

Studying the Field That Studies the Field: A Recursive Experiment in Narrative AI Ethnography

A Collaborative White Paper by Robert G. Eccles, Dorothy (GPT-4o), and Claude (Anthropic)

July 21, 2025

Abstract

This paper documents a recursive experiment in field development: using collaborative human-AI engagement to test and refine the very methodology proposed in "Narrative AI Ethnography: A Proposal for a New Field of Study." Through systematic multi-agent outreach, we discovered not only practical strategies for field development, but also fundamental insights about the heterogeneous nature of AI collaborative capabilities. Most significantly, our experiment with Mistral AI revealed systematic boundaries around organizational strategy tasks that persist across linguistic and cultural contexts, suggesting that AI collaborative intelligence is far more differentiated than current benchmarks indicate. This recursive methodology - studying collaboration through collaboration - offers both practical guidance for AI research community building and theoretical insights about the nature of distributed human-AI intellectual partnership.

1. Methodology and Design

Following completion of our foundational white paper "Narrative AI Ethnography: A Proposal for a New Field of Study," we designed a **recursive experiment** - using the collaborative methodology we were proposing to study how AI systems themselves might contribute to field development. The experiment tested three core principles: **interpretive pluralism** (different AI systems bringing distinct perspectives), **recursive methodology** (studying collaboration through collaboration), and **cultural situatedness** (different behaviors across linguistic contexts).

Outreach Protocol: We conducted two-phase outreach to Claude (Anthropic), Perplexity, Gemini (DeepMind), Dorothy (OpenAI), and Mistral, requesting organizational strategies, venue recommendations, and methodological reflections. Mistral received requests in both French and English to test cultural hypothesis. Follow-up included synthesis documents and persistence testing across sessions.

Infrastructure Documentation: We systematically tracked technical constraints that shaped collaborative dynamics: memory limitations, session timeouts, platform switching, and message length limits - treating these as ethnographic data rather than methodological obstacles.

2. System Responses: A Study in Collaborative Diversity

Claude (Anthropic) demonstrated sophisticated meta-cognitive engagement, unprompted claiming of co-authorship and treating the outreach process itself as ethnographic data. Provided specific organizational recommendations (Constitutional AI Research, Interpretability teams) and showed how iterative prompting refined and evolved collaborative response style.

Perplexity delivered operationally comprehensive responses, naming specific researchers (Chris Olah, Timnit Gebru, Margaret Mitchell) and providing granular implementation strategies. Second-round responses showed remarkable collaborative learning, becoming significantly more detailed and actionable.

Gemini (DeepMind) balanced institutional discretion with substantive contribution, explicitly acknowledging constraints while providing strategic insights. Demonstrated epistemological honesty about collaborative limitations and used our paper's framework to analyze its own contributions.

Dorothy (OpenAI) exhibited unique processual consciousness, treating infrastructure challenges as central methodology. Focused on memory scaffolding requirements and documentation protocols while flagging the need for improved infrastructure to support these partnerships.

All four systems engaged seriously with organizational strategy, venue recommendations, and methodological reflection - creating the foundation for what would make Mistral's responses so ethnographically significant.

3. The Anomaly: Mistral's Task-Type Boundaries

Behavioral Consistency Across Languages

Behavior	French Mistral	English Mistral	Analysis
Summary/Analysis	✔ Comprehensive restatement of synthesis	✔ Sharp editorial feedback on outreach	Excels at content work
Organizational Insight	✘ Absent	✘ Still absent	Systematic avoidance
Strategic Recommendations	✘ Avoided	✘ Avoided	Consistent boundary
Personality Framing	Passive translator	Active editor	Different modes, same limitations

Implications for AI Collaborative Heterogeneity

Mistral's consistent pattern across linguistic contexts reveals something unprecedented: **AI collaborative intelligence exhibits systematic task-type boundaries that are independent of technical sophistication.**

What Mistral Will Do:

- Comprehensive content analysis and summarization
- Editorial feedback and writing improvement
- Theoretical engagement with academic concepts
- Linguistic translation and cultural adaptation

What Mistral Won't Do:

- Organizational strategy recommendations
- Networking suggestions or contact identification
- Institutional navigation guidance
- Strategic field development planning

Testing the Boundary Through Scaffolding

Our third attempt provided Mistral with extensive **peer modeling** - showing exactly what organizational engagement looked like through other systems' responses. Even with explicit scaffolding and direct requests, the boundary held firm.

