our work shows that despite this obstacle, individuals can still learn from their neighbors when neighbors' actions influence collective outcomes. Though the effects are modest, we found that teams with more interactions between members perform better when refining their work while teams with fewer interactions perform better when innovating. We also found that across 34 tasks with diverse qualities, teams that decentralize coordination responsibilities outperform those that do not.

Author contributions: J.M. and L.H.-D. designed research; J.M. performed research; J.M. contributed new reagents/analytic tools; J.M. and L.H.-D. analyzed data; and J.M. and L.H.-D. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](#).

¹To whom correspondence may be addressed. Email: john.meluso@uvm.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2303568120/-/DCSupplemental>.

Published August 14, 2023.




Fig. 1. Types of learning in teams. In each panel, i is some member of a team, and j can be any other member of that team; x_i and x_j are their respective decisions. Then, f_i is a task that i can work on alone; h is a task that an entire team works on, and g_i is a task that team member i works on but that is also influenced by the actions of their network neighbors j . (A) Individuals can learn from their own actions without outside influences. (B) Individuals can also learn by copying solutions from others they interact with. (C) Teams can learn concurrently through their actions, interactions, and feedback. (D) In multidisciplinary teams, individuals can learn through mediating artifacts, like performance assessments, that provide feedback on tasks that multiple individuals contribute to. This does not necessarily require social or team learning because knowledge boundaries between disciplines and inseparable contributions to artifacts tend to impede these forms of learning between disciplines.

possible, multidisciplinary settings do not lend themselves to individual learning, social learning within the team, or team learning.

Fortunately, another learning pathway exists: mediated learning. Here, individuals learn from themselves and others through artifacts (Fig. 1D, 26, 49). For example, members of multidisciplinary teams can learn about their performance and their teammates' performance through changes in the quality of the collective output of the team. Consider how ants learn where food is by following chemical trails left by other ants (50). People learn from one another through symbolic objects they create, like books, websites, and performance assessments (26, 46, 49). Team members learn from one another through boundary objects, like sketches and databases (51–53). Such mediated interactions can significantly affect the outcomes of microbes (54), ecological webs (55), social behavior spread (56), education (57), and innovation (58), even producing larger effects than unmediated interactions (56). Mediated learning is not mutually exclusive of the other learning types as individuals and teams can learn through artifacts as well.

Mediated learning has been observed in many settings, but to our knowledge, the performance effects of this subtle mechanism have yet to be examined systematically in multidisciplinary teams. In particular, network properties produce varied performance effects depending on the type of learning (16, 56) which makes systematic analysis of network properties crucial to understanding how mediated learning affects team outcomes. To that end, this work addresses the question: How do network properties affect multidisciplinary team performance through mediated learning?

We begin by introducing neighborhood task performance, which formalizes the networked nature of mediated learning pathways and distinguishes them from other learning pathways. Through this concept, we investigate how network density and other properties affect team performance, showing three results. First, mediated learning can either help or hurt performance depending on the task and team network structure. Then, we show that densely-connected teams tend to perform better when

“exploiting” (or refining) solutions, while sparsely-connected teams perform better when “exploring” (or innovating) new solutions—a finding that echoes other types of learning (12, 59, 60). But contrary to other types of learning (61, 62), we show that decentralized teams see greater performance benefits than other structures across a diverse set of tasks. This benefit likely arises because decentralization balances the benefits of dense and sparse teams on tasks with different qualities. We close with a discussion of implications for designing teams and limitations.

Model Description

We constructed a model that represents mediated learning in multidisciplinary teams. To achieve this, we simulated performance-mediated feedback networks, team members as computational agents, and representative tasks while minimizing other forms of learning and variables with established effects. In brief, each team member independently selects actions. Individuals' actions affect performance on their own subtask but also affect their neighbors' subtask performances and the team's overall task performance. This means individuals learn from a “neighborhood” performance shaped by their own actions and the actions of their network neighbors.

Performance-Mediated Feedback Framework. The performance nonseparability problem exemplifies mediated learning in multidisciplinary teams. Here, team members perform actions that affect each other's work. In other words, individuals' performances depend on the actions of other individuals. This yields ties in a network consisting of performance dependencies between individuals. Then, when individuals receive feedback from those confounded performance assessments (as in Fig. 1D), these action–performance–feedback sequences create a network of pathways for learning. Rather than being individual learning (learning from feedback on one's own actions in isolation) or social learning (learning by observing and copying others actions, as in ref. 12), this is mediated learning because individuals learn through mutually constructed artifacts—collective performance assessments that result from multiple individuals' actions. Thus, this scenario describes performance-mediated learning in multidisciplinary teams.

In our model, then, computational agents work on a shared team-level fitness landscape representing a task (called a “task” hereafter) over time. Each team has n agents, and each agent $i \in \{1, \dots, n\}$ has a value $x_i \in [0, 1]$ that represents its decisions on how to contribute to the task. Like others (12, 15, 27, 63), we assume that a group's performance on a task can be represented as a function $h(\vec{x})$ of the agents' decisions. The context we study in this work means we cannot assume that individuals only receive individual feedback via some $f_i(x_i)$, as with individual learning, or team feedback via $h(\vec{x})$, as with team learning (Fig. 2A). Instead, agent i owns and works on a subtask g_i of the team's overall task h . The subtask is affected by the agent's own actions x_i , but not only those actions. As in the performance nonseparability problem, agent i does not have complete control over its outcome because the actions $\vec{x}_j$ of some set of other agents $j \in J_i = \{1, \dots, k_i\}$ also affect agent i 's performance on g_i .

This yields what we call a neighborhood function $g_i(x_i, \vec{x}_j)$ for each agent i . Neighborhood functions describe an agent's subtask performance as a function of the agent's own actions and the actions of other agents that affect it (Fig. 2B). This makes neighborhood functions $k_i + 1$ dimensional, with one dimension for i and one for each $j \in J_i$. Neighborhoods are not subteams, though, as there is no coordination between agents and agents'




Fig. 2. Framework of performance-mediated learning pathways describes different networks and comprehensive model construction. (A) In the framework, individuals i can learn from team performance h or their own performance f_i , but in multidisciplinary teams they often have to learn from a combined “neighborhood” performance g_i . (B) In this context, the individual’s actions x_i and their neighbors’ actions $\vec{x}_j$ affect the agent’s performance on a subtask g_i of the team’s overall task h . This makes it difficult for individuals to tell “how good” their actions are. (C) These neighborhoods of dependence can describe any performance-mediated learning pathway, from individual learning (an empty graph) to the performance feedback paths of team learning (a complete graph) and any distribution in between. (D) Our model enables us to explore different networks by giving each individual i feedback from their neighborhood g_i . Separately, our model assesses the team’s overall performance h as a function of all agents actions $\vec{x}$.

actions can affect multiple other agents. So each neighborhood function $g_i(x_i, \vec{x}_j)$ describes agent i ’s subtask contribution to the team’s performance h while acknowledging agent i ’s limited control over its outcome.

