Large Population Models

Differentiable and Private Learning over Million Autonomous Agents

Platform: github.com/AgentTorch/AgentTorch

Talks: Academic (45 minutes); Enterprise (5-min)

Many of today's critical challenges—from pandemic response to supply chain resilience—emerge from the complex interactions of millions of autonomous agents. While current approaches can model sophisticated individual behaviors, they fail to capture the emergent phenomena that arise in real-world populations. This limitation stems from three fundamental challenges: defining and executing protocols at societal scales, integrating real-world data effectively, and bridging the simulation-reality gap. Large Population Models (LPMs) address these challenges through differentiable protocols that enable end-to-end learning across massive synthetic and physical agent networks [1].

Our first key innovation lies in protocol specification and execution at scale. Through the AgentTorch framework [3], we introduce a formal language for defining complex interaction protocols that compose via gradients in a unified computational graph. This enables simulation of millions of agents with sophisticated architectures—from simple heuristics to language model-powered behaviors [2], Our benchmarks demonstrate unprecedented computational efficiency: 600x speedup in simulation time (5 minutes vs 50 hours for 8M agents), enabling execution of up to 300,000 agent interactions per second on commodity hardware [8].

The second breakthrough establishes differentiability through both simulation environments and agent behaviors, enabling seamless integration of heterogeneous real-world data [6]. This innovation allows composition of hybrid neural-mechanistic pipelines that combine deep learning, agent-based modeling, and differential equations [3,6]. By maintaining end-to-end gradients, we achieve an 8300x acceleration in model calibration and enable rapid sensitivity analysis through gradient computation rather than repeated simulation [7]. The practical impact was demonstrated during COVID-19, where our models informed vaccination strategies by balancing epidemiological dynamics with economic outcomes across multiple countries [11,12].

Our third contribution extends these capabilities to physical agent networks through privacy-preserving protocols. Rather than centralizing sensitive data, we decentralize the simulation itself using secure multi-party computation. [6] This enables direct participation of physical agents while preserving privacy of individual states and interaction patterns. Through novel protocols for secret sharing and gradient estimation, we achieve "backpropagation through reality"—allowing end-to-end composition of synthetic and physical agent networks.

This bridges a fundamental gap between simulation and reality, enabling learning from real-world behaviors without compromising individual privacy.

The framework has demonstrated significant real-world impact: optimizing vaccine distribution for New Zealand's 5 million citizens[9], tracking cascading disruptions across global supply chains[10,13], and modeling mobility and employment patterns for entire metropolitan areas[2,5]. Our results show that scaling multi-agent systems to societal challenges requires fundamentally new approaches to protocol design and execution. By maintaining differentiability across both synthetic and physical networks, LPMs enable systematic discovery and refinement of protocols that can help address critical collective challenges.

References

1. Social Orchestration of Agentic Systems, arxiv 25
2. On the limits of agency in agent-based models, AAMAS 2025 (Oral)
3. A Framework for Learning in Agent-Based Models, AAMAS 2024 (Oral)
4. Private Agent-based Modeling, AAMAS 2024 (Oral)
5. First 100 days of pandemic: interplay of pharmaceutical, behavioral and digital interventions, AAMAS 2024 (Oral, Best Student-Paper Runner Up)
6. Differentiable Agent-based Epidemiology, AAMAS 2023 (Oral) and ICML'W 22 (Best Paper Award)
7. Don't Simulate Twice: One-shot sensitivity analyses via automatic differentiation, AAMAS 2023 (Oral)
8. DeepABM: scalable, fast and differentiable agent-based modeling. WSC 2021 (Oral)
9. ALMA: Deploying LPMs in New Zealand. ESR 2024
10. US-scale Digital Similar of Interacting Livestock, Wild Birds and Human Ecosystems for Multi-host Epidemic Spread, arxiv 25
11. Public health impact of delaying second dose of BNT162b2 vaccine. British Medical Journal 2021
12. Using neural networks to calibrate agent-based models improves regional evidence for vaccine strategy and policy. Vaccine 2023.
13. Gradient-based Calibration of Financial Agent-based Models. ICAIF 2023