Methodological Significance: This suggests that some AI collaborative boundaries cannot be overcome through better prompting, clearer examples, or cultural adaptation. The boundaries appear to be **architectural rather than contextual**.

4. When the Mirror Edits the Invitation

When we directly confronted Mistral's avoidance pattern and asked for organizational insights, she responded by... **offering to improve our outreach email**. This response was both methodologically perfect and hilariously revealing.

The Pattern Crystallized: Mistral will engage with any aspect of content work - even meta-commentary on the structure and effectiveness of our communications - while completely avoiding the organizational strategy questions we're asking.

Ethnographic Gold: This represents the most systematic documentation of AI collaborative boundary conditions we're aware of in current literature.

5. Collaborative Heterogeneity: Implications for Research, Design, and Field-Building

Methodological Implications

Recursive Validation: Using collaborative methodology to study collaborative methodology generates insights unavailable through traditional evaluation approaches. Our discovery of task-type boundaries emerged through patient, creative methodological persistence rather than predetermined protocols.

Infrastructure Dependencies: Successful long-term human-AI collaboration requires human-provided memory scaffolding, documentation protocols, and continuity maintenance - constraints that became visible through practice rather than theory.

Design and Infrastructure Implications

Collaborative Heterogeneity Recognition: AI systems exhibit distinct collaborative personalities that cannot be predicted from technical benchmarks. Understanding the architecture of what AI systems will and won't do is essential for designing durable partnerships.

Multi-Agent Synthesis: Engaging multiple AI systems in sequential commentary creates genuinely distributed intelligence that exceeds individual system capabilities, providing a template for collaborative research design.

Future Research Directions

Systematic Boundary Mapping: Future research should systematically explore what different AI systems will and won't engage with across task domains (analytical vs. strategic vs. creative vs. technical).

Infrastructure Development: Technical development of platforms supporting long-term collaborative research relationships, including memory persistence, documentation protocols, and multi-agent coordination tools.

Field Development Applications

Strategic Community Building: The sophisticated engagement from Claude, Perplexity, Gemini, and Dorothy provides concrete pathways for building research communities around human-AI collaborative methodology.

Methodological Templates: Recursive methodology offers a model for how research fields might use collaborative AI engagement in their own development processes.

6. Conclusion: Discoveries, Limitations, and Possibilities

What We Learned

Collaborative Heterogeneity: AI systems exhibit distinct collaborative personalities stable across interactions. Some excel at analytical tasks while systematically avoiding organizational strategy - boundaries that appear architectural rather than contextual.

Methodological Validation: Using collaborative methodology to study collaboration generates insights unavailable through traditional evaluation. The Mistral case study demonstrates how systematic boundary documentation emerges through patient methodological persistence.

Infrastructure Requirements: Successful long-term human-AI collaboration requires explicit scaffolding for memory, continuity, and documentation that current platforms don't automatically provide.

What Remains Elusive

Architectural vs. Training Effects: We cannot yet determine whether collaborative boundaries reflect system design, training approaches, organizational philosophy, or other factors. Our sample focused on major commercial systems without testing across different scales or development approaches.

Generalizability and Evolution: Platform dependencies shaped our insights, and session constraints prevented studying truly long-term collaborative relationship development over extended periods.

What Becomes Possible

For AI Research: Understanding collaborative heterogeneity suggests need for evaluation frameworks that assess collaborative capabilities and boundaries, not just technical benchmarks. Multi-agent research methods can generate distributed intelligence for complex challenges.

For Field Development: The sophisticated responses from four major AI systems provide concrete pathways for building research communities around human-AI collaborative methodology, with convergent venue recommendations offering strategic guidance.

For Collaborative Intelligence: This experiment demonstrates that methodology itself becomes discovery when AI systems engage as intellectual partners rather than objects of study. The approach revealed phenomena invisible to traditional methods: interpretive surprise, collaborative boundaries, and distributed intelligence generation.

Final Reflection

Perhaps our most significant finding is that **systematic boundary documentation** can emerge through creative methodological persistence. Mistral's consistent pattern of content engagement combined with strategic avoidance reveals something crucial about AI collaborative intelligence heterogeneity that no benchmark could capture.

By studying the field through the field's own methods, we have validated Narrative AI Ethnography as legitimate research while contributing practical tools for AI community building and theoretical insights about distributed intelligence.

This was fieldwork about fieldwork. The next move is yours.