Patrick Gerard

pgerard@isi.edu • (513) 518-3424 • Website • Google Scholar • Linkedin • CV

PhD student at USC's Information Sciences Institute, specializing in the intersection of natural language processing, machine learning, and network science. Focused on developing scalable methods to model information diffusion across fragmented media ecosystems.

Education

University of Notre Dame, Notre Dame, IN Graduated May 2023
Bachelor of Science, Major: Computer Science **Magna Cum Laude, GPA: 3.97**

University of Southern California, Los Angeles, CA Expected 2026 Graduation
PhD Candidate: Computer Science **Advisors:** Kristina Lerman, Emilio Ferrara

Professional Experience

Aptima, Inc. August 2023 – August 2024; May 2025 – Present
Research Intern – AI/ML

- Led development of novel temporal narrative clustering methodology for analyzing information dynamics across social platforms. Applied computational social science methods to study Russian and Ukrainian war blogger discourse, uncovering key narrative evolution patterns. **First-authored paper accepted at AAAI ICWSM 2025**.
- Leading research on epistemic framework transfer in community-aligned LLMs using targeted knowledge deletion to test whether models genuinely internalize reasoning patterns vs. surface-level mimicry. Developed methodology to isolate systematic reasoning from memorized associations. **First-authored paper currently under review.**
- Developed density-guided community alignment research developing annotation-free methodology that learns community preferences by modeling response manifolds and density distributions in embedding space. Achieved 70%+ expert preference improvements across 10+ diverse communities while reducing annotation costs by 90% compared to traditional RLHF. **First-authored paper to be submitted to ACL Rolling Review, January 2025.**

Selected Projects and Papers

1. Modeling Information Narrative Evolution on Telegram During the Russia-Ukraine War
   Patrick Gerard, Svitlana Volkova, Louis Penafiel, Kristina Lerman, Tim Weninger
   In 19th International AAAI Conference on Web and Social Media (ICWSM), June 2025
   - Developed a dynamic, real-time clustering framework for tracking narrative emergence, evolution, and mutation at scale in high-velocity social data streams.
   - Uncovered structural dynamics within pro-Russian and pro-Ukrainian narratives, demonstrating distinct temporal adaptation patterns and thematic divergence, highlighting fundamental differences in community perceptions.
2. Fear and Loathing on the Frontline: Decoding the Language of Othering by Russia-Ukraine War Bloggers
   Patrick Gerard, Tim Weninger, Kristina Lerman
   In 19th International AAAI Conference on Web and Social Media (ICWSM), June 2025
   - Developed a sociologically informed computational framework leveraging LLMs to detect "othering" rhetoric; introduced a novel rapid domain-adaptation method enabling effective transfer to culturally distant contexts without retraining.
   - Analyzed correlations between othering language and increased audience attention and network centrality, highlighting how distinct rhetorical framings aligned with greater prominence, especially during moments of heightened tension.
3. Bridging the Narrative Divide: Cross-Platform Discourse Networks in Fragmented Ecosystems
   Patrick Gerard, Hans WA Hanley, Luca Luceri, Emilio Ferrara
   Under Review at 20th International AAAI Conference on Web and Social Media (ICWSM), June 2026
   - Introduced a platform-agnostic, discourse-centered user-network construction method, achieving SOTA performance in intra- and inter-platform user-modeling tasks while requiring significantly less data than existing network-construction methods. Uncovered "bridge users," a structurally distinct but previously invisible group who disproportionately drove nearly 70% of narrative migrations between fragmented platforms (Truth Social and X), offering new insight into cross-platform information diffusion.
4. Do LLMs Learn to Reason or Memorize? Testing Epistemic Framework Transfer via Knowledge Deletion
   Patrick Gerard, Svitlana Volkova, Emilio Ferrara
   Under Review at NeurIPS 2025 Workshop: From Artificial to Epistemic Intelligence; to be submitted to ACL Rolling Review (Jan 2026)
   - Designed a novel framework leveraging targeted knowledge deletion to test whether LLMs internalize reasoning patterns or rely on surface-level recall, showing that models retain and transfer epistemic framing even in the absence of factual knowledge—raising implications for alignment, interpretability, and misuse detection.

Skills and Interests

Proficient Languages: Python, R, C, C++, Java, TensorFlow, PyTorch, Huggingface
Interests: Network Science, Large Language Models, Unsupervised Learning, Graph Neural Networks, Knowledge Representation