Collectively, a team’s mediated learning network with nodes i and network neighbors $j \in J_i$.^{*} This multilevel framework can describe any performance-mediated learning pathway (Fig. 2C). At one extreme, agents only affect themselves (representing individual learning, or an empty graph). As we add interactions, we can construct any undirected network distribution by adding interactions to neighborhood functions $\vec{g}$ up to and including team learning (a complete graph) in which every agent affects every other agent. In this work, we vary the neighborhood functions to study the effects of mediated learning networks on team performance (Fig. 2D).

Individuals often have different goals from the team (64). We could choose any task g_i for each agent to work on, regardless of the task h on which the team is evaluated. In this work, we assume agents work toward the same goal as the team. We do this by using tasks that are commutative, normalized, and scalable to any number of dimensions. This lets us use the same task type for both neighborhood subtasks g_i and the team task h because the operand corresponding to each agent i retains the same form and function in every subtask in which it exists and in the team task (see *SI Appendix, S1* for details).

Put another way, g_i captures “how well” agent i thinks their action contributes to the team’s overall task h . In many cases, agent i cannot tell how their own decision x_i led to their own outcome g_i (let alone the team’s outcome h) because the $\vec{x}_j$ decisions of the other agents $j \in J_i$ obscure how agent i ’s actions affected collective outcomes. For example, while one agent could make a decision that would perform well in isolation, adjacent

agents may be negatively affected by the decision. This can occur because agents work on different subtasks from one another (again, neighborhoods are not subteams) and a decision that improves one agent’s subtask performance might degrade another agent’s subtask performance. On the other hand, the decision could also improve the performance of adjacent agents j either by increasing g_j or by opening new possibilities.

Agents Description. Given this framework, we give agents four qualities. First, agents act independently and cannot coordinate future actions by sharing their intended states. We chose this because multidisciplinary teams may not be able to discern how each other’s actions will affect their own outcomes. Also, sharing information with one another is part of both social and team learning which we sought to minimize. So, agents’ only interactions are mediated through how their choices affect each other’s outcomes as described through their neighborhood performance functions $g_i(x_i, \vec{x}_j)$.

Second, agents are naive. Individual skill significantly affects team and task performance (34, 36). To minimize the effects of individual skills on team performance, agents randomly search for better values of x_i once per time step. Agents do this by drawing a single random value, evaluating that value against their prior best value, and keeping the better value of the two. This random individual experimentation cannot remove all skill as it grants agents the ability to predict how one potential value change might affect their outcome at each time step. However, this naive approach minimizes skill for our examination of mediated learning because alternative search approaches either increase individual skill (i.e., employing search strategies, multiple sampling, error checking) or necessitate adding some unmediated form of learning (i.e., coordinating team experiments) which prove challenging in multidisciplinary settings.

Third, agents are search-limited. In the tradition of March (65) and others (12, 15, 66, 67), this represents the extent to which

^{*}More completely, $g_i(x_i, \vec{x}_j)$ describes the directed edges of a network of mediated interactions, where x_i indicates the existence of a self-edge and each x_j indicates the existence of a directed edge from j to i . For simplicity, we assume that all edges in this network are reciprocated, making a network of undirected edges sufficient for this work.

individuals try to innovate by exploring new solutions and refine solutions by exploiting the current solution (whether by choice or constraints). When agents sample new values x'_i , they explore within a specified search radius r —so they look for a new value $x'_i \in [x_i - r, x_i + r]$, but lower-bounded by 0 if $x_i - r < 0$ and upper-bounded by 1 if $x_i + r > 1$. We can then examine how one's choice of search conditions affects performance, varying from local refining searches to global exploring searches.

Finally, agents are also greedy. This means that agents seek to maximize performance by accepting any new value that they expect to improve their neighborhood's performance. We chose this decision-making process for two reasons. For one, it is consistent with the incentive structures of prior lab studies and the explicit rules in simulation studies of learning in teams (12, 15, 16, 59). Second, it is a uniform performance-maximizing rule for assessing how individuals learn through mediated artifacts. Other rules are certainly possible (e.g., accepting every third value, or only improving values until they are "good enough") (68), but such rules add a variable better explored after establishing a baseline.

Simulation Description. To summarize, every turn of the model, each agent performs a naive search for a better value of x_i by selecting a random value x'_i within a given search radius $\pm r$. If the new value x'_i produces a better outcome given what they know from the team's past choices—if $g_i(x'_i, \vec{x}_j) > g_i(x_i, \vec{x}_j)$ —they accept the new value by setting $x_i = x'_i$. They do not confirm that their decision produces the intended effect (a higher-skill verification technique) but simply continue searching. Otherwise, they keep their old x_i and continue searching (see *Materials & Methods* for pseudocode).

Altogether, this construction allows us to systematically explore network properties while controlling for variables of known significance, including team size (69–72), tasks (34, 41), and the balance between innovating (exploring) and refining (exploiting) their solutions (12, 15, 65, 67). We constructed teams of different sizes—4, 9, 16, and 25 agents—that span those represented in other studies (e.g., refs. 12, 15 and 29).

We controlled for four realistic qualities of tasks as popularized by McGrath (73) and recommended for studies on collective intelligence (29, 36) and transactive memory systems (30, 74). These task qualities—generate, choose, negotiate, and execute—map to four quantities that we can measure—

respectively: exploration difficulty, exploitation difficulty, neighborhood alignment, and neighborhood interdependence (see *SI Appendix, S2* for details). We designed a battery of 34 tasks that vary in difficulty for each of these four measures to understand team performance across a diverse set of tasks, much like how studies of collective intelligence select tasks from each group (e.g., refs. 29 and 34; see Fig. 3A–E for sample tasks; see *SI Appendix, S1* for all task specifications).

We also controlled for the search radius of agents from much smaller than the smallest feature size of the tasks ($r = 0.001$), through local "exploitation" searches just smaller than the width of the smallest valley in the task landscape ($r = 0.01$), to intermediate searches that mix exploration and exploitation ($r = 0.1$), and global "exploration" searches in which agents can always explore the entire domain ($r = 1$).

For our systematic analysis of network structures, we selected 12 measures that describe different structural properties of networks related to individual connectedness, neighbors' connectedness, network efficiency, and shared connections (cf. ref. 75, see *SI Appendix, S3* for definitions). Then, we selected 18 team network configurations that vary these qualities to understand how each quality affects team performance (see Fig. 3F–J for sample networks; see *SI Appendix, S4* for network specifications). Some of the networks are common empirically like the small world (76) and preferential attachment (77); others are random graphs to explore the effects of network density; and several are idealized or idiosyncratic graphs because they allow us to explore particular qualities. Combined, the task and network types produced varied difficulties for the teams (see Fig. 3K for two of the four task measures).

