GENERATIVE MODELING FOR HUMAN BEHAVIOR IN AUTONOMOUS SYSTEMS

by

Juanwu Lu

A Dissertation

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy



Lyles School of Civil and Construction Engineering

West Lafayette, Indiana

August 2026

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL

Dr. Ziran Wang, Chair

Lyles School of Civil and Construction Engineering, Purdue University

Dr. Joseph G. Makin

Elmore Family School of Electrical and Computer Engineering, Purdue University

Dr. Ruqi Zhang

Department of Computer Science, Purdue University

Dr. Maria Laura Delle Monache

Department of Civil and Environmental Engineering, University of California, Berkeley

Approved by:

Dr. Buck Doe

To my family and friends, whom I love, admire, and follow.

ABSTRACT

The influx of successful autonomous systems deployments has created environments in which machines must continuously interact with humans. As a result, it is essential for modeling human behavior in the physical world for machines to better anticipate and respond to interactions. In transportation, for instance, drivers approaching intersections, changing lanes, or yielding to pedestrians do not adhere to predetermined trajectories. Instead, their behavior results from latent intentions, environmental constraints, and social interactions. Accurately modeling such behavior for forecasting and simulation requires generative models, which are statistical or machine-learning frameworks capable of producing diverse outcomes. These models must capture multimodal futures, quantify uncertainty, adapt in a closed loop, and provide interpretable behavioral factors for manipulation and evaluation.

In this dissertation, we approach generative modeling of human behavior using probabilistic models that learn to maximize the likelihood of plausible futures. We begin by presenting *deep variational Bayesian methods*, which approximate intractable posterior with neural networks, to learn the conditional distributions of destinations and trajectories from log-replayed motion data. These models generate interpretable multimodal forecasts and provide predictive uncertainty estimates, facilitating the identification of potential distribution shifts. Building on these foundations, we present a hierarchical model for *multi-agent simulation* that organizes factors such as intentions and driving profiles into different layers. This approach transitions from open-loop to closed-loop generation, enabling simulated agents to respond recursively within scenarios. We further extend the idea in building controllable simulations, in which agent behavior is represented by *reward-space latent variables*, allowing users to guide generative models that reflect varying degrees of efficiency, safety, and map compliance. Finally, we reflect on the existing evaluations of multi-agent simulation, noting that standard marginal metrics may overlook *coupling among interactive agents*. We conclude the work with a discussion of open questions and potential in physical-world modeling, highlighting both the promise and limitations of current generative approaches to modeling human behavior in complex interactive environments.