

Generative AI for High-Stakes Decisions with Societal Impact

Addressing today's most critical societal problems, such as pandemic response and biodiversity conservation, requires intelligent systems capable of making reliable decisions under evolving uncertainty. For example, during the COVID-19 pandemic, governments had to allocate limited vaccines and hospital beds urgently while facing rapidly changing infection rates. These decisions had immediate consequences for public health, and mistakes could lead to significant loss of life. While optimization and reinforcement learning (RL) are commonly applied to one-time and sequential decisions, their deployment in societal contexts presents key challenges: (1) societal decisions are high-risk and require accurate uncertainty quantification in dynamic environments, which is difficult due to the large spatial and temporal scales involved; (2) timely decisions are essential for effective interventions, but traditional optimization methods are costly to re-solve in large action spaces; and (3) data collection is expensive, yet RL relies on extensive interactions to learn effective policies.

I believe that recent advances in generative AI can fundamentally enhance decision-making in societal applications. First, generative models learn full probability distributions over high-dimensional data, making them well-suited for modeling uncertainty in large-scale, dynamic environments. Second, by reframing optimization as a generative process, these models can directly learn a probabilistic mapping from problem parameters to high-quality decisions, providing an efficient alternative to repeated solver calls. Third, they capture complex data patterns that serve as data-driven priors to improve sample efficiency in future decisions.

Building on these foundations, my research focuses on **advancing generative AI with optimization and reinforcement learning** to enable reliable decision-making in high-stakes societal settings. The key technical contributions of my work are summarized below (Fig. 1):

- **Uncertainty-aware forecasting with generative models for robust decisions.** I integrated *flow matching* with domain knowledge to generate uncertainty-aware forecasts that remain reliable even under sparse data and noisy observations [1]. I further proposed the *first robust optimization framework* in which the uncertainty set is induced by *diffusion models* and supported by theoretical guarantees [2].
- **Scalable and expressive decision-making through generation.** I used *energy-based models* to reformulate optimization as a generation task. This perspective yielded up to $100\times$ *speedup* in model training [3], and enables optimization over constraints that are difficult to specify explicitly [4].
- **Generative models as priors for sample-efficient RL.** To improve the sample efficiency of RL, I leveraged *flow matching* to extract prior knowledge from pre-collected offline data. By exploiting its connection to optimal transport, I developed a method that selectively incorporates useful prior knowledge during training, even when the offline data exhibit distributional shifts in environmental dynamics [5].

To ensure real-world impact, my work addresses pressing societal challenges in **environmental sustainability** and **public health**. I work closely with practitioners and deploy models in collaboration with WCS and WWF, co-designing predictive patrol strategies to reduce poaching of endangered species. This includes the **first application of generative AI for poaching risk prediction** [1]. I also partner with ARMMAN, a large-scale maternal and child health program that reaches over 40 million beneficiaries across India, to increase care engagement and reduce maternal and infant mortality through mobile health interventions [6]. My research has also been honored with the Otto & Jenny Krauss Fellowship from Georgia Tech.

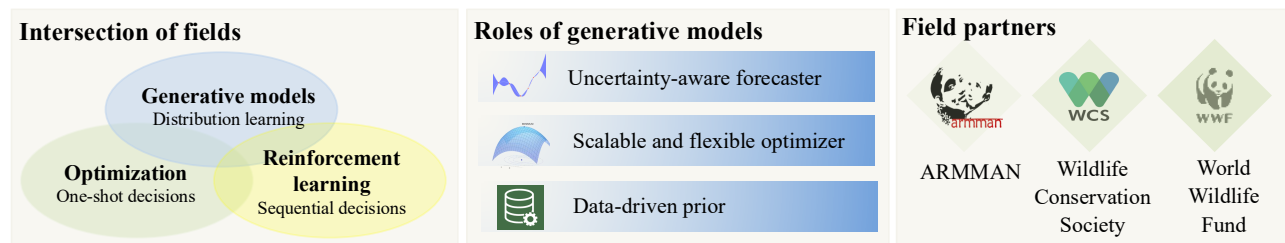


Figure 1: My research advances generative AI to enable reliable decision-making in high-stakes societal settings. I work closely with field partners to translate algorithmic innovations into real-world impact.