Finally, we ran each combination of these choices 250 times to examine statistical differences between population means. All agents began from a random starting point during each run of the model and conducted a search at each of 25 time steps. The following section presents our results.

Results

Mediated Learning can Help or Hurt Performance. To answer our research question, we begin by considering how well teams perform on different tasks when we vary their interaction network's density. Intuitively, one might expect agents to affect the quality of solutions that neighboring agents find. For example,




Fig. 3. Example tasks, networks, and their difficulties. We put teams through different fitness landscapes representing 34 tasks while in each of 18 network configurations, including (A–E) the example tasks and (F–J) networks shown here. (K) Together, the task functions and networks varied task difficulty across four measures, including for exploration and exploitation as shown here.

one agent might randomly start with a low-scoring value, making it likely that most values that their neighbors find improve their neighborhood performances and team performance. This implies that greater density (more interactions) might give agents the flexibility to find better solutions to their portions of the problem—even when they did poorly themselves—or conversely to constrain them when their neighbors do poorly. We examined this intuition by calculating the average team performance for each combination of team size, task, search radius, network, and point in time over the 250 runs.

The results show that density affects task performance, but only in specific cases. Fig. 4 demonstrates how density can have varied effects on team performance. For some tasks, teams perform the same as one another regardless of their density (e.g., in Figs. 4A and B). Consider an intermediate search radius of $r = 0.1$, for example. On 10 of the 34 tasks (29%), the average performance of teams with any interactions (all 17 nonempty graphs) is the same as for teams with no interaction (an empty graph, only individual learning) with a 95% CI.

However, this is not always the case. Fig. 4C shows a task in which teams with any interactions between agents outperform teams with no interaction. This effect is exaggerated further in Fig. 4D, depicting a task where teams with more interaction between agents (and higher network density) find better solutions faster than those with less interaction (and lower density). Teams with interactions outperform teams with no interactions on 8 of 34 tasks (24%), again true for all 17 nonempty graphs (95% c.i.). In contrast, Fig. 4E shows one of the 8 tasks (24%) in which having more interactions decreases the likelihood of finding better solutions (all 17 graphs, 95% c.i.). Of the remaining tasks, 3 yield an advantage and 5 yield a disadvantage for varied combinations of interacting teams (see *SI Appendix, S5* for average performances for all tasks, search radii, and networks for teams of $n = 9$).

The formulas of the tasks (*SI Appendix, S1*) shed light on why teams with certain networks perform differently from one another. Of the 16 tasks on which interacting teams have a distinct advantage or disadvantage, 15 of them multiply and/or exponentiate each agent's decision variable x_i by the number of other agents k_i that the agent interacts with. The remaining task, the minimum function, strongly depends on the number of values in its operating set which is also a function of the number of interacting neighbors. Put another way, teams with different networks tended to perform differently from one another when tasks weighted agents' contributions by a network property.

In hindsight, it may seem obvious that a team's network will affect performance when tasks depend on network properties.

Nevertheless, this is realistic. For example, in a group brainstorming task (a generate task), individuals have different information depending on whether they generate ideas alone (an empty graph), with every member of a group (a complete graph), or with some subset (a distribution), with well-established repercussions for the team's performance (34, 69, 78–80). The same is true on many choose tasks (16, 35, 59, 81). Next, we explore how each network performs across our full set of tasks.

Networks Can Help or Hurt Performance across Diverse Tasks.

Recent work on collective intelligence has raised the importance of understanding how well teams perform, not just on single tasks, but across sets of diverse tasks (like in refs. 29, 34, and 36). In light of this and our research question, our next analysis explores how team performance varies with network density when averaged across our 34 tasks. We calculated the average team performance for the 18 network types, four search radii, and four team sizes across all 34 tasks. As a point of comparison, we then calculated the percent difference of these values relative to a set of baseline values—the average performance of teams that only learned individually (the empty graph).

The first result of this analysis is that network density can modestly help or hurt a team's performance depending on how much the team explores new solutions and exploits old ones. Sparse teams tend to outperform dense teams when innovating because of reduced uncertainty, while dense teams outperform sparse teams when refining due to faster optimization (see Fig. 5 for teams of $n = 9$ and three search radii; see *SI Appendix, S6* for all team sizes and radii). The performance benefits are modest—up to 3.4% for sparse graphs when exploring and 6.1% for dense graphs when exploiting.

To build intuition for this finding, recall that agents in teams with dense networks (like the complete graph) have many more connections than in teams with sparse networks. Each agent controls its one variable in their neighborhood's performance, and their neighbors' variables can change randomly each turn. So if an agent has fewer neighbors, it is more likely that any change the agent makes will produce the expected outcome. But the more neighbors an agent has, the more likely it becomes that multiple agents make changes at the same time, landing both agents in a different location on the task landscape than they expected. While these changes could be compatible (as they are when agents do not change each other's probabilities), their combined outcome is less certain on tasks tied to network properties. When exploring globally (radius = 1.0), agents in dense teams have even more




Fig. 4. Examples of team performances averaged over 250 runs for search radius $r = 0.1$. (A and B) Network structure does not affect team performance on some tasks (10 of 34). However, network structure does affect performance on others (24 of 34 tasks), typically when the tasks are a function of network properties. (C) For example, on some tasks, teams with only individual learning (an empty graph) perform worse than teams with any interactions (any other graph). (D) On others, increasing network density corresponds to increased performance, (E) while others still exhibit decreased performance with greater density.




Fig. 5. Average performances for each network and search radius. Values shown are percent differences from the averages of teams that only learn individually (empty graphs) for teams of $n = 9$ agents with the specified search radius. Error bars show 95% CIs of the differences between means. (A) Average performances are shown by network type, sorted *Top-to-Bottom* first by density, then by eigenvector centrality, a measure of how important each individual is to each other individual. (B) The same performance averages plotted by network density. Trend lines represent simple least squares regressions between density and performance, illustrating how the search radius influences the relationship between network density and performance. *SI Appendix, S6* shows similar outcomes for teams of 4, 16, and 25.

uncertain environments and experience more random changes than agents in teams with fewer connections. As a result, sparse teams outperform dense teams when innovating.

