

## Fan Huang

[fanhuan@iu.edu](mailto:fanhuan@iu.edu)

[muyuhuatang.github.io](https://muyuhuatang.github.io)

Luddy School of Informatics, Computing, and Engineering  
Indiana University Bloomington

## Education

- 2023–\*      **Ph.D.** in Informatics (Complex Networks and Systems), Indiana University Bloomington, USA  
Advisors: Prof. Jisun An and Prof. Haewoon Kwak
- 2022–23      **Ph.D.** in Computer Science, Singapore Management University, Singapore
- 2020–22      **M.Sc.** in Information Systems, Nanyang Technological University, Singapore  
Thesis: *Public Opinion Trend Analysis on Nuclear Energy via Aspect-based Sentiment Analysis Methods*
- 2014–18      **B.Sc.** in Computer Science & Technology, Central South University, China

## Research Experience

- 2023–      Research Assistant, Indiana University Bloomington  
Advisors: Prof. Jisun An and Prof. Haewoon Kwak  
Collaborators: Prof. Yong-Yeol Ahn, Prof. Filippo Menczer, and Prof. Filippo Radicchi
- 2024      Quantitative Research Intern, Core Research Team, QuantLab Financial, Houston (Jun–Aug 2024)
- 2024      Research Summer Internship, Indiana University Bloomington, supervised by Prof. Katy Börner (May–Jun 2024)
- 2022–23      Research Assistant, Singapore Management University

## Publications

 [Google Scholar](#)

† → Equal contribution

## Preprints / Under Review

- J1. **Huang, Fan.** Catalyst Papers in Artificial Intelligence Research: A Landscape on ICLR from 2017 to 2025. *arXiv preprint* (2026).
- J2. **Huang, Fan.** Reasoning Topology Matters: Network-of-Thought for Complex Reasoning Tasks. *arXiv preprint arXiv:2603.20730* (2026).
- J3. **Huang, Fan.** ReFlect: An Effective Harness System for Complex Long-Horizon LLM Reasoning. *arXiv preprint* (2026).

---

\*Expected 2027.

- J4. **Huang, Fan.** Simulating Hate Speech Cascades with Multi-LLM Agents: Empirical Grounding, Modeling Fidelity, and Intervention Strategies. *arXiv preprint* (2026).
- J5. **Huang, Fan & Kim, M.** LLMs as Implicit Probabilistic World Models: Distributional Stability, Convergence, and Scale Discordance. *OpenReview* (2026).
- J6. **Huang, Fan & Kim, M.** Can LLMs Predict Academic Collaboration? Topology Heuristics vs. LLM-Based Link Prediction on Real Co-authorship Networks. *arXiv preprint arXiv:2604.01379* (2026).
- J7. **Huang, Fan, Tong, D., Chen, C., Kwak, H. & An, J.** Instruction-Conditioned Motion Control (ICMC) for Robotic Arms: VLA Decomposition and Reinforcement Learning Tuning. *OpenReview* (2026).
- J8. **Huang, Fan, Zhang, S., Kwak, H. & An, J.** CogBias: Measuring and Mitigating Cognitive Bias in Large Language Models. *arXiv preprint arXiv:2604.01366* (2026).
- J9. **Huang, Fan & Zhou, H.** Comparative Evaluation of LLM-Based Review Paradigms for Academic Paper Quality Assessment. *OpenReview* (2026).
- J10. **Qi, W., Huang, Fan, An, J. & Kwak, H.** A Cross-Cultural Comparison of LLM-based Public Opinion Simulation: Evaluating Chinese and US Models on Diverse Societies. *arXiv preprint arXiv:2506.21587* (2025).
- J11. **Seckin, O. C., Nascimento Silva, F., Truong, B. T., Kim, S., Huang, Fan, Liu, N., Flammini, A. & Menczer, F.** The Rise of Bluesky. *arXiv preprint* (2025).
- J12. **Yin, L. & Huang, Fan.** Dif: A Framework for Benchmarking and Verifying Implicit Bias in LLMs. *arXiv preprint arXiv:2505.10013* (2025).
- J13. **Huang, Fan, Kwak, H. & An, J.** Token-Ensemble Text Generation: On Attacking the Automatic AI-Generated Text Detection. *arXiv preprint* (2024).
- J14. **Huang, Fan.** Highlight Timestamp Detection Model for Comedy Videos via Multimodal Sentiment Analysis. *arXiv preprint*. <https://arxiv.org/abs/2106.00451> (2021).
- J15. **Huang, Fan.** Network Activities Recognition and Analysis Based on Supervised Machine Learning Classification Methods Using J48 and Naïve Bayes Algorithm. *arXiv preprint*. <https://arxiv.org/abs/2105.13698> (2021).

