

RESEARCH NOTE

From Coordination Topology to Organization

*A Cross-Domain View of Collective Problem Solving in Humans and AI***Kai Wang**

Independent Researcher, Australia | ORCID: 0009-0000-5018-9305 | Version 1.0, September 2026

Abstract

Recent experiments on human groups and language-model agents converge on a narrow but important point: collective performance cannot be inferred from node capability or communication topology alone. This note situates Kim et al. (2026) within a longer experimental lineage in human collective problem solving. Leavitt (1951) showed that communication pattern changes group behavior; Mason and Watts (2012) isolated topology while holding network size and local degree fixed; Bernstein, Shore, and Lazer (2018) showed that changing the timing of interaction changes collective search; and Almaatouq et al. (2021) showed that the balance between synergy and coordination loss depends on task complexity. Kim et al. extend this problem into modern AI systems under matched prompts, tools, and compute, showing large positive and negative returns from different coordination architectures and identifying single-agent baseline capability as the most robust predictor of system performance. Their earlier MDAgents work already made collaboration structure task-conditioned. Taken together, this literature suggests a broader research object: operating organization - the time-varying relation among node capabilities, available connections, roles or local policies, actual information-flow dynamics, and retained organizational history. The next question is therefore not merely which fixed topology wins, but what organization should form, when, and why.

1. The recurring empirical problem

Collective problem solving is often framed as a comparison between the capability of individual members and the advantages or costs of collaboration. That framing is incomplete. A group can contain the same participants yet behave differently when communication structure changes. It can retain the same potential connections yet behave differently when interaction is intermittent rather than continuous. The same collaborative arrangement can also switch from net cost to net benefit as the task becomes more complex. In AI systems, the same underlying model family can likewise produce very different system-level performance under different coordination architectures.

The common empirical question is therefore not simply whether groups outperform individuals. It is which properties of the operating arrangement determine whether distributed capability becomes useful collective capability rather than redundancy, delay, or error propagation. The human literature has progressively isolated several such variables. Kim et al. (2026) add a modern AI counterpart by jointly varying coordination architecture and model capability under a controlled system-level compute budget.

2. A seventy-five-year experimental lineage

Communication structure: Leavitt (1951)

Leavitt experimentally placed comparable five-person groups in different communication patterns, including the wheel, chain, Y, and circle. The communication pattern affected accuracy, total activity, member satisfaction, leadership emergence, and the organization of the group. The durable lesson is not that centralization is always superior. It is that holding the task and participant population broadly comparable while altering permitted communication relations can change collective behavior.

Topology under controlled local degree: Mason and Watts (2012)

Mason and Watts made the topology question cleaner in a networked search experiment. Groups of 16 participants searched the same problem landscape. Every network had 16 nodes and every participant had exactly three neighbors, while eight network topologies varied in global structure and information-diffusion efficiency. Networked groups outperformed equal-sized independent searchers on average, and performance varied with network efficiency. This isolates an important structural result: even when node count and local degree are fixed, who is connected to whom can change collective search.

Group process beyond member ability: Woolley et al. (2010)

A complementary literature shifted attention from network structure to group-level interaction. Woolley et al. reported a general collective-intelligence factor across diverse tasks. It was not strongly explained by the average or maximum individual intelligence of group members; instead, social sensitivity and more equal conversational turn-taking were associated with higher collective intelligence. The result does not identify a universal mechanism, but it reinforces the distinction between member capability and the interaction process through which member capability is organized.

Temporal activation of interaction: Bernstein, Shore, and Lazer (2018)

Bernstein, Shore, and Lazer changed a different variable: when interaction was available. Three-person groups repeatedly solved a traveling-salesperson problem under constant interaction, intermittent interaction, or no interaction. Constant interaction improved average performance by diffusing good solutions but reduced independent exploration; no interaction preserved exploration but prevented diffusion. Intermittent interaction combined much of both. The implication is explicitly temporal: a connection may exist as a potential relation without being continuously active as an organizational relation. The timing of interaction can change collective capability even when the participants and broad collaborative setting remain comparable.

Task-relative collaboration: Almaatouq et al. (2021)

Almaatouq et al. varied task complexity while comparing individuals with interacting groups. In a preregistered study of 1,200 participants, groups were as fast as the fastest individual and more efficient than the most efficient individual when tasks were complex, but not when tasks were simple. By separately quantifying synergistic gains and process losses, the study showed that the sign of collaboration value is task-relative rather than intrinsic to collaboration itself.