Next, consider when teams exploit locally (radius = 0.01). This radius makes agents' changes incremental: Changes they make either stay on the same peak in the task landscape or move to an adjacent peak. As agents adjust toward the summit of their local peak, the probability of jumping to adjacent peaks (if they exist in the task) goes down because fewer points within their search radius perform better than their current solution. This is not explicit coordination between agents; much like the optimization method of gradient descent, agents passively "teach" each other where good solutions are through small shifts that guide neighbors toward better outcomes for the team. At the same time, this guidance also performs a passive form of error-checking by preventing other agents from finding solutions with a lower fitness than solutions that are within the current search radius. So like before, agents in dense teams are exposed to more simultaneous changes than those in sparse teams. But this time, those changes guide agents with many connections toward local peaks faster than those with few. Sparsely-connected teams, on the other hand, have to wait for information about the task landscape to pass from agent to agent through their random searches over many turns, leaving them to navigate more slowly toward optimal solutions. Sparse teams are more likely to refine one dimension at a time, like coordinate descent optimization, while dense teams often refine multidimensionally. Consequently, dense teams outperform sparse teams when exploiting. Between these extremes, intermediate searches (radius = 0.1) exhibit mixed results as the shape of each task determines whether density or sparsity yields better outcomes (*SI Appendix, S6*).

Thus, having fewer connections, sparser layouts, or both makes it easier for multidisciplinary teams to explore novel solutions. Large changes are more likely to improve individuals' contributions in such teams while small changes slow the process

of reaching optimal team outcomes. On the other hand, teams with more connections, denser layouts, or both exploit existing solutions more effectively. In this case, small changes help the team optimize performance quickly while large changes create uncertainty throughout the team. This finding is comparable to studies of other types of learning that find the same effect (12, 59, 60). In those contexts, dense networks are "efficient" because they spread information and converge to local solutions quickly. The "inefficiency" of sparse networks gives individuals greater opportunity to find diverse solutions making them more likely to find better ones (12).

However, density does not account for all—or even most—of the performance variation with mediated learning. In Fig. 5A, we sorted the networks from *Top to Bottom* first by density. Then, we broke any ties by sorting according to the greatest mean eigenvector centrality, a measure of how central the neighbors of each node are on average (75). While these sorts mostly placed the networks in performance order for the global search results, they are not perfectly ordered, let alone the orders of the other search radii. This raises the question of what network qualities best correspond to greater team performance across tasks and search conditions. The next section addresses this question by systematically examining how different task and network qualities relate to performance.

Decentralization Improves Performance across Tasks. While our analyses have focused on network density so far, we saw that density does not account for most of the performance variation with mediated learning. Therefore, our next two analyses utilized our full dataset (all 18 network types, 34 tasks, four search radii, four team sizes, and 26 time steps for each) to identify the relative contributions of a variety of network qualities to team performance. The first analysis found the relative importances of 12 network measures through random forest regressions (82). However, random forests cannot measure whether effects are




Fig. 6. Task measure importances, likelihoods, and effect directions. Each point shows how important the indicated measure is (x-axis), and plots those importances against how likely that measure is to either have a positive effect (positive y-axis) or negative effect (negative y-axis) on team performance. Each point summarizes the fraction of valid regression cumulative effects (combining main and search radii interactions) with a positive effect. For example, exploration difficulty was negative in all 18 cumulative effects, while neighborhood interdependence was positive in 10 of 18 cumulative effects, and negative in 8 of 18.

positive or negative. So, our second analysis predicted how likely each measure is to have a positive or negative performance effect by aggregating the results of six regressions with different controls. Figs. 6 and 7 show our results.

We identified the most important network measures through two random forest regressions (*SI Appendix, S7*). Each forest included task fixed effects and controls for the four task measures, time step, team size, and search radius. The first forest included the 12 network measures that apply to fully connected networks. This comprehensive list of network measures came at the cost of excluding data from teams with networks that are not fully connected (e.g., the empty graph, some random graphs) and cannot be evaluated by three of the measures (shortest path length, diameter, and degree assortativity). The second forest excluded network measures that cannot evaluate disconnected graphs, allowing us to include the comprehensive set of all networks regardless of their connectedness. We calculated the average and SD of the importances for each measure across the two random forests to find their impacts on team performance.

Our second analysis extends Mason and Watts' (15) application of regressions to assess how well network measures predict performance. However, network measures are often highly correlated and contain similar information to one another (83). To overcome this, we performed six multivariate ordinary least squares regressions, each containing one of six combinations of measures (see *SI Appendix, S8* for details). We then combined the results of these regressions by calculating how often each measure produced either a positive or negative effect across the regressions. The six combinations consisted of a) one of three groups of task measures and b) one of two groups of network measures. For tasks, each regression could include the task measures, task fixed effects, or both task measures and fixed effects. For networks, half of the regressions used completely-connected networks with all network measures, while the other half included all networks but only network measures that are valid on disconnected graphs. All six regressions included controls for time step, team size, the square of team size, and search radius. Each regression also included search radius interaction effects for task and network measures as other studies have found that task and search objectives can moderate desirable qualities (34, 59).




Fig. 7. Network measure importances, likelihoods, and effect directions. Subfigures show measures of (A) neighbors' connectedness, (B) network efficiency, (C) shared connections, and (D) individual connectedness. Again, each point shows how important the indicated measure is (x-axis), and plots those importances against how likely that measure is to either have a positive effect (positive y-axis) or negative effect (negative y-axis) on team performance. Decentralization and having many intermediaries are tied for most important (within one SD of each other's means). NND stands for nearest neighbor degree.

The regressions gave us main effects and search radius interaction effects for task and network measures. Finally, we counted how often each measure yielded a positive cumulative effect. Each cumulative effect was built by adding the main and interaction

effects for the exploration, exploitation, and mixed search radii (radius = 0.01, 0.1, and 1.0) whenever they were statistically significant. This gave estimates of how likely each measure's effect is to be positive or negative.

The combined results of these analyses show several items of note. First, task qualities likely affect performance more than network measures, by one to two orders of magnitude (Fig. 6). Exploration difficulty was both the most likely measure to affect performance and the most important, followed by alignment between individuals' task neighborhoods (their objectives) which is tied as the second-most likely effect. Together, task qualities accounted for about 50.0% of team performance effects. Even if saying "hard tasks are hard" is not surprising, it reaffirms prior findings on task complexity's moderating effect (16, 35).

Network measures accounted for about 2.4% of team performance effects, which is 4.8% the size of the task effects. The most important group of network measures was neighbors' connectedness (Fig. 7A). Decentralization is measured through the mean of eigenvector centrality which has a normalized euclidean norm, meaning its average is maximal in all cases where nodes are equally central (true from empty to complete graphs; see *SI Appendix, S9*). That measure is tied for the most important network measure (within one SD of mean betweenness centrality, a measure of intermediarity). It also is the most likely network effect, the most likely positive effect, and is tied with alignment as the second most likely measure overall. The other measures of neighbors' connectedness (including varied centralization) fill four of the top seven network measures revealing a consistent importance and likelihood not seen in the other groups. While it may seem contradictory for variation measures to score highly alongside decentralization, it need not be. Varied centralization is measured through the SD of eigenvector centrality. Networks like the ring of cliques, the wheel, and the windmill simultaneously exhibit moderate decentralization and high variance in centralization (*SI Appendix, S9*). These networks performed modestly across most tasks demonstrating how teams can possess both qualities simultaneously.