Figure 2: Predicting Poaching Risk with Noisy Detection. We integrate flow matching with an ecological detection framework to account for imperfect detection in poaching data. (a) True underlying poaching risk estimated via flow matching. (b) Probability of actual detection by rangers, accounting for detection uncertainty. (c) Rangers conducting patrols in the national park. Photo credit: Wildlife Conservation Society (WCS).

1 Uncertainty-aware Forecasting with Generative Models for Robust Decisions

Proactive decision-making often requires first forecasting future environmental conditions and then optimizing actions accordingly. For example, conservation agencies must predict poaching risk to allocate limited patrol resources effectively. Such a societal challenge involves high stakes, making accurate uncertainty quantification essential for effective risk management [7]. My research advances this goal by developing generative models for uncertainty-aware forecasting under noisy and limited data, and by designing robust optimization frameworks that leverage these forecasts to support reliable decision-making.

Adversarial behavior prediction through flow matching for conservation. Adversarial behaviors in conservation, such as poaching and illegal fishing, are complex and adaptive. While flow matching provides a powerful approach for modeling high-dimensional behavior distributions, real-world conservation data pose two key challenges: imperfect detection and limited data. To address these, I integrated flow matching with an ecological occupancy model and trained the flow in latent space to better infer true adversarial activity. I further incorporated domain knowledge by initializing the flow with a linear ecological prior to enhance generalization with limited data. In collaboration with the Uganda Wildlife Authority and WCS, our method improved poaching prediction accuracy by **10%** across Murchison Falls and Queen Elizabeth National Parks. *To our knowledge, this is the first application of generative AI for forecasting adversarial behavior in conservation [1].*

Probabilistic time-series forecasting for public health. Time-series forecasting plays a critical role in public health decision-making. In collaboration with the CDC, we developed a non-parametric framework that combines deep sequential models with neural stochastic processes to produce well-calibrated probabilistic forecasts of influenza [8]. I further disentangled epistemic uncertainty (arising from limited data) and aleatoric uncertainty (arising from inherent randomness) [9], and demonstrated with ARMMAN that this separation improved predictions of maternal engagement in mobile health programs [10].

To support structured probabilistic forecasting in more complex settings, I extended these methods to handle complex data types, including spatial graphs [11, 12], hierarchical geospatial structures [13], and multi-modal time-series data with heterogeneous signals [14, 15].

Robust optimization with diffusion models. To translate generative forecasts into reliable decisions, *I proposed the first robust optimization framework that constructs uncertainty sets directly from diffusion models, with formal guarantees [2].* The problem is framed as a zero-sum game between a decision-maker and nature, where nature selects worst-case forecasts from a neighborhood around the diffusion model’s forecasted distribution. A key challenge is that diffusion models represent complex, high-dimensional distributions without closed-form expressions, rendering standard robust optimization inapplicable. To address this, I developed a novel algorithm that integrates diffusion-based sampling with twisted sequential Monte Carlo to efficiently approximate the worst-case distribution. I proved convergence to an ϵ -optimal solution for the decision-maker with high proba-

bility, even under limited samples, training steps, and continuous action spaces. Applied to patrol optimization in Murchison Falls National Park, this approach reduced regret by **68%**. We later introduced a lightweight score function estimator to avoid differentiating through the diffusion sampling process, significantly accelerating the training of the diffusion forecaster while preserving alignment with the downstream optimization objective [16].

2 Scalable and Expressive Decision-Making through Generation

Societal decision-making often involves extremely large action spaces. In public health, for example, allocating limited medical resources such as hospital beds across large populations and over multiple time periods results in a combinatorial number of possibilities. Traditional optimization methods solve each instance independently, which is computationally expensive at scale. Moreover, in many real-world applications, constraints are often implicit or difficult to formalize analytically, complicating direct optimization. To address these challenges, I have developed methods that transform optimization into a generative modeling task, enabling scalable decision-making and allowing optimization over constraints that are otherwise hard to specify explicitly.