## Peer-reviewed Conference Proceedings

- C1. **Huang, Fan.** LLM Belief Atlas in InfoVis 2026 (Poster) (2026).
- C2. **Huang, Fan, Kwak, H. & An, J.** Understanding Moral Reasoning Trajectories in Large Language Models: Toward Probing-Based Explainability in Findings of the Association for Computational Linguistics (EMNLP Findings 2026) (2026).
- C3. **Huang, Fan, Kwak, H. & An, J.** Vulnerability of LLMs' Stated Belief? LLMs Belief Resistance Check Through Strategic Persuasive Conversation Interventions in Findings of the Association for Computational Linguistics (ACL Findings 2026) (2026).
- C4. **Qi, W., Huang, Fan, Muralidharan, R., An, J. & Kwak, H.** XChoice: Explainable Evaluation of AI-Human Alignment in LLM-based Constrained Choice Decision Making in Findings of the Association for Computational Linguistics (AACL-IJCNLP Findings 2026) (2026).
- C5. **Qi, W., Huang, Fan, An, J. & Kwak, H.** Is DeepSeek a New Voice Among LLMs in Public Opinion Simulation? in *Proceedings of the IEEE International Conference on Data Mining (ICDM 2025)* (2025).

- C6. **Huang, Fan**, Kwak, H., Park, K. & An, J. *ChatGPT Rates Like Human: Towards a Better Alignment of Text Explanation Quality Assessments* in *Proceedings of the Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)* (2024).
- C7. **Huang, Fan**, Kwak, H. & An, J. *Chain of Explanation: New Prompting Method to Generate Higher Quality Natural Language Explanation for Implicit Hate Speech* in *Companion Proceedings of the Web Conference (WWW 2023)* (2023).
- C8. **Huang, Fan**, Kwak, H. & An, J. *Is ChatGPT Better than Human Annotators? Potential and Limitations of ChatGPT in Explaining Implicit Hate Speech* in *Companion Proceedings of the Web Conference (WWW 2023)* (2023).
- C9. Liaw, S., **Huang, Fan**, Benevenuto, F., Kwak, H. & An, J. *YouNIcon: YouTube’s CommuNity of Conspiracy Videos* in *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM 2023)* (2023).

## Teaching

### Indiana University Bloomington – Luddy School of Informatics, Computing, and Engineering

|      |                                                                                                                                        |
|------|----------------------------------------------------------------------------------------------------------------------------------------|
| 2026 | Teaching Assistant, INFO-I 201 Mathematical Foundations of Informatics (Fall 2026)                                                     |
| 2026 | Lecture Presentation, Vision Language Models (VLM), INFO-I 590 Foundation and Application for Large Language Models (Apr 2026)         |
| 2026 | Lecture Presentation, Retrieval Augmented Generation (RAG), INFO-I 590 Foundation and Application for Large Language Models (Mar 2026) |
| 2026 | Teaching Assistant, INFO-I 590 Foundation and Application for Large Language Models (Spring 2026)                                      |
| 2025 | Teaching Assistant, INFO-I 590 Foundation and Application for Large Language Models (Fall 2025)                                        |
| 2025 | Teaching Assistant, INFO-I 590 Foundation and Application for Large Language Models (Spring 2025)                                      |

## Research Advising

Master’s Student: Sathvik Bandaru (co-advising)

## Presentations

### Talks

- T1. **Huang, Fan**. *AI-Assisted Research* Invited talk, Prof. Thai Le’s lab, Indiana University Bloomington. Apr. 2026.
- T2. **Huang, Fan**. *LLM Belief Atlas* Poster presentation at InfoVis 2026 (Boston). 2026.
- T3. **Huang, Fan**. *LLM Moral Reasoning Trajectory Analysis* Invited talk, Prof. Han Xue’s lab, Donghua University, Shanghai. June 2026.

