Title: What can we learn from a billion agents?

github.com/AgentTorch/AgentTorch

Abstract: From pandemic response to supply chain resilience, today's critical challenges emerge from the complex interactions of millions of autonomous agents. Traditional agent-based simulations struggle to scale beyond small populations while preserving rich agent behaviors and interaction patterns, making it difficult to model these societal-scale phenomena. This talk introduces Large Population Models (LPMs), a novel methodology that enables high-performance, privacy-preserving modeling of massive agent populations.

LPMs enable societal-scale modeling through three key innovations. First, they scale to millions of synthetic agents while preserving rich agent behaviors—from simple heuristics to language model-powered interactions—achieving unprecedented efficiency: simulating 8M agents in 5 minutes versus 50 hours in traditional approaches. Second, LPMs maintain end-to-end differentiability of the simulation dynamics, enabling use of gradients to automatically calibrate against heterogeneous real-world data sources, compose with neural networks and perform rapid sensitivity analysis without repeated simulation—accelerating model tuning by 8300×. Third, LPMs bridge simulation and reality through secure multi-party computation, enabling real-world agents to participate in decentralized simulations while preserving the privacy of their states and interactions - effectively “backpropagating through reality”.

These capabilities are unified in AgentTorch, our open-source framework for composable agent simulation at scale. We demonstrate AgentTorch's impact through high-stakes applications: optimizing vaccine distribution across populations of 5+ million, safeguarding billion-dollar food supply chains, and modeling disease spread in dense urban environments of 8+ million residents. As we scale towards a world with billions of agents, LPMs establish a foundation for next-generation high-performance computing platforms that can help address national-scale challenges while preserving data privacy and security.

About Me: I am a PhD student at MIT, where I lead research on Large Population Models (lpm.media.mit.edu) under Prof. Ramesh Raskar's supervision. My work bridges theoretical advances in multi-agent AI with real-world impact - reaching over 20 million people through deployments across multiple countries and garnering coverage in global press. I have published over 50 papers in leading conferences and journals, including AAMAS, CVPR, KDD, British Medical Journal; earning best paper awards and contributing to 25 patents. My research experience spans academic and industry labs, including Mayo Clinic and JP Morgan AI Research. Prior to MIT, I was a scientist at Adobe where I received the Outstanding Young Engineer Award for my work on collaborative machine learning. I have co-organized workshops and tutorials on multi-agent systems at ICLR (2021, 2023) and AAMAS (2024). I earned my MS from Massachusetts Institute of Technology and BE from Delhi College of Engineering.