A computational bridge: topology x local policy

Barkoczi and Galesic (2016) provide a useful computational bridge rather than a human-subject experiment. In simulations, the performance ranking of efficient versus inefficient networks reversed when agents changed social-learning strategy. The result formalizes a point already suggested by the experiments above: topology has no context-free performance value. Its effect depends on the policies operating over the network and on the task environment.

3. Where Kim et al. fit in this lineage

Kim and colleagues have already moved through two distinct stages of this problem. In MDAgents (2024), collaboration structure was made task-conditioned: medical queries were classified by complexity and routed to solo or group collaboration, with the number and roles of collaborating agents adjusted to the task. This was already a move away from treating multi-agent collaboration as a universal default.

The 2026 Nature Machine Intelligence study broadens and systematizes the question. Across 260 controlled configurations spanning six benchmarks, five architectures, and three LLM families, the experiment matched task prompts, tool interfaces, and per-system compute ceilings while varying coordination structure and model capability. Performance changes relative to single-agent baselines ranged from +80.8% on structured financial reasoning to -70.0% on sequential planning. No architecture dominated across domains. Finance benefited from decomposable parallel analysis; PlanCraft degraded when coordination consumed resources needed for sequential state tracking.

The paper also identifies single-agent baseline performance as the most robustly supported predictor of absolute agent-system performance. A separately fitted selection rule found that baselines above roughly 45% tended to predict zero-to-negative multi-agent gains, while a held-out within-domain model selected the best architecture in 87% of configurations. These results do more than show that topology matters. They turn architecture selection into a measurable regime problem conditioned by task structure and node capability.

That gives the work a specific position in the longer history. Leavitt and Mason and Watts demonstrate that communication structure matters. Bernstein shows that temporal interaction policy matters. Almaatouq shows that task complexity changes the balance between synergy and process loss. Barkoczi and Galesic show that local node policy can reverse topology effects. Kim et al. add model capability and modern agentic task structure under matched computational constraints. The accumulated result is not a universal law of collaboration. It is a progressively richer map of variables that determine collective performance.

Experimental progression at a glance

Study	Nodes	Primary variable	Key control	Interpretive contribution
Leavitt, 1951	Human groups	Communication pattern	Same group task under wheel / chain / Y / circle	Communication structure changes group behavior

Study	Nodes	Primary variable	Key control	Interpretive contribution
Woolley et al., 2010	Human groups	Group interaction process	Diverse group tasks	Collective performance is not reducible to best-member ability
Mason & Watts, 2012	Human networks	Global topology	N = 16, degree = 3	Topology changes collective search under fixed local degree
Barkoczi & Galesic, 2016	Simulation	Topology x social-learning policy	Same search family, varied policies/networks	The sign of topology effects can reverse
Bernstein et al., 2018	Human groups	Interaction timing	Constant / intermittent / none	Temporal activation changes exploration-exploitation balance
Almaatouq et al., 2021	Human groups	Task complexity	Individuals vs 3-person groups	Synergy versus process loss changes with complexity
Kim et al., 2026	LLM agents	Architecture x model capability	Matched prompts, tools, compute	Coordination gains are regime-dependent; baseline capability is highly informative

4. From topology to operating organization

The combined literature suggests that topology is one component of organization rather than organization itself. A graph tells us which nodes can interact. It does not determine which links are active now, what information crosses them, how interactions are sequenced, which node proposes or validates a step, how strongly different paths are weighted, whether communication is continuous or intermittent, or how previous outcomes alter later coordination.

Connection Theory uses this distinction explicitly. Its sequence from interaction to feedback to connection to organization treats a stable connection inventory as insufficient to specify system-level organization. Organization additionally concerns the operational coordination of relations through routing, timing, synchronization, sequencing, amplification, combination, inhibition, and reuse over time. This is consistent with the human experimental record: Bernstein et al. can alter performance by changing temporal activation without requiring a fundamentally different node population, while Barkoczi and Galesic can reverse structural effects by changing the local learning policy operating on the same network family.