Scalable optimization with energy-based models. A key challenge in end-to-end decision-making is the mismatch between forecasting and optimization objectives: forecasting models minimize prediction error, while optimization aims to improve decision quality. Existing predict-then-optimize methods address this by repeatedly differentiating through an optimization solver using KKT conditions, which is computationally expensive at scale. To address this, I developed SO-EBM [3], a scalable alternative that reframes optimization as a generative modeling task. SO-EBM uses an energy-based model to learn the conditional distribution over decisions given problem parameters by maximizing the likelihood of precomputed optimal actions, with a regularization term to shape the decision landscape. This eliminates solver-based differentiation and can achieve up to **100× faster training**. I further extended this approach to settings with extreme uncertainty [17].

Handling unknown constraints in optimization with diffusion models. In many real-world applications, the constraints of an optimization problem are difficult to specify explicitly, yet violating them may result in unsafe or infeasible solutions. Existing methods often rely on derivative-free optimization combined with a constraint oracle, but these approaches are hard to scale in high-dimensional spaces and may not be applicable when an oracle is unavailable. To address this, I proposed DiffOPT, a diffusion-based method that learns the manifold of valid solutions directly from high-dimensional data [4]. The constrained optimization problem is reformulated as a sampling task from the product of a Boltzmann distribution over the objective and the learned manifold. This formulation enables robust decision-making in large action spaces without relying on hand-crafted constraints, thereby extending generative optimization to more realistic and complex settings.




Figure 3: DiffOPT optimizes the objective $h(x)$ under constraints C that are difficult to specify analytically, by learning them from data using a *diffusion model*.

3 Generative Models as Priors for Sample-Efficient Reinforcement Learning

Training sequential decision policies using RL often requires extensive interaction with the environment to support effective exploration. In societal settings, such interactions are costly. For example, in conservation planning, national parks may have only a small number of rangers available for patrol. A common strategy to improve sample efficiency is to incorporate offline data collected under prior policies into training. However, such data may come from different time periods or regions, and shifts in environmental dynamics can limit its applicability. To use offline data effectively, it is essential to estimate the discrepancy between offline dynamics (captured in historical data) and online dynamics (encountered during deployment) and filter out mismatched samples. This remains challenging, as online dynamics are difficult to learn accurately from limited interactions.

To address this challenge, I proposed using flow matching as an informative prior to estimate the dynamics gap more data-efficiently [5]. Instead of initializing from a random base distribution (e.g., a standard Gaussian), my method first trains a flow model on offline data to serve as an informative base, which is then used to train a second flow that models the online dynamics. *I theoretically showed that the learned base flow captures richer prior knowledge, improving the generalization of online dynamics.* Furthermore, training the second flow under an optimal transport coupling allows its average displacement to reflect the Wasserstein distance between offline and online dynamics. This enables the selective incorporation of useful prior knowledge from shifted offline data and significantly enhances sample efficiency. Finally, we introduced a decision-aware measure of the dynamics gap that captures both distributional differences and their impact on decision quality [18].

4 Future Vision

My prior work has laid the foundation for connecting generative models with decision-making. Building on this, my long-term goal is to develop generative AI systems capable of autonomously supporting complex decision-making in collaboration with public-sector partners. Such collaborations often require repeated efforts to contextualize challenges and co-design tailored solutions. Developing AI systems capable of navigating this process autonomously could significantly accelerate impact, particularly in low-resource settings where expert capacity is limited. To advance this vision, I plan to pursue the following directions:

Multiagent systems for autonomous collaboration with the public sector. Generative foundation models, such as large language models (LLMs), offer new opportunities for autonomous interaction with stakeholders through natural language. Building on our recent work on LLM agents [6, 19], I aim to advance from single-agent systems to coordinated teams of specialized agents that support complex societal decision-making. In this framework, agents take on roles such as coordinator, planner, analyst, and verifier. The coordinator manages stakeholder interaction and inter-agent communication; the planner formulates and decomposes decision problems; the analyst simulates outcomes; and the verifier ensures feasibility and compliance. I am currently working on learning role-specific policies using multi-agent RL. Looking ahead, I will design communication protocols and dynamic role-allocation mechanisms that reassign roles at runtime to enhance coordination.