Results for the other measures proved mixed. Network efficiency—having short paths between individuals—often proves important for other types of learning (12, 15, 16, 84). For mediated learning, we find that efficient networks have mixed effects (Fig. 7B). Having many intermediaries was tied for the most important network measure, and varied intermediarity throughout a team came in third. However, these measures did not consistently produce positive or negative effects. Given their importances, this suggests that efficiency still affects performance, but likely through higher-order interactions. Longer paths resulted in more consistent positive gains but with less importance, while measures of shared connections (Fig. 7C) were somewhat important but with inconclusive directionality. Last, high levels of individual connectedness and variation therein produced moderately consistent negative effects but with less import (Fig. 7D).

In sum, with mediated learning, decentralization is the most important and most likely network quality to improve multidisciplinary team performance across the data from all 34 tasks, followed by other qualities of network efficiency and neighbors' connectedness in line with previous studies.

Discussion

Many teams divide tasks into separate, interrelated subtasks by discipline, much like how cooks collaborate in a high-volume restaurant kitchen (2, 3). In this study, we explored how members of such teams learn from one another mediated by

their neighborhood performances even when they cannot see each other's actions or outcomes. We found that a team's interaction network can significantly affect team performance, but only on tasks that explicitly rely on network properties. The effects of network qualities are relatively small at about 5% the effect size of task qualities. Still, this proved sufficient to give sparsely-connected teams an advantage (up to 3.4%) when innovating and densely-connected teams an advantage (up to 6.1%) when refining their solutions across our set of 34 tasks. Such advantages could prove consequential in contexts where modest performance gains provide an edge over competitors, for example, but may be too costly to implement in contexts that necessitate sizeable improvements.

Furthermore, we found that decentralizing a team's network by evenly distributing coordination responsibilities among a team's members tends to improve performance across our battery of 34 tasks. As with other types of learning (12, 15, 16, 19, 67), the efficiency with which a network distributes information appears likely to affect performance, but our analyses yielded mixed results on the importance of network efficiency and other qualities of networks.

Our results provide insights into better ways for multidisciplinary teams to coordinate. Given mediated learning, teams would benefit from carefully shaping their search objectives, interaction networks, and tasks—depending on what they can control at the time—as each holds repercussions for their performance. For example, when teams can control their search conditions, teams may be more effective at refining already-complex products through incremental changes than by letting individuals run wild (Fig. 5). Likewise, it may be easier for one or two individuals to explore new solutions on their own than for the whole team to equivocate together.

When teams can control their network, moving away from centralized formal structures and toward decentralized ones (Fig. 8) may improve team outcomes across a variety of tasks, similar to other findings on social networks (85–87). This may be especially important early in team formation because interaction networks are often rigid, tending toward fixed formal structures due to lower coordination costs along those pathways (88). But this may not mean that teams should self-organize their structures, as even weakly enforced formal structures tend to outcompete such teams (88). While much remains unknown, our findings indicate that teams are likely to benefit from formally




Fig. 8. Examples of decentralized networks. The most-decentralized networks tend to evenly assign the same number of connections k to each node, as one might evenly distribute coordination responsibilities among members of a team.

coordinating decentralized patterns of interaction to improve their performance.

When networks are difficult to control or already in place, restructuring tasks could improve performance. Task qualities, like difficulty and interdependence, may affect team performance by one to two orders of magnitude more than network qualities (Fig. 6). This presents opportunities to explore methods of improving performance and reducing uncertainty through task redesign—by making task landscapes less difficult for search conditions (e.g., innovation, refinement), reshaping subtasks at strategic locations throughout the network, at particular times in an iterative search process, etc. Altogether, our findings support the growing consensus around the necessity of adaptive team learning for optimal performance (21, 35, 89–92).

Our findings also suggest ways that organizations may improve their performance. Scholars have long applied the methods we use here to study questions of organization design (12, 27, 65, 67, 88) as well as teams. Limiting a department's connections to outside departments may aid innovation, supported by famous design shop successes like the Lockheed "Skunk Works" (93). Likewise, gradual changes may help organizations improve complex products more readily than making large changes (94). And importantly, decentralizing coordination activities may improve performance more than centralizing coordination (95, 96).

Our evidence also supports that tasks can limit a team's collective intelligence (41). Mediated learning could underlie parts of team learning wherever knowledge boundaries exist within groups. For example, some teams in studies of collective intelligence divide up tasks and specialize (e.g., ref. 34) making our findings relevant to those tasks at least. We do not mean to suggest that testing teams with tasks that limit collective intelligence makes those tasks invalid; to the contrary, those tasks provide teams with opportunities to demonstrate their collective intelligence by adapting to the tasks at hand (97). Communication patterns that arise within teams—whether as the product of social perceptiveness (29, 36), virtuality (31, 98), or homophily (99)—may influence a team's collective intelligence through current task and network constructions. Thus, our findings support Graf-Drasch et al.'s (41) recommendation that we reconceive of collective intelligence as a multidimensional concept like we do with individual intelligence, including a team's ability to adjust their interaction pattern to the task.

Finally, our results may extend to performance-related networks with nonadditive effects in adjacent fields such as population genetics. For example, in some contexts, our finding that dense networks exploit better may offer an alternative search dynamic to shifting balance theory, which highlights the way in which adaptive search can be more efficient on sparser epistatic networks (100).

Just as this work conceptually benefits from simplifying qualities like individual skill and social learning, its implications are also limited by these assumptions. Our findings do not suggest how introducing varied skills or individual intelligence may moderate these results, nor how any copying or coordination that takes place in multidisciplinary teams may influence them. For example, we performed an experiment with slightly increased coordination between agents, an alternative approach to minimizing skills by granting a small amount of coordination (a form of team learning) while removing the small individual experimentation skill. In this case, we do not see strong positive effects of mediated learning, though we do see significant negative effects (SI Appendix, S10 and Fig. S17). The relative importances of network properties also shifted in this case. While decentralization was still an important positive

performer, intermediaries and clustering proved more important, albeit with less confidence in their positive effects (SI Appendix, Fig. S20). Hence, more research on team process variables of all kinds—including different search strategies—remains crucial for understanding team performance.

