Research Statement

Zonghan Zhang

Overview and Vision

Artificial intelligence systems increasingly support high-stakes decisions in public health, information ecosystems, infrastructure, and collective behavior. These systems operate in complex socio-technical environments where data are limited, observational, and structurally dependent. In such settings, prediction alone is insufficient. Reliable decision-making requires models that quantify uncertainty, expose structural dependencies, and remain interpretable under intervention.

My research develops uncertainty-aware and interpretable machine learning frameworks for modeling and intervening in complex networked systems. I focus on how information, risk, influence, and resources propagate through interconnected environments, and how principled modeling can support robust and accountable decision-making. Rather than treating graph-based processes as isolated algorithmic problems, I view them as structured dynamical systems whose behavior must be understood, validated, and controlled under uncertainty.

Networked propagation processes provide a fundamental testbed for this vision. Epidemics, misinformation cascades, mobility dynamics, and biological interactions all exhibit directional influence, nonlinear accumulation, and interdependent feedback across layers. Yet existing machine learning approaches often rely on black-box approximations or domain-specific heuristics, limiting interpretability and transferability across related systems.

My work addresses this gap through three complementary perspectives:

- **Uncertainty-aware modeling**: developing probabilistic graph-based methods that quantify structural and epistemic uncertainty in high-dimensional networked systems.
- **Interpretable structural representations**: designing tensorized and decomposition-based frameworks that expose higher-order interactions and cross-layer dependencies.
- **Mechanistic and physically informed learning**: integrating conservation principles, forcing effects, and structured inductive biases into graph-based AI models to improve reliability and generalization.

My earlier research established algorithmic and theoretical foundations in influence control, Bayesian optimization for inverse problems on graphs, and higher-order decomposition of spreading processes. I have also developed principled sensitivity analysis tools grounded in information decomposition theory, demonstrating how uncertainty and interaction effects can be rigorously quantified in complex models. Building on this foundation, my dissertation advances physically informed diffusion models, multilayer tensor representations, and unified flow formulations that connect replicable processes, such as epidemics, with non-replicable transport dynamics, such as traffic.

My long-term vision is to establish a principled foundation for reliable and interpretable AI systems that model, infer, and guide interventions in complex socio-technical environments.

Past Works

Uncertainty-Aware Inference and Decision in Networked Systems

A central theme of my research is the design of probabilistic methods that enable reliable inference and intervention under limited, partially observed data.

My early work examined how constrained interventions shape collective cascades in networked diffusion processes. This research revealed that global outcomes are highly sensitive to structural dependencies that are not captured by isolated node-level metrics, motivating a shift from purely algorithmic optimization toward principled uncertainty-aware modeling.

To address the computational burden of intervention design, I contributed to Neural Tangent Bayesian Optimization methods for influence maximization (ICTAI 2024), replacing simulation-heavy estimation with kernel-based surrogate modeling. In a first-author AAAI 2024 paper, I developed a graph Bayesian optimization framework for multiple-source localization from a single snapshot of spread. By defining Gaussian process priors over graph domains and incorporating structure-aware kernels, the method performs efficient inference in high-dimensional combinatorial spaces. These contributions establish probabilistic graph modeling and Bayesian decision-making as core tools in my broader agenda of uncertainty-aware AI.

Structural Decomposition and Attribution Theory for Interpretable AI

Reliable AI systems must not only predict accurately but also clarify how structural interactions shape outcomes. A second thread of my work develops formal tools for decomposing and quantifying feature-level contributions in the presence of interaction and redundancy.

In a first-author paper published at IEEE BigData 2025, I analyzed Sobol's total indices and Shapley values through the framework of Partial Information Decomposition (PID). I established a principled definition of true performance loss under feature exclusion in terms of unique and synergistic information, and proved that Sobol's total indices more faithfully approximate this loss by excluding redundancy while fully accounting for interaction effects. This work connects uncertainty quantification, variance-based sensitivity analysis, and information-theoretic decomposition, providing a rigorous foundation for interpretable and reliable machine learning in high-dimensional settings.

Complementing this theoretical analysis, my SDM 2023 paper introduced a higher-order decomposition framework that separates first-order effects from synergistic interactions in diffusion systems. Together, these contributions advance a unified objective: designing AI models whose structural dependencies can be formally decomposed, whose uncertainty can be quantified, and whose decisions remain interpretable under intervention.

Mechanistic and Structured Learning in Scientific and Societal Systems

To ensure that uncertainty-aware and interpretable methods remain grounded in real-world impact, I have applied structured graph learning to scientific and spatial systems.

In EMBC 2024, I contributed to graph neural network models for antimicrobial resistance prediction, demonstrating how domain-informed structure improves generalization. In SIGSPATIAL 2025, I co-developed a graph symbolic regression framework to derive interpretable expressions for epidemic spread, illustrating how data-driven learning can recover mechanistic patterns from spatial propagation data.

