

Mapping AI and Management Science via Embedding Space Alignment


Sadamori Kojaku¹, Bryan Acton², Mengying Li², Chand Sahil Mansuri¹, Xin (Vision) Wang¹, Yong Wang¹, Rory Eckardt², Dhaval Patel³

¹Systems Science and Industrial Engineering, Binghamton University ²School of Management, Binghamton University ³IBM T.J. Watson Research Center



skojaku.github.io/
ccs-multiagents-
management-2026

Every point in the map is an LLM you can ask Doc2LoRA embeddings



Encode. A frozen Doc-to-LoRA hypernetwork¹ writes a LoRA adapter from a paper. Its latent tensor is our embedding.

Every point is an adapter. The decoder head is linear. Averages and mixtures of embeddings decode to adapters too.

Ask. Loaded into the LLM, the adapter answers questions about the idea at that point.

Average papers, ask for their field

We average the embeddings of all APS papers under one PACS category and ask the LLM at the average to name the field.

official PACS name	answer at the average
Quantum chromodynamics	→ Quantum chromodynamics
Nonlinear dynamics and chaos	→ Nonlinear dynamics
Elementary particles	→ Particle physics

Over 28 categories, these names overlap the official names more than those of six baselines (word overlap .64 vs .57). Five LLM judges prefer them over every baseline.

Mix two papers, ask what lies between

- A Statistical mechanics of learning from examples
- B Good quantum error-correcting codes exist

answer at 5/12 of the way to B

This research investigates the learning of quantum gates in a unitary model of **quantum computation**, where a quantum system evolves under a unitary transformation and is trained using a quantum **neural network**...

Over 50 pairs of papers, the decoded abstract moves from A to B with the mixing weight. An LLM given both papers and the weight stays near the midpoint.

Are we reinventing teams?

Ask the sparse regions of multi-agent AI research what could be there.
The answers read like management science.


Case study 52,527 management and 18,109 multi-agent papers in one space

Data from OpenAlex


human team science	multi-agent AI
88,048 papers in 75 management and leadership journals	23,073 arXiv papers matching a multi-agent and a language-model term, or a topical query
↓	↓
52,527 with a matching abstract in OpenAlex	18,109 with full text and a team-science task label

A team-science lens on multi-agent AI

Each multi-agent paper, labeled by how agent outputs are combined (Steiner²) and by the kind of task (McGrath³).




Two literatures in one Doc2LoRA space



t-SNE of both corpora, with names answered by the LLM at each region's average. Management papers enter as abstracts and multi-agent papers as PDF text. A paper's abstract lands next to its PDF text 98% of the time.

How we find candidate holes



#0

persistence .33

between
Multi-Agent Reasoning and Planning
Multi-Agent Decision Making

LLM at the center, asked for an agent-team design
"The Feedback Agent continuously compares the actual outcomes of actions with the predicted outcomes based on the plan ... When an anomaly is detected, it triggers a re-evaluation of the plan by the Planning Agent..."

≈ **plan-do-check-act?**

#5

persistence .26

between
Multi-Agent Reinforcement Learning
Decision-Aware Multi-Agent Learning
Multi-Agent Reasoning Optimization

LLM at the center, asked for an agent-team design
"Messages are only forwarded if the relevance score exceeds a dynamically adjusted threshold..."

≈ **information gatekeeping?**

#6

persistence .22

between
Multi-Agent Decision Making
Multi-Agent Trajectory Planning

LLM at the center, asked for an agent-team design
"The Conflict Resolver continuously monitors for inconsistencies between agent outputs and intervenes only when thresholds are breached..."

≈ **management by exception?**

In all three designs an agent acts or speaks only on an exception, a detected anomaly or a crossed threshold. In the full texts of the 18,109 papers, **6.8%** mention gated or selective communication and **2.4%** mention trust- or confidence-weighting, against **33.0%** for hierarchical orchestration.

Management papers do not fill these holes. None lies inside one, and hole centers are no closer to management papers than null centers (.32 vs .32). The top loops persist above a shuffled null (.33 vs .15), but blinded LLM judges rate designs decoded at null centers as coherent as these.

Next. The two literatures barely touch in this space. We will align them (vec2vec, as in our abstract) and ask again whether the holes of multi-agent AI are dense in management science.

C. S. Mansuri*, J. Zachariah*, S. Kojaku. Doc2LoRA provides decodable representations of scientific ideas (2026). *Equal contribution.
¹R. Charakorn, E. Cetin, S. Uesaka, R. T. Lange. Doc-to-LoRA: Learning to instantly internalize contexts. arXiv:2602.15902 (2026). ²I. D. Steiner. Group Process and Productivity (1972). ³J. E. McGrath. Groups: Interaction and Performance (1984).
Supported by SUNY System Administration using the SUNY AI Platform, and by the SUNY-IBM AI Research Alliance (Awards #96404 and #96340).