Human-AI collaboration in high-stakes societal domains. Although generative AI has made significant progress, human participation remains essential in societal decision-making, particularly when safety is a major concern. Humans may intervene when needed or provide input during interactions. Building on my prior work in uncertainty quantification [7, 20], I will explore principled deferral policies that delegate decisions to humans when model confidence is low or potential risk is high. Equally important is modeling human behavior under bounded rationality. Diffusion models offer a promising direction, as they can capture complex, high-dimensional probabilistic behavior. Together, these efforts aim to develop AI systems that serve not as replacements, but as reliable and trustworthy collaborators in high-stakes environments.

Aligning generative decisions with pluralistic values of society. As generative AI systems increasingly influence decisions affecting broad populations, a key challenge is ensuring they reflect society’s diverse values. Existing methods often optimize for average preferences, overlooking critical differences across communities and stakeholders. To truly serve the public good, generative models must reconcile multiple, and often conflicting, objectives. In domains such as public policy and resource allocation, effective solutions must balance economic efficiency, environmental sustainability, and societal well-being—including health, education, and community stability. Building on my prior work aligning generative models with human preferences via RL [21], I will further explore how to incorporate pluralistic values into decision-making. One promising direction is to generate a portfolio of policies that captures a spectrum of stakeholder preferences, as in my recent work [22].

While much work remains to be done, I believe generative AI can be a transformative tool for high-stakes decision-making. **Its impact can extend far beyond commercial success to meaningfully benefit society.**

