Research Statement

Hanjia LYU
School of Computing and Information Systems, Singapore Management University
Tel : (65) 6828 0686 ; Email : hjlyu@smu.edu.sg
25 May 2026

Background

When outcomes hinge on human actions or interpretations, achieving effective and trustworthy AI performance requires systems that can reason about human behavior, communication, and social context. For example, policy communication tools that understand how people interpret information can improve public engagement and trust, while recommendation systems that model users' evolving preferences and goals can provide more meaningful, transparent guidance. Likewise, knowledge systems that understand how humans connect and apply information can deliver dynamic, context-aware insights rather than static answers. Building such human-aligned intelligence demands AI that not only processes data but also comprehends the richness of human expression and the structure of human interactions. To this end, my research develops methods for constructing human-aligned AI systems that understand and reason in human-like ways through augmenting and structuring data representations across modalities to capture what humans express through language, visuals, and behavior, and across graphs to represent how humans and information are connected within social and semantic networks. By leveraging and enhancing Large Language Models (LLMs) and their multimodal variants (MLLMs) as core reasoning engines, my work aims to create AI systems that not only integrate diverse sources of information but also reason in ways that are interpretable, adaptive, and socially grounded.

Research Areas

1. Modality-Based Augmentation

Humans communicate through diverse modalities that together convey meaning, intent, and emotion. For AI systems to become truly human-aligned, they need to not only process these signals but also understand their semantic richness and contextual nuance. Yet real-world data are often noisy, sparse in meaningful signals, and imbalanced. To address these challenges, my research advances modality-based augmentation along two directions. At the content level, I design methods to refine and enrich unimodal and multimodal data representations, improving the quality and informativeness of learned representations. At the distribution level, I design approaches to mitigate data imbalance by adaptively amplifying informative samples, guiding models toward efficient learning and improved performance. These efforts establish a progressive framework for augmenting perceptual data, enhancing its quality and completeness.

<u>Refining Unimodal and Multimodal Representations</u> A central challenge in modality-based learning lies in the limited semantic richness of individual modalities. Human-generated text, for example, often carries incomplete or ambiguous descriptions that

obscure the behavioral intent. To address this, I develop methods that enrich representations to make them more expressive and behaviorally informative.
For unimodal representations, my work addresses the challenge of improving text-based recommendation systems, where item descriptions are often sparse or uninformative. To enrich these descriptions, I employ LLMs to generate additional, semantically meaningful content through four structured prompting strategies. This enrichment substantially increases the density of descriptive keywords and attributes, outperforming traditional content- and knowledge-based augmentation methods as well as state-of-the-art LLM-based augmentation baselines in recommendation quality. For multimodal representations, I develop methods that bridge textual and visual representations to enhance the behavioral interpretability of multimodal data. I analyze both textual and visual modalities from GoFundMe campaigns to capture the factors that influence donor behavior. Inspired by how humans reason across multiple cues, I design a framework that enables MLLMs to reason jointly over text and images. These works deepen multimodal understanding and support AI reasoning in environments that encompass semantically rich and heterogeneous information.

<u>Amplifying Meaningful Signals under Imbalance</u> Real-world human data are often imbalanced. Some behaviors, opinions, and modalities are richly represented, while others are sparse or overlooked. Such imbalance biases model perception and weakens generalization across diverse human contexts. To address this, I extend modality-based augmentation from improving data content to optimizing its distribution, developing strategies that selectively amplify the information most critical for understanding human behavior. For example, when studying opinions about COVID-19 vaccines before they were available, many online posts contain vaccine-related keywords but do not actually indicate whether the author is willing to take the vaccine. Directly annotating all such posts is inefficient. To address this, I develop a human-guided machine learning framework that employs active learning to retrieve the most informative posts, raising the proportion of relevant posts from 16% to 40% and improving F1 from 0.45 to 0.60 in just five annotation rounds of 500 samples each, striking a balance between efficiency and accuracy. More broadly, this line of work demonstrates that selectively amplifying meaningful signals can substantially improve both the efficiency and effectiveness. Beyond performance gains, it contributes a general framework for mitigating data imbalance by combining human guidance with machine learning to prioritize information that truly reflects behavioral intent.

2. Graph-Based Augmentation

Human behavior is also inherently relational, shaped by social ties, shared knowledge, and contextual dependencies among people, ideas, and actions. To align AI systems with human cognition, I move beyond isolated representations toward models that understand and leverage these relationships. At the content level, I augment existing AI systems with structured representations, enabling them to incorporate dynamic knowledge and follow explicit reasoning paths. At the distribution level, I develop a mixture-of-experts framework that adaptively routes data between graph-based and non-graph-based models, balancing structured and unstructured learning. These efforts advance graph-based augmentation as a means to capture how humans relate, infer, and reason within networks of knowledge and interaction.

