BTUT | Kernel-Weighted Mean-Field Game Simulator

Scalable Multi-Agent Equilibrium Computation via Forward Fokker–Planck Dynamics

DARPA Mathematical Challenge 13 — Scalable Game Theory Beyond Classical PDE Methods

PI: Diren Kumaratilleke, University of North Carolina at Chapel Hill

Live System: `btut.ai`

$O(N)$	$\epsilon < 0.07$	**10K+**	**Forward**
Per-Step Complexity	Empirical Nash Approximation	Validated Agent Population	Only PDE Solver

1 PROBLEM STATEMENT

DARPA Mathematical Challenge 13 asks: *“What new scalable mathematics is needed to replace the traditional PDE approach to differential games?”* Classical N-player game solvers—including iterative best-response and fictitious play—require $O(N^2)$ pairwise evaluations per step, making large populations intractable. Mean-field game (MFG) theory (Lasry & Lions 2007; Huang, Caines & Malhamé 2006) replaces individual tracking with a population density, but existing MFG solvers typically require a coupled forward–backward PDE system that is itself computationally expensive. BTUT contributes a **forward-only** Fokker–Planck formulation that achieves $O(N)$ per-step complexity and converges to approximate Nash equilibrium without a backward Hamilton–Jacobi–Bellman pass.

2 TECHNICAL APPROACH

Core Formulation. Agent populations are modeled as a probability density $\rho(x,t)$ evolving under a forward Fokker–Planck PDE:

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (v[\rho]\,\rho) + \frac{\sigma^2}{2}\Delta\rho$$

where the drift $-\nabla \cdot (v[\rho]\,\rho)$ drives exploitation toward higher-payoff strategies and the diffusion $+\frac{\sigma^2}{2}\Delta\rho$ enforces stochastic exploration to prevent premature convergence.

Kernel-Weighted Interaction. Pairwise interactions are replaced by local field sensing via a Gaussian RBF kernel, preserving $O(N)$ scaling:

$$K(x,y) = \exp\left(-\frac{\|x-y\|^2}{2\sigma^2}\right)$$

Each agent evaluates cost against a kernel-smoothed density field rather than individual opponent states, eliminating the N^2 interaction bottleneck.

3 KEY RESULT: FORWARD-ONLY CONVERGENCE

Forward-only MFG solvers—those omitting the backward HJB pass—are known to exhibit oscillatory herd dynamics that prevent stable convergence (Achdou & Capuzzo-Dolcetta 2010). BTUT stabilizes the forward pass through **(i)** momentum decay that attenuates inter-step oscillations and **(ii)** density-dependent anti-crowding corrections that redistribute mass from over-saturated regions. Empirically, the solver converges to ϵ-Nash equilibrium ($\epsilon < 0.07$) across six tested domains at populations up to $N{=}10{,}000$ agents. The $O(N)$ complexity profile is architecturally preserved at higher N; backend deployment targets $N > 10^5$.

4 CROSS-DOMAIN VALIDATION

All simulations ingest real-time data via public APIs (drift coefficients, diffusion noise scaling). Tested on $N{=}10{,}000$ agents in-browser; $O(N)$ architecture supports backend scaling beyond current validation ceiling.

Domain	Tested Scale	Live Data Sources	Capability vs. Conventional Approach
Financial Markets Portfolio Equilibria	10K agents $\epsilon{=}0.07$	Yahoo Finance, Frankfurter (ECB)	$O(N)$ vs. $O(N^2)$ scaling; rapid convergence vs. ~ 500 iter (fictitious play)
Drone Swarms Aerial Coordination	10K units	Open-Meteo, OpenSky Network	Nash-consistent collision avoidance (ORCA + MFG) vs. heuristic RVO
Traffic Networks Urban Flow	10K vehicles	Open-Meteo	City-scale adaptive routing vs. fixed-timing signal control
Smart Grid Energy Dispatch	10K nodes	Open-Meteo	Distributed MFG dispatch ($<$100 ms) vs. centralized SCADA (1–5 min)
AI Agent Systems Multi-Agent RL	10K agents	Wikipedia REST, arXiv	ϵ-Nash convergence in $\sim$10K episodes vs. $\sim$1M (MARL baseline)
Supply Chain Global Logistics	10K nodes	World Bank, REST Countries, ECB	Bullwhip variance $\sim$1.1 vs. $\sim$3–5; replanning $<$1 min vs. hours

Live interactive demonstration: `btut.ai` | Six domains | Real-time API ingestion | Full source inspection