Across these projects, my research emphasizes that interpretability and uncertainty quantification are not auxiliary features, but essential properties of AI systems deployed in public health, biological, and socio-technical environments.

Current and Ongoing Work: Toward Reliable and Structured AI for Complex Dynamics

My dissertation advances a unified perspective on learning and controlling complex dynamics in socio-technical systems. This work integrates mechanistic priors, structured representations, and uncertainty-aware reasoning to improve the reliability and interpretability of AI models operating under limited data and cross-system dependencies.

Reliability-Aware Dynamic Modeling under Uncertainty

Classical diffusion and propagation models often rely on simplified stochastic assumptions that ignore propagation speed, delayed response, or structural constraints. I am developing physically informed graph-based models that incorporate structured inductive biases, such as conservation principles, forcing effects, and attenuation dynamics, into learnable operators.

Rather than treating the physical structure as a rigid constraint, I treat it as a regularizing prior that improves stability, interpretability, and generalization under distributional shifts. These models interpolate between purely data-driven learning and mechanistic formulations, enabling uncertainty-aware prediction and intervention in epidemic spread, misinformation dynamics, and other networked processes.

Interpretable Cross-Layer Representation Learning

Many socio-technical systems are inherently multilayered: communication influences epidemic spread, mobility affects resource allocation, and information dynamics shape behavioral response. I am developing tensorized representations that preserve layer-specific structure while learning cross-layer dependencies through alignment operators and low-rank decompositions.

These structured representations expose higher-order interaction effects across layers, enabling interpretable analysis of how uncertainty propagates between interconnected systems. This approach supports decision-making in settings in which interventions in one network affect outcomes in another, such as epidemic–communication coevolution.

Unified Modeling Principles for Robust Intervention

A broader objective of my work is to construct modeling principles that generalize across different types of propagation processes. I am developing a unified formulation that introduces production and attenuation terms to bridge replicable processes, such as epidemics, and non-replicable transport processes, such as traffic flow.

This unified view is not merely conceptual; it provides a shared mathematical language for designing learning algorithms that remain stable, interpretable, and uncertainty-aware across domains. By formalizing common structural patterns, this framework supports transferable modeling and robust intervention design in complex socio-technical environments.

Six-Year Research Agenda: Foundations of Reliable and Interpretable AI for Socio-Technical Systems

Over the next six years, I will develop a cohesive research program centered on uncertainty-aware and interpretable machine learning for complex socio-technical systems. This agenda advances along three tightly integrated directions that move from foundational modeling principles to cross-layer reasoning and, ultimately, decision-centric deployment.

First, I will develop learnable and structured operators for reliable AI systems. A central scientific challenge in modern machine learning is ensuring stability and robustness under distribution shift and partial observability. Building on my work in physically informed diffusion and structured modeling, I will design graph-based operators that integrate mechanistic priors, such as conservation and attenuation principles, with neural parameterizations. Rather than treating these principles as rigid constraints, I will use them as structured inductive biases that improve generalization, interpretability, and uncertainty calibration in epidemic, information, and infrastructure dynamics.

Second, I will study uncertainty propagation and cross-layer interaction in interconnected socio-technical systems. Many real-world environments involve tightly coupled layers, such as communication and epidemic networks or mobility and resource systems. Extending my tensorized multilayer framework, I will develop interpretable representation-learning methods that quantify higher-order cross-layer dependencies and formally characterize how uncertainty propagates across interacting systems. This line of work supports robust intervention planning in settings where actions in one layer alter outcomes in another.

Finally, I will translate these modeling advances into decision-centric AI tools for public health and societal resilience. Through collaboration with domain experts, I will apply structured and uncertainty-aware models to epidemic surveillance, misinformation mitigation, and infrastructure planning. These applications will serve both as validation platforms and as drivers of new theoretical questions about accountability, robustness, and transparency in high-stakes AI systems.

Together, these directions form a unified program that connects theoretical foundations, structured learning, and societally grounded impact, positioning my research at the intersection of trustworthy AI and complex networked systems.

Lab Vision and Broader Impact

As a faculty member, I plan to establish a research group focused on uncertainty-aware and interpretable AI for networked socio-technical systems. The lab will integrate graph learning, probabilistic modeling, tensor methods, and structured operator design. Students will be trained to connect theoretical rigor with societal relevance, developing AI systems that are transparent, accountable, and robust under uncertainty.

More broadly, my research seeks to advance trustworthy AI systems capable of operating responsibly in high-stakes socio-technical environments. By integrating structural modeling with formal uncertainty quantification, I aim to reduce overconfidence in automated decision systems and improve transparency in intervention planning. These advances have implications for public health response, information integrity, and infrastructure resilience. Through interdisciplinary collaboration and student training, my work contributes to the development of AI systems that are not only technically powerful, but also reliable, interpretable, and aligned with societal needs.