For analysis, an operating organization at time t can be decomposed as $O_t = (N_t, G_t, Pi_t, Phi_t, H_t)$, where N_t denotes nodes and their current capabilities; G_t the available connection topology; Pi_t functional roles or local operating policies; Phi_t actual information-flow dynamics such as routing, timing, sequencing, validation, synchronization, and coordination intensity; and H_t retained organizational history or state. This tuple is an analytical decomposition, not a claim that every organization must be implemented or mathematically represented in exactly this form.

The narrower cross-domain conclusion is therefore: collective capability is not a property of nodes considered alone, and it is not determined by topology in isolation. It is a property of an operating arrangement whose value is conditional on task, environment, node capability, and the dynamics through which relations are actually used. Human groups and LLM agents are not mechanistically equivalent; the convergence is structural, at the level of the organizational variables that controlled experiments show can matter.

Topology tells us who can interact. Organization tells us how the system actually operates.

The next question is not simply which topology wins, but what organization should form, when, and why.

5. The next question: organization as a state variable

Most current multi-agent evaluations compare a small set of predefined coordination architectures. MDAgents already takes an important step beyond a fixed default by selecting collaboration structure according to task complexity. The next scientific step is broader: can the organization itself be characterized as a changing state whose roles, connectivity, information-flow timing, coordination intensity, and specialization vary with task structure, node capability, feedback, and environment?

This question is prospective rather than retrospective. It is not enough to observe that one architecture happened to win on one benchmark and label it well organized. A stronger science of AI organization would seek measurable conditions that predict when differentiated collaboration will exceed coordination cost, when a strong single node is sufficient,

when a topology-policy combination will amplify rather than contain error, and when the organization should change because the operating regime has changed.

The historical sequence can therefore be read as a progression in the explanatory object: individual capability -> communication structure -> topology -> local policy and interaction timing -> task-relative collaboration -> model-capability x task x coordination -> operating organization. Kim et al. provide a particularly important AI anchor in this progression because their controlled design turns large performance reversals into a quantitative architecture-selection problem rather than an anecdotal claim that "more agents" are better.

Connection Theory treats this as part of a more general problem of adaptive organization through time. The present note deliberately stops at the theory-level research question. It does not describe engineering mechanisms for forming, selecting, updating, preserving, or governing organizational states.

Further reading

AI-specific essay: [Why an AI Organization Is More Than a Collection of Models](#)

Underlying framework: [Connection Theory](#)

Contact: wangkai@connection-theory.org

References

- Almaatouq, A., Alsobay, M., Yin, M., & Watts, D. J. (2021). Task complexity moderates group synergy. *Proceedings of the National Academy of Sciences*, 118(36), e2101062118. <https://doi.org/10.1073/pnas.2101062118>
- Barkoczi, D., & Galesic, M. (2016). Social learning strategies modify the effect of network structure on group performance. *Nature Communications*, 7, 13109. <https://doi.org/10.1038/ncomms13109>
- Bernstein, E., Shore, J., & Lazer, D. (2018). How intermittent breaks in interaction improve collective intelligence. *Proceedings of the National Academy of Sciences*, 115(35), 8734-8739. <https://doi.org/10.1073/pnas.1802407115>
- Kim, Y., Park, C., Jeong, H., et al. (2024). MDAgents: An Adaptive Collaboration of LLMs for Medical Decision-Making. *Advances in Neural Information Processing Systems* 37. <https://arxiv.org/abs/2404.15155>
- Kim, Y., Gu, K., Park, C., et al. (2026). Capable language models can outgrow the benefits of collaboration. *Nature Machine Intelligence*, 8, 1157-1172. <https://doi.org/10.1038/s42256-026-01268-y>
- Leavitt, H. J. (1951). Some effects of certain communication patterns on group performance. *Journal of Abnormal and Social Psychology*, 46(1), 38-50. <https://doi.org/10.1037/h0057189>
- Mason, W., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, 109(3), 764-769. <https://doi.org/10.1073/pnas.1110069108>
- Wang, K. (2026). Connection Theory: A General Framework for Adaptive Network Dynamics - From Temporally Extended Existence to Emergent Dynamic Properties. Version 4.6. <https://connection-theory.org/>
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686-688. <https://doi.org/10.1126/science.1193147>

Scope note. This research note addresses empirical positioning and theory-level questions only. It does not disclose implementation mechanisms for adaptive AI organization.