Structuring Representations through Graph Integration Human reasoning depends not only on recognizing concepts but also on understanding how those concepts are related, who interacts with whom, which events influence others, and how beliefs evolve within a social or political context. Yet, the AI systems that only rely on unstructured textual data lack the explicit relational scaffolding that humans use to interpret the world. This limitation is especially evident in domains such as political or social analysis, where relationships among entities are dynamic, causal, and hierarchically organized. To address this, I develop methods that structure model representations through graph integration, allowing AI systems to incorporate curated relational knowledge and follow transparent reasoning paths. I augment LLMs with a multi-view political knowledge graph that connects entities, events, and attributes through both local and global evidence. The model first retrieves entity-linked triples and filters them via dense retrieval to select salient facts. It then aggregates these into global representations using both embedding-based and LLM-based reasoning modules, which are injected back into the model's inference pipeline. Empirically, the framework improves performance on political reasoning and forecasting tasks, achieving an average absolute accuracy gain of 3.16% over state-of-the-art baselines. This research advances the integration of structured knowledge and reasoning, enabling AI systems to reason not only about what is known but also about how pieces of information are connected. By combining graph-based relational structure with the contextual flexibility of LLMs, this work contributes to the development of interpretable reasoning frameworks that mirror the relational nature of human cognition. Beyond improving political analysis tasks, it offers a general pathway toward AI systems capable of understanding complex, interconnected domains with greater fidelity, and explanatory power.

Balancing Structured and Unstructured Learning In many real-world graphs, connected nodes do not necessarily share common properties, a phenomenon known as heterophily. For instance, people who are closely related to each other may still hold very different opinions or behaviors. Traditional graph neural networks (GNNs) assume homophily, where adjacent nodes tend to share labels or attributes, and their message passing aggregates neighbor features to smooth node representations. Under heterophily, this aggregation mixes conflicting signals, dilutes node-specific information, leading to oversmoothing and degraded predictions. The core challenge is to benefit from structure when it is informative while ignoring misleading neighborhoods. I design a mixture-of-experts model that combines GNNs with MLPs. By selectively routing data, the framework effectively augments each expert's training signal: the GNN learns enriched representations when neighbor information is useful, while the MLP specializes on cases where local features dominate. This approach consistently outperforms state-of-the-art baselines across all evaluated datasets, yielding an average absolute accuracy gain of 2.16% over the corresponding GNN backbone. By dynamically balancing structured and unstructured signals, the framework demonstrates how AI systems can benefit from relational information without being constrained by it. Beyond improving predictive performance, this approach contributes to a deeper understanding of how to design models that reason effectively in complex networks where relationships can be informative, neutral, or even misleading.

3. Integrating Modalities and Graphs

This branch of my work integrates modalities and graphs for augmentation, capturing both what people express and how they are socially and informationally connected. At the content level, I study user representations on heterogeneous graphs by augmenting textual posts with visual cues and retweeting relationships, enabling models to infer not only what individuals say but also how their perspectives spread and polarize through social networks. At the distribution level, I address the sparsity of user–item interactions by clustering users with similar behavioral patterns to form informative priors that enrich user representations. This approach effectively augments behavioral modalities through graph-guided clustering, allowing models to learn from community-level dynamics rather than isolated individuals.

<u>Bridging Perceptual and Relational Signals</u> I introduce a bipartite heterogeneous graph neural framework that fuses multimodal post content with user–post relational structure. Nodes are of two types (users and posts), and edges represent interactions (post, retweet). Each post node is enriched with both text embeddings and visual embeddings when available, and each user node includes profile and social features. The model applies heterogeneous GNN layers that sample and aggregate neighbors via random walks, use a BiLSTM with an attention mechanism to combine intra-type and cross-type neighbor information, and integrate these signals with attention weights over "self" vs neighbor embeddings. A specialized social-platform negative sampling strategy ensures the model emphasizes behaviorally informative user–tweet relationships. Without requiring any ideology labels, this method learns meaningful representations that outperform unimodal, early/late fusion, and homogeneous GNN baselines by a wide margin (~9% AUROC gain). The resulting framework allows AI systems to infer behavioral and ideological patterns directly from how people express and connect, rather than relying on manually defined labels. Beyond achieving substantial empirical gains, this work contributes to the broader understanding of how information, behavior, and influence co-evolve across social networks, offering more interpretable, context-aware, and socially responsible analysis of human communication.