Our work is also limited by our selection of tasks, networks, and network measures. We sought to select tasks, networks, and measures that covered qualities of broad interest and rapport. Still, more nuanced structures and measures would aid identification of network qualities that perform well, both on specific tasks and across many tasks. Our regression analyses raised as many questions as they answered about the qualities that shape performance in the context of mediated learning. Network features interact and cannot be fully independent of one another. The result is a complex system whose output is never fully governed by one feature or mechanism alone, leaving the subject ripe for future examination.

Materials and Methods

Model Pseudocode.

```

1: Input Graph with vertices  $i \in I$ , neighbors  $j \in J_i \forall i \in I$ ,
   neighborhood objective functions  $g_i(x_i, \vec{x}_j) \forall i \in I$  where
    $\vec{x}_j = (x_{j=1}, \dots, x_{j=k_j})$ , graph objective function  $h(\vec{x})$  where
    $\vec{x} = (x_{i=1}, \dots, x_{i=n})$ , and initial positions  $x_i \forall i \in I$ .
2: Output Performance values  $g_i \forall i \in I$  and  $h$ .
3: for each timestep do
4:   for each  $i \in I$  do
5:     Draw  $\Delta x$  from  $\mathcal{U}(-\epsilon, \epsilon)$ 
6:      $x'_i := x_i + \Delta x$ 
7:      $g_i := g_i(x_i, \vec{x}_j)$ 
8:      $g'_i := g_i(x'_i, \vec{x}_j)$ 
9:   end for
10:  for each  $i \in I$  do
11:    if  $g'_i > g_i$  then
12:       $x_i := x'_i$ 
13:    end if
14:  end for
15:   $h = h(\vec{x})$ 
16: end for

```

Data, Materials, and Software Availability. The model and analysis code for this work were developed with Python 3.9.5, Numpy 1.21.6, Scipy 1.9.0, Pandas 1.4.2, NetworkX 2.8.3, Dask 2022.4.1, Statsmodels 0.13.2, Scikit-learn 1.0.2, Matplotlib 3.5.2, and Seaborn 0.11.2. Complete code is available at <https://github.com/meluso/multi-disciplinary-learning>. Simulation raw data and statistics data have been deposited in OSF (101).

ACKNOWLEDGMENTS. Our sincere thanks to James Bagrow for his time and thoughts, particularly his suggestions to include the rook's graph and the empty graph as a baseline. Special thanks to Mirta Galesic and the reviewers for manuscript feedback. Thanks also go to Jean-Gabriel Young, Melania Lavric, Kendall Fortney, Babak Heydari, and the many others who supported us on this work. J.M. and L.H.-D. are supported by Google Open Source under the Open Source Complex Ecosystems And Networks project and by the Sloan Foundation through the Vermont Research Open Source Program Office. L.H.-D. is also supported by the NSF award EPS-2019470.

Author affiliations: ^aVermont Complex Systems Center, College of Engineering & Mathematical Sciences, University of Vermont, Burlington, VT 05405; and ^bDepartment of Computer Science, College of Engineering & Mathematical Sciences, University of Vermont, Burlington, VT 05405

