Create a Digital Population at Scale with Large Language Models (LLMs) and Generative AI through the CrowdLLM Framework

BIOGRAPHY

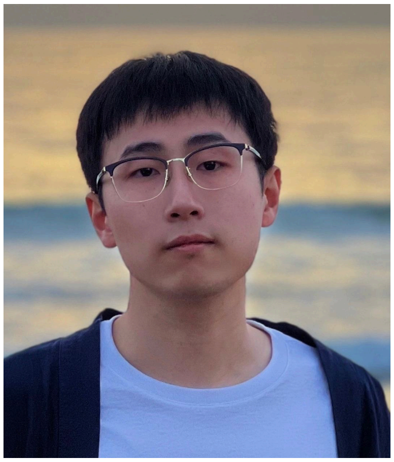

Ryan Lin is currently a Ph.D. candidate (advised by Prof. Shuai Huang) in the department of Industrial and Systems Engineering at the University of Washington. His research interests lie at the intersection of artificial intelligence (AI), statistics, and quality science, where he develops computational methodologies for model personalization, uncertainty quantification and design of experiments, with applications to healthcare, medicine and transportation. He aims to build trustworthy and personalized AI systems that integrate human consideration and expert knowledge into system-level models to achieve quality user-system interactions. Ryan received a B.S. degree in Statistics, and an M.S. degree in Computer Science, from the University of Science and Technology of China. He is a recipient of Meta AI Fellowship and finalists of the Best Paper Competition at INFORMS Quality, Statistics, and Reliability (QSR) in 2023 and 2025.

ABSTRACT

Many applications, such as crowdsourcing, voting, and surveys, have long relied on the collective intelligence of human crowds. Harnessing the collective intelligence from a diverse population is often costly and logistically challenging. This raises a natural question: Can we build digital populations that not only simulate human decision-making at a microscopic level but also scale to exhibit macroscopic collective intelligence? Such synthetic populations hold transformative potential across domains including social simulation, behavioral science, marketing, and recommender systems by dramatically reducing costs and mitigating the challenges of recruiting and managing human participants. Recent advances in large language models (LLMs) offer new opportunities for constructing such populations, but existing approaches relying solely on LLMs often fall short in capturing the nuanced accuracy, diversity, and emergent dynamics characteristic of real human collectives. To address these limitations, we present **CrowdLLM**, a framework that combines pretrained LLMs with generative modeling techniques to more faithfully reproduce heterogeneous and human-like decision-making tailored for realistic applications at scale. Through extensive evaluations across diverse decision-making tasks, CrowdLLM consistently outperforms state-of-the-art baselines, achieving substantial improvements in a range of performance metrics. These results highlight the promise of CrowdLLM as a foundation for scalable, collective intelligence systems and as a powerful tool for studying and harnessing digital populations in many real-world applications where the collective intelligence of a diverse population is needed.