- T4. **Huang, Fan.** *Understanding Moral Reasoning Trajectories in Large Language Models: Toward Probing-Based Explainability* Parallel talk at IC2S2 2026 (Burlington, Vermont). 2026.
- T5. **Huang, Fan.** *Vulnerability of LLMs' Stated Belief? LLMs Belief Resistance Check Through Strategic Persuasive Conversation Interventions* Presentation at MSLD 2026 (UIUC). Mar. 2026.
- T6. **Huang, Fan.** *ChatGPT Rates Like Human: Towards a Better Alignment of Text Explanation Quality Assessments* Oral presentation at MSLD 2024 (Ann Arbor, Michigan, USA). 2024.
- T7. **Huang, Fan.** *Is ChatGPT Better than Human Annotators? & Chain of Explanation* Poster presentation at The Web Conference (WWW) 2023 (Austin, Texas, USA). 2023.
- T8. **Huang, Fan.** *Natural Language Explanation: From XAI to Human-AI Interaction* YY Group, CNetS, Indiana University Bloomington. 2023.

## Workshops

- WK1. **Huang, Fan.** *Participant, Summer Institute of Computational Social Science (SICSS)* NUS, Singapore. 2022.
- WK2. **Huang, Fan.** *Co-organizer, Workshop "Introduction to NLP"* IEEE NTU Student Branch, NTU, Singapore. Oct. 2021.

## Academic Service

### Program Committee Member

|      |                                  |
|------|----------------------------------|
| 2026 | WISE 2026                        |
| 2026 | IC2S2 2026                       |
| 2026 | WWW 2026: Web4Good               |
| 2026 | 7th CySoc Workshop at ICWSM 2026 |
| 2025 | 6th CySoc Workshop at ICWSM 2025 |
| 2025 | WebSci 2025                      |

### Conference Reviewer

|      |                                                                          |
|------|--------------------------------------------------------------------------|
| 2026 | NeurIPS 2026                                                             |
| 2026 | KDD 2026                                                                 |
| 2026 | COLM 2026                                                                |
| 2026 | IUI 2026                                                                 |
| 2026 | ACL Rolling Review (ARR) 2026                                            |
| 2026 | ICWSM 2026                                                               |
| 2025 | IMX 2025 (ACM International Conference on Interactive Media Experiences) |
| 2025 | ACL Rolling Review (ARR) 2025                                            |
| 2025 | CHI 2025                                                                 |
| 2025 | ICWSM 2025                                                               |
| 2025 | IDC 2025                                                                 |
| 2022 | WSDM 2022                                                                |

## Journal Reviewer

|         |                                     |
|---------|-------------------------------------|
| 2025    | PeerJ Computer Science              |
| 2024–25 | Online Social Networks and Media    |
| 2024    | Information Processing & Management |

## Organized Activities / Volunteer

|         |                                                                                                     |
|---------|-----------------------------------------------------------------------------------------------------|
| 2026    | IC2S2 2026: Science of Science (AI) Parallel Talk Session, Session Chair                            |
| 2025    | SODA Lab Research Talk: ‘Multimodal NLP Research Insights’, Dr. Henry Wei Han, Organizer (Nov 2025) |
| 2024–   | LLM Innovators Hub @ Soda Lab, Organizer (since Oct 2024)                                           |
| 2023–24 | CNetS @ Luddy School Gathering Event, Organizer (Aug 2023–Oct 2024)                                 |
| 2023    | ISSI 2023, Student Volunteer                                                                        |

## Technical Skills

|                       |                                                                                                    |
|-----------------------|----------------------------------------------------------------------------------------------------|
| <b>Programming</b>    | Python, R, SQL, Java                                                                               |
| <b>ML/NLP</b>         | PyTorch, Hugging Face Transformers, LangChain, OpenAI API, scikit-learn, spaCy, NLTK               |
| <b>LLM Tools</b>      | Fine-tuning (LoRA, PEFT), Retrieval-Augmented Generation (RAG), Prompt Engineering, LLM Evaluation |
| <b>Data &amp; Viz</b> | Pandas, NumPy, Matplotlib, Seaborn, NetworkX                                                       |
| <b>Platforms</b>      | Linux/Unix, Git, Docker, HPC/Slurm, Jupyter                                                        |

---

Last updated: September 8, 2026