References

- [1] **Lingkai Kong***, Haichuan Wang*, Charles A Emogor, Vincent Börsch-Supan, Lily Xu, and Milind Tambe. Generative ai against poaching: Latent composite flow matching models for wildlife conservation. *Innovative Applications of Artificial Intelligence (IAAI)*, 2026.
- [2] **Lingkai Kong**, Haichuan Wang, Yuqi Pan, Cheol Woo Kim, Mingxiao Song, Alayna Nguyen, Tonghan Wang, Haifeng Xu, and Milind Tambe. Robust optimization with diffusion models for green security. *Uncertainty in Artificial Intelligence (UAI)*, 2025.
- [3] **Lingkai Kong**, Jiaming Cui, Yuchen Zhuang, Rui Feng, B Aditya Prakash, and Chao Zhang. End-to-end stochastic optimization with energy-based model. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [4] **Lingkai Kong***, Yuanqi Du*, Wenhao Mu*, Kirill Neklyudov, Valentin De Bortoli, Dongxia Wu, Haorui Wang, Aaron M Ferber, Yi-An Ma, Carla P Gomes, and Chao Zhang. Diffusion models as constrained samplers for optimization with unknown constraints. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2025.
- [5] **Lingkai Kong***, Haichuan Wang*, Tonghan Wang, Guojun Xiong, and Milind Tambe. Composite flow matching for reinforcement learning with shifted-dynamics data. *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
- [6] Shresth Verma, Alayna Nguyen, Niclas Boehmer, **Lingkai Kong**, and Milind Tambe. Priority2reward: Incorporating healthworker preferences for resource allocation planning. *AAAI Conference on Artificial Intelligence (AAAI)*, 2025.
- [7] **Lingkai Kong**, Harshavardhan Kamarthi, Peng Chen, B Aditya Prakash, and Chao Zhang. Uncertainty quantification in deep learning. *Conference on Knowledge Discovery and Data Mining (KDD)*, 2023.
- [8] Harshavardhan Kamarthi, **Lingkai Kong**, Alexander Rodriguez, Chao Zhang, and B Aditya Prakash. When in doubt: Neural non-parametric uncertainty quantification for epidemic forecasting. *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [9] **Lingkai Kong**, Jimeng Sun, and Chao Zhang. Sde-net: Equipping deep neural networks with uncertainty estimates. *International Conference on Machine Learning (ICML)*, 2020.
- [10] Sarah Martinson, **Lingkai Kong**, Cheol Woo Kim, Aparna Taneja, and Milind Tambe. Llm-based agent simulation for maternal health interventions: Uncertainty estimation and decision-focused evaluation. *AAMAS Workshop on Autonomous Agents for Social Good (AAMAS-AASG)*, 2025.
- [11] **Lingkai Kong***, Haotian Sun*, Yuchen Zhuang, Haorui Wang, Wenhao Mu, and Chao Zhang. Two birds with one stone: Enhancing uncertainty quantification and interpretability with graph functional neural process. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2024.
- [12] Yinghao Li, **Lingkai Kong**, Yuanqi Du, Yue Yu, Yuchen Zhuang, Wenhao Mu, and Chao Zhang. Muben: Benchmarking the uncertainty of molecular representation models. *Transactions on Machine Learning Research (TMLR)*, 2024.
- [13] Harshavardhan Kamarthi, **Lingkai Kong**, Alexander Rodríguez, Chao Zhang, and B Aditya Prakash. When rigidity hurts: Soft consistency regularization for probabilistic hierarchical time series forecasting. *Conference on Knowledge Discovery and Data Mining (KDD)*, 2023.
- [14] Harshavardhan Kamarthi, **Lingkai Kong**, Alexander Rodríguez, Chao Zhang, and B Aditya Prakash. Camul: Calibrated and accurate multi-view time-series forecasting. *The Web Conference (WWW)*, 2022.
- [15] Haoxin Liu, Shangqing Xu, Zhiyuan Zhao, **Lingkai Kong**, Harshavardhan Prabhakar Kamarthi, Aditya Sasanur, Megha Sharma, Jiaming Cui, Qingsong Wen, Chao Zhang, et al. Time-mmd: Multi-domain multimodal dataset for time series analysis. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [16] Zihao Zhao, Christopher Yeh, **Lingkai Kong**, and Kai Wang. Diffusion-dfl: Decision-focused diffusion models for stochastic optimization. *Under review*, 2026.
- [17] **Lingkai Kong**, Wenhao Mu, Jiaming Cui, Yuchen Zhuang, B. Aditya Prakash, Bo Dai, and Chao Zhang. DF²: Distribution-free decision-focused learning. *Uncertainty in Artificial Intelligence (UAI)*, 2025.
- [18] Paula Rodriguez-Diaz, **Lingkai Kong**, Kai Wang, David Alvarez-Melis, and Milind Tambe. What is the right notion of distance between predict-then-optimize tasks? *Uncertainty in Artificial Intelligence (UAI)*, 2025.

- [19] Haotian Sun, Yuchen Zhuang, **Lingkai Kong**, Bo Dai, and Chao Zhang. 2023. Adaplaner: Adaptive planning from feedback with language models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [20] **Lingkai Kong**, Haoming Jiang, Yuchen Zhuang, Jie Lyu, Tuo Zhao, and Chao Zhang. Calibrated language model fine-tuning for in-and out-of-distribution data. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [21] **Lingkai Kong**, Haorui Wang, Wenhao Mu, Yuanqi Du, Yuchen Zhuang, Yifei Zhou, Yue Song, Rongzhi Zhang, Kai Wang, and Chao Zhang. Aligning large language models with representation editing: A control perspective. *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [22] Cheol Woo Kim, Jai Moondra, Shresth Verma, Madeleine Pollack, **Lingkai Kong**, Milind Tambe, and Swati Gupta. Navigating the social welfare frontier: Portfolios for multi-objective reinforcement learning. *International Conference on Machine Learning (ICML)*, 2025.