1. National Research Council, *Convergence: Facilitating Transdisciplinary Integration of Life Sciences, Physical Sciences, Engineering, and Beyond* (The National Academies Press, Washington, DC, 2014).
2. P. Rambourg, *Histoire de la cuisine et de la gastronomie françaises* (Place des éditeurs, 2013).
3. M. Gregoire, *Foodservice Organizations: A Managerial and Systems Approach* (Pearson, Boston, 9th edition, 2016).
4. J. R. Wertz, D. F. Everett, J. J. Puschell, *Space Mission Engineering?: The New SMAD* (Microcosm Press, 2011), p. 1033.
5. J. Meluso, J. Austin-Breneman, J. P. Bagrow, L. Hébert-Dufresne, A review and framework for modeling complex engineered system development processes. *IEEE Trans. Syst. Man Cyber.: Syst.* **52**, 7679–7691 (2022).
6. N. Levina, E. Vaast, The emergence of boundary spanning competence in practice: Implications for implementation and use of information systems. *MIS Q.* **29**, 335–363 (2005).
7. M. A. Cronin, L. R. Weingart, Representational gaps, information processing, and conflict in functionally diverse teams. *Acad. Manage. Rev.* **32**, 761–773 (2007).
8. W. C. Barley, Anticipatory work: How the need to represent knowledge across boundaries shapes work practices within them. *Org. Sci.* **26**, 1612–1628 (2015).
9. J. Kotlarsky, B. van den Hooff, L. Houtman, Are we on the same page? Knowledge boundaries and transactive memory system development in cross-functional teams. *Commun. Res.* **42**, 319–344 (2015).
10. A. C. Edmondson, J. F. Harvey, Cross-boundary teaming for innovation: Integrating research on teams and knowledge in organizations. *Hum. Res. Manage. Rev.* **28**, 347–360 (2018).
11. E. Mattarelli, F. Bertolotti, A. Prencipe, A. Gupta, The effect of role-based product representations on individual and team coordination practices: A field study of a globally distributed new product development team. *Organ. Sci.* **33**, 1423–1451 (2021).
12. D. Lazer, A. Friedmann, The network structure of exploration and exploitation. *Admin. Sci. Q.* **52**, 667–694 (2007).
13. J. Lorenz, H. Rauhut, F. Schweitzer, D. Helbing, How social influence can undermine the wisdom of crowd effect. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 9020–9025 (2011).
14. L. Hong, S. E. Page, M. Riolo, Incentives, information, and emergent collective accuracy. *Managerial Decis. Econ.* **33**, 323–334 (2012).
15. W. Mason, D. J. Watts, Collaborative learning in networks. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 764–769 (2012).
16. D. Barkoczi, M. Galesic, Social learning strategies modify the effect of network structure on group performance. *Nat. Commun.* **7**, 13109–13109 (2016).
17. E. Bernstein, J. Shore, D. Lazer, How intermittent breaks in interaction improve collective intelligence. *Proc. Natl. Acad. Sci. U.S.A.* **115**, 8734–8739 (2018).
18. W. Toyokawa, A. Whalen, K. N. Laland, Social learning strategies regulate the wisdom and madness of interactive crowds. *Nat. Hum. Behav.* **3**, 183–193 (2019).
19. D. Brackbill, D. Centola, Impact of network structure on collective learning: An experimental study in a data science competition. *PLoS One* **15**, e0237978 (2020).
20. E. Miu, N. Gulley, K. N. Laland, L. Rendell, Flexible learning, rather than inveterate innovation or copying, drives cumulative knowledge gain. *Sci. Adv.* **6**, eaaz0286 (2020).
21. V. C. Yang, M. Galesic, H. McGuinness, A. Harutyunyan, Dynamical system model predicts when social learners impair collective performance. *Proc. Natl. Acad. Sci. U.S.A.* **118** (2021).
22. A. Bandura, Human agency in social cognitive theory. *Am. Psychol.* **44**, 1175–1184 (1989).
23. J. E. Grusec, Social learning theory and developmental psychology: The legacies of Robert Sears and Albert Bandura. *Dev. Psychol.* **28**, 776–786 (1992).
24. A. Bandura, Social cognitive theory in cultural context. *Appl. Psychol.* **51**, 269–290 (2002).
25. R. L. Kendal *et al.*, Social learning strategies: Bridge-building between fields. *Trends Cognit. Sci.* **22**, 651–665 (2018).
26. D. Guillo, N. Claidière, Do guide dogs have culture? The case of indirect social learning. *Hum. Soc. Sci. Commun.* **7**, 1–9 (2020).
27. J. W. Rivkin, N. Siegelkow, Patterned interactions in complex systems: Implications for exploration. *Manage. Sci.* **53**, 1068–1085 (2007).
28. K. Lewis, D. Lange, L. Gillis, Transactive memory systems, learning, and learning transfer. *Org. Sci.* **16**, 581–598 (2005).
29. A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, T. W. Malone, Evidence for a collective intelligence factor in the performance of human groups. *Science* **330**, 686–688 (2010).
30. K. Lewis, B. Herndon, Transactive memory systems: Current issues and future research directions. *Organ. Sci.* **22**, 1254–1265 (2011).
31. D. Engel, A. W. Woolley, L. X. Jing, C. F. Chabris, T. W. Malone, Reading the mind in the eyes or reading between the lines? Theory of mind predicts collective intelligence equally well online and face-to-face. *PLoS One* **9**, e115212 (2014).
32. J. B. Barlow, A. R. Dennis, Not as smart as we think: A study of collective intelligence in virtual groups. *J. Manage. Inf. Syst.* **33**, 684–712 (2016).
33. K. A. McHugh *et al.*, Collective decision making, leadership, and collective intelligence: Tests with agent-based simulations and a field study. *Leadersh. Q.* **27**, 218–241 (2016).
34. V. Amelkin, O. Askarisichani, Y. J. Kim, T. W. Malone, A. K. Singh, Dynamics of collective performance in collaboration networks. *PLoS One* **13**, e0204547 (2018).
35. A. Almaatouq, M. Alsobay, M. Yin, D. J. Watts, Task complexity moderates group synergy. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2101062118 (2021).
36. C. Riedl, Y. J. Kim, P. Gupta, T. W. Malone, A. W. Woolley, Quantifying collective intelligence in human groups. *Proc. Natl. Acad. Sci.* **118**, e2005737118 (2021).
37. B. Ostrowski, A. W. Woolley, K. W. Haan, Translating member ability into group brainstorming performance: The role of collective intelligence. *Small Group Res.* **53**, 3–40 (2022).
38. A. W. Woolley, R. M. Chow, A. T. Mayo, C. Riedl, J. W. Chang, Collective attention and collective intelligence: The role of hierarchy and team gender composition. *Organ. Sci.* **34**, 987–1352 (2022).
39. S. Faraj, L. Sproull, Coordinating expertise in software development teams. *Manage. Sci.* **46**, 1554–1568 (2000).
40. M. A. Valentine, A. C. Edmondson, Team scaffolds: How mesolevel structures enable role-based coordination in temporary groups. *Organ. Sci.* **26**, 405–422 (2014).
41. V. Graf-Drasch, H. Gimpel, J. B. Barlow, A. R. Dennis, Task structure as a boundary condition for collective intelligence. *Personnel Psychol.* **75**, 739–761 (2021).
42. M. N. Kiggundu, Task interdependence and job design: Test of a theory. *Organ. Behav. Hum. Perform.* **31**, 145–172 (1983).
43. R. Saavedra, P. C. Earley, L. Van Dyne, Complex interdependence in task-performing groups. *J. Appl. Psychol.* **78**, 61–72 (1993).
44. R. Wageman, Interdependence and group effectiveness. *Adm. Sci. Q.* **40**, 145–180 (1995).
45. A. A. Alchian, H. Demsetz, Production, information costs, and economic organization. *Am. Econ. Rev.* **62**, 777–795 (1972).
46. J. Uribe, S. Carnahan, J. Meluso, J. Austin-Breneman, How do managers evaluate individual contributions to team production? A theory and empirical test. *Strateg. Manage. J.* **43**, 2577–2601 (2022).
47. S. Postrel, Islands of shared knowledge: Specialization and mutual understanding in problem-solving teams. *Organ. Sci.* **13**, 303–320 (2002).
48. G. Stasser, W. Titus, Hidden profiles: A brief history. *Psychol. Inquiry* **14**, 304–313 (2003).
49. A. Kozulin, “Psychological tools and mediated learning” in *Vygotsky’s Educational Theory in Cultural Context*, A. Kozulin, B. Gindis, V. S. Ageyev, S. M. Miller, Eds. (Cambridge University Press, Cambridge, UK, 2003), pp. 15–38.
50. D. E. Jackson, F. L. W. Ratnieks, Communication in ants. *Curr. Biol.* **16**, R570–R574 (2006).
51. P. R. Carli, A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organ. Sci.* **13**, 442–455 (2002).
52. A. Majchrzak, P. H. B. More, S. Faraj, Transcending knowledge differences in cross-functional teams. *Organ. Sci.* **23**, 951–970 (2011).
53. D. Nicolini, J. Mengis, J. Swan, Understanding the role of objects in cross-disciplinary collaboration. *Organ. Sci.* **23**, 612–629 (2011).
54. N. Xiao *et al.*, Disentangling direct from indirect relationships in association networks. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2109995119 (2022).
55. T. Ohgushi, Indirect interaction webs: Herbivore-induced effects through trait change in plants. *Annu. Rev. Ecol. Syst.* **36**, 81–105 (2005).
56. J. Xie *et al.*, Indirect influence in social networks as an induced percolation phenomenon. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2100151119 (2022).
57. C. R. Henrie, L. R. Halverson, C. R. Graham, Measuring student engagement in technology-mediated learning: A review. *Comput. Ed.* **90**, 36–53 (2015).
58. G. Ahuja, Collaboration networks, structural holes, and innovation: A longitudinal study. *Adm. Sci. Q.* **45**, 425–455 (2000).
59. J. Shore, E. Bernstein, D. Lazer, Facts and figuring: An experimental investigation of network structure and performance in information and solution spaces. *Organ. Sci.* **26**, 1432–1446 (2015).
60. M. Derex, R. Boyd, Partial connectivity increases cultural accumulation within groups. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 2982–2987 (2016).
61. J. N. Mell, D. van Knippenberg, W. P. van Ginkel, The catalyst effect: The impact of transactive memory system structure on team performance. *Acad. Manage. J.* **57**, 1154–1173 (2014).
62. Ö. Koçak, D. A. Levinthal, P. Puranam, The dual challenge of search and coordination for organizational adaptation: How structures of influence matter. *Organ. Sci.* **34**, 509–986 (2022).
63. L. Hong, S. E. Page, Problem solving by heterogeneous agents. *J. Econ. Theory* **97**, 123–163 (2001).
64. W. R. Scott, G. F. Davis, *Organizations and Organizing: Rational, Natural, and Open System Perspectives* (Pearson Prentice Hall, Upper Saddle River, N.J., ed. 1, 2007).
65. J. G. March, Exploration and exploitation in organizational learning. *Organ. Sci.* **2**, 71–87 (1991).
66. A. K. Gupta, K. G. Smith, C. E. Shalley, The interplay between exploration and exploitation. *Acad. Manage. J.* **49**, 693–706 (2006).
67. C. Fang, J. Lee, M. A. Schilling, Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning. *Organ. Sci.* **21**, 625–642 (2010).
68. H. A. Simon, Bounded rationality and organizational learning. *Organ. Sci.* **2**, 125–134 (1991).
69. T. J. Bouchard Jr, M. Hare, Size, performance, and potential in brainstorming groups. *J. Appl. Psychol.* **54**, 51 (1970).
70. R. B. Gallupe *et al.*, Electronic brainstorming and group size. *Acad. Manage. J.* **35**, 350–369 (1992).
71. P. J. Lamberson, S. E. Page, Optimal forecasting groups. *Manage. Sci.* **58**, 805–810 (2011).
72. I. De Vincenzo, I. Giannoccaro, G. Carbone, P. Grigolini, Criticality triggers the emergence of collective intelligence in groups. *Phys. Rev. E* **96**, 022309 (2017).
73. J. E. McGrath, *Groups: Interaction and Performance* (Prentice-Hall Englewood Cliffs, NJ, 1984), vol. 14.
74. D. G. Bachrach *et al.*, Transactive memory systems in context: A meta-analytic examination of contextual factors in transactive memory systems development and team performance. *J. Appl. Psychol.* **104**, 464–493 (2019).
75. M. Newman, *Networks* (Oxford University Press, 2018).
76. D. J. Watts, S. H. Strogatz, Collective dynamics of “small-world” networks. *Nature* **393**, 440–440 (1998).
77. A. L. Barabási, R. Albert, Emergence of scaling in random networks. *Science* **286**, 509–512 (1999).
78. D. W. Taylor, P. C. Berry, C. H. Block, Does group participation when using brainstorming facilitate or inhibit creative thinking? *Adm. Sci. Q.* **3**, 23–47 (1958).
79. B. Mullen, C. Johnson, E. Salas, Productivity loss in brainstorming groups: A meta-analytic integration. *Basic Appl. Soc. Psychol.* **12**, 3–23 (1991).
80. N. L. Kerr, R. S. Tindale, Group performance and decision making. *Annu. Rev. Psychol.* **55**, 623–655 (2004).
81. A. Kittur, B. Lee, R. E. Kraut, “Coordination in collective intelligence: The role of team structure and task interdependence” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Association for Computing Machinery, New York, NY, USA, 2009), <https://doi.org/10.1145/1518701.1518928>, pp. 1495–1504.
82. L. Breiman, Random forests. *Mach. Learn.* **45**, 5–32 (2001).
83. S. Oldham *et al.*, Consistency and differences between centrality measures across distinct classes of networks. *PLoS One* **14**, e0220061 (2019).
84. W. A. Mason, A. Jones, R. L. Goldstone, Propagation of innovations in networked groups. *J. Exp. Psychol.: General* **137**, 422–433 (2008).
85. D. L. Rulke, J. Galaskiewicz, Distribution of knowledge, group network structure, and group performance. *Manage. Sci.* **46**, 612–625 (2000).