<u>Enhancing Modality Representations through Graph-Guided Clustering</u> I develop a graph-based generative recommendation framework that augments user behavioral representations through graph-guided clustering. Instead of modeling users as independent data points, this framework constructs an item–item co-engagement graph that captures implicit connections among users through shared interactions. Using a graph clustering algorithm, the model identifies latent interest clusters that summarize recurring patterns of co-engagement. This process effectively amplifies informative behavioral signals by clustering users with similar patterns, enriching their modality representations with graph-derived priors. The method has been deployed across multiple product surfaces at one of the largest social media platforms. Product launches with my approach lead to >0.03% in-session improvement and exhibit consistent gains (0.06% to 0.19%) across key engagement metrics such as total time spent, likes, and shares. This research advances recommendation modeling by moving from isolated user representations to relational, structure-aware reasoning. Beyond its empirical success and large-scale deployment, it demonstrates how

graph-structured learning can translate fundamental advances in representation learning into measurable, real-world impact.

My long-term research goal is to advance artificial intelligence from human-aligned to human-centered systems. While human-aligned AI focuses on ensuring that an AI's understanding and behavior remain consistent with human values and intentions, human-centered AI emphasizes the design of systems that genuinely understand, collaborate with, and empower humans in real-world contexts. The key gap between the two lies in how alignment principles can be transformed into meaningful, usable, and trustworthy human experiences. Two fundamental aspects that characterize this gap are (i) understanding complexity, which refers to the ability of AI to reason about rich and dynamic scenarios involving multiple entities and long-term interactions, and (ii) explainability, which concerns the ability to make AI's decision-making processes transparent and interpretable to humans.

Selected Publications and Outputs

- Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Chris Leung, Jiajie Tang, and Jiebo Luo. LLM-Rec: Personalized recommendation via prompting large language models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 583–612, Mexico City, Mexico, June 2024. Association for Computational Linguistics.
- Xupin Zhang, Hanjia Lyu, and Jiebo Luo. What contributes to a crowdfunding campaign's success? evidence and analyses from GoFundMe data. *Journal of Social Computing*, 2(2):183–192, 2021.
- Hanjia Lyu, Junda Wang, Wei Wu, Viet Duong, Xiyang Zhang, Timothy D. Dye, and Jiebo Luo. Social media study of public opinions on potential covid-19 vaccines: informing dissent, disparities, and dissemination. *Intelligent Medicine*, 2(1):1–12, 2022.
- Hanjia Lyu*, Jinsheng Pan*, Zichen Wang*, and Jiebo Luo. Computational assessment of hyperpartisanship in news titles. In *Proceedings of the International AAAI Conference on Web and Social Media*, 18(1):999–1012, May 2024.
- Hanjia Lyu*, Jinfa Huang*, Daoan Zhang*, Yongsheng Yu*, Xinyi Mou*, Jinsheng Pan, Zhengyuan Yang, Zhongyu Wei, and Jiebo Luo. GPT-4V(ision) as a social media analysis engine. *ACM Trans. Intell. Syst. Technol.,* 16(3), April 2025.
- Weihong Qi, Hanjia Lyu, and Jiebo Luo. Representation bias in political sample simulations with large language models. In *Companion Proceedings of the ACM on Web Conference 2025, WWW '25*, page 1264–1267, New York, NY, USA, 2025. Association for Computing Machinery.
- Hanjia Lyu, Jiebo Luo, Jian Kang, and Allison Koenecke. Characterizing bias: Benchmarking large language models in simplified versus traditional chinese. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency, FAccT '25*, page 2815–2846, New York, NY, USA, 2025. Association for Computing Machinery.
- Xinyi Mou, Zejun Li, Hanjia Lyu, Jiebo Luo, and Zhongyu Wei. Unifying local and global knowledge: Empowering large language models as political experts with knowledge graphs. In *Proceedings of the ACM Web Conference 2024*, *WWW '24*, page 2603–2614, New York, NY, USA, 2024. Association for Computing Machinery.

- Hanqing Zeng*, Hanjia Lyu*, Diyi Hu, Yinglong Xia, and Jiebo Luo. Mixture of weak & strong experts on graphs. In *The Twelfth International Conference on Learning Representations*, ICLR, Vienna, Austria, May, 2024.
- Hanjia Lyu and Jiebo Luo. Understanding political polarization via jointly modeling users, connections and multi-modal contents on heterogeneous graphs. In *Proceedings of the 30th ACM International Conference on Multimedia, MM '22*, page 4072–4082, New York, NY, USA, 2022. Association for Computing Machinery.
- Chenwei Lin, Hanjia Lyu, Xian Xu, and Jiebo Luo. INS-MMBench: A comprehensive benchmark for evaluating lvlms' performance in insurance. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, Honolulu, HI, USA, October 2025. IEEE Computer Society.