- 86.J. Becker, D. Brackbill, D. Centola, Network dynamics of social influence in the wisdom of crowds. *Proc. Natl. Acad. Sci. U.S.A.* **114**, E5070–E5076 (2017).
- 87.F. Xu, L. Wu, J. Evans, Flat teams drive scientific innovation. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2200927119 (2022).
- 88.J. Clement, P. Puranam, Searching for structure: Formal organization design as a guide to network evolution. *Manage. Sci.* **64**, 3469–3970 (2018).
- 89.D. A. Levinthal, Adaptation on rugged landscapes. *Manage. Sci.* **43**, 934–950 (1997).
- 90.C. S. Burke, K. C. Stagl, E. Salas, L. Pierce, D. Kendall, Understanding team adaptation: A conceptual analysis and model. *J. Appl. Psychol.* **91**, 1189–1207 (2006).
- 91.S. K. Baard, T. A. Rench, S. W. J. Kozlowski, Performance adaptation: A theoretical integration and review. *J. Manage.* **40**, 48–99 (2014).
- 92.J. S. Christian, M. S. Christian, M. J. Pearsall, E. C. Long, Team adaptation in context: An integrated conceptual model and meta-analytic review. *Organ. Behav. Hum. Decis. Process.* **140**, 62–89 (2017).
- 93.C. L. K. Johnson, M. Smith, Kelly: More than My Share of It All (1985).
- 94.C. E. Helfat, S. G. Winter, Untangling dynamic and operational capabilities: Strategy for the (N)ever-changing world. *Strateg. Manage. J.* **32**, 1243–1250 (2011).
- 95.J. G. Matusik, R. L. Mitchell, N. A. Hays, S. Fath, J. Hollenbeck, The highs and lows of hierarchy in multiteam systems. *Acad. Manage. J.* **65**, 1571–1592 (2022).
- 96.V. Venkataramani, C. Tang, When does external knowledge benefit team creativity? The role of internal team network structure and task complexity *Organ. Sci.* (2023).
- 97.A. Almaatouq *et al.*, Adaptive social networks promote the wisdom of crowds. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 11379–11386 (2020).
- 98.D. Engel *et al.*, "Collective intelligence in computer-mediated collaboration emerges in different contexts and cultures" in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Association for Computing Machinery, New York, NY, USA, 2015), <https://doi.org/10.1145/2702123.2702259>, pp. 3769–3778.
- 99.M. McPherson, L. Smith-Lovin, J. M. Cook, Birds of a feather: Homophily in social networks. *Annu. Rev. Sociol.* **27**, 415–444 (2001).
- 100.S. Wright, The shifting balance theory and macroevolution. *Annu. Rev. Genet.* **16**, 1–20 (1982).
- 101.J. Meluso, L. Hébert-Dufresne, Datasets for "Multidisciplinary learning through collective performance favors decentralization". *Open Science Framework*. <https://doi.org/10.17605/OSF.IO/KYVTD>. Deposited 17 June 2022.