

The City College of New York

CSC 59867-E: Senior Project II

AI Agents for Decision Making in the Real World

Quantum Control Mechanisms and Non-Deterministic AI

Saptarashmi Bandyopadhyay

By Saptarashmi Bandyopadhyay
Email: sbandyopadhyay@ccny.cuny.edu
Assistant Professor of Computer Science
City College of New York and Graduate Center at the City University of New York

October 14, 2026 CSC 59867 - E

Lecture Agenda

Distributed Swarm Kinetics and Kinematics

01

- Reynolds/Olfati-Saber flocking control
- Dynamic algebraic graph connectivity (Fiedler value)
- State augmentation under RF transmission latency

Autonomous Transportation and AVs (self-driving cars, drones, ships etc)

02

- Dec-POMDP formulation of unsignalized crossings
- High-speed platoon string stability limits
- Control Barrier Functions (CBFs) for zero-collision safety

AR/VR/XR and Hardware Agents

03

- OpenXR spatial tele-autonomy & retinal HUD vectors
- Sim-to-Real domain randomization (mass, friction, wind)
- Edge deployment under 15W SWaP budgets (Jetson Orin)

Quantum Control vis-a-vis Deterministic AI

04

- QAOA/QUBO combinatorial fleet task allocation
- Parameterized Quantum Circuits (PQCs) for policy search
- Escaping coordination traps via quantum entropy exploration

Saptarshmi Bandyopadhyay

Safety-Critical Action Pruning via Control Barrier Functions

Control Barrier Function (CBF)

- Define safe set $C = \{x \mid h(x) \geq 0\}$ (e.g. inter-vehicle distance margin)
- Forward invariance condition: $\dot{h}(x, u) \geq -\alpha(h(x))$
- Guarantees the system never leaves safe set C under valid controls
- Acts as a real-time mathematical safety envelope around the policy

QP Safety Filter & Masking

- CBF-QP Filter: $u^* = \text{argmin } 0.5||u - u_{RL}||^2$ s.t. $L_f h + L_g h \cdot u \geq -\alpha(h)$
- Projects raw RL policy outputs to closest collision-free control
- Discrete Action Masking: Pre-softmax logits set to $-\infty$ for unsafe maneuvers
- Eliminates exploratory collisions during both training and deployment

Saptarashmi Bandyopadhyay

3-Layer Hierarchical MARL Architecture for Fleets

Strategic Manager (1 Hz)

- Global digital-twin traffic awareness
- Macro corridor allocation and formation selection
- Assigns tactical sub-goals to individual vehicles

Tactical Coordinator (10 Hz)

- Multi-Agent Graph Attention Policy (GAT)
- Local conflict resolution, gap acceptance, and lane change timing
- Outputs reference acceleration and steering targets

Execution Controller (100 Hz)

- Onboard MPC / differential flatness controller
- High-frequency reference tracking on vehicle dynamics
- Enforces CBF safety limits against sudden obstacles

Cooperative Training Paradigms & Sim-to-Real Transfer

Multi-Agent Algorithms

- MAPPO (Centralized Training, Decentralized Execution): Central critic evaluates joint state; actors execute on local sensors
- QMIX with Spatial Attention: Factorizes joint action-values under monotonicity ($\partial Q_{tot} / \partial Q_i \geq 0$)
- Breaks multi-agent policy symmetry to prevent deadlocks

Sim-to-Real Robustness

- Domain Randomization: Vehicle mass (±20%), tire friction ($\mu \in [0.3, 0.9]$), wind gust perturbations
- Latency Injection: Stochastically delays observations (10–60ms) during simulation training
- Prevents policy overfitting to idealized simulator dynamics

AR/VR/XR Agents; Tackling Non-Determinism with Quantum Control

Saptarashmi Bandyopadhyay

AR/VR/XR Spatial Computing & Hardware Edge Agents

Spatial XR Human-Swarm Teaming

Bidirectional Digital Twin: Live telemetry streamed via WebRTC/ROS2 to OpenXR headsets (Vision Pro, Quest 3)

Retinal HUD Visualizations: Render invisible planned trajectories, LiDAR fields, and CBF safety boundaries

Spatial Tele-Autonomy: Gesture-based 3D waypoint placement and virtual perimeter containment

Edge Hardware Constraints

Hardware-in-the-Loop (HIL) deployment on NVIDIA Jetson AGX Orin & Hailo-8

Size, Weight, and Power (SWaP) budgets limit total power to ≤15W on drones

INT8 Quantization & TensorRT engine compilation

Guarantees deterministic <10ms onboard inference loop execution

Sattarashmi Bandyopadhyay

Quantum Control Mechanisms & Non-Deterministic AI

Bandyopadhyay

Combinatorial Task Allocation

- Fleet routing mapped to Quadratic Unconstrained Binary Optimization (QUBO)
- Solved via Quantum Approximate Optimization Algorithm (QAOA) on QPUs
- Structural polynomial speedups for high-density multi-depot vehicle routing
- Hybrid quantum-classical optimization loops (VQE/QAOA)

Quantum RL & Exploration

- Parameterized Quantum Circuits (PQCs) represent policy: $|\psi(s;\theta)\rangle = U(\theta)U_{enc}(s)|0\rangle$
- Actions sampled from Pauli expectation values: $a_i = \langle\psi|\hat{Z}_i|\psi\rangle$
- Quantum tunneling escapes local traps in non-convex multi-agent reward spaces
- Intrinsic quantum entropy sampling handles deep uncertainty

EVALUATION

Benchmark Performance: Efficiency, Safety & Latency

Intersection Delay

- -34.2%
- Reduction in crossing delay vs actuated traffic lights
- Zero vehicle deadlocks observed

Safety Violations

- 0.00%
- Zero collisions across 10^5 steps using CBF filtering
- Unconstrained PPO incurs 12.4% crash rate

Platoon Spacing

- +21.8%
- Convoy energy savings via 0.6m spacing
- String stability preserved across 10 vehicles

Edge Latency

- 6.2 ms
- Inference runtime on Jetson Orin Nano
- Comfortably within 20ms real-time control limit

Benchmark Performance: Efficiency, Safety & Latency

Intersection Delay
- -34.2%
- Reduction in crossing delay vs actuated traffic lights
- Zero vehicle deadlocks observed

Safety Violations
- 0.00%
- Zero collisions across 10^5 steps using CBF filtering
- Unconstrained PPO incurs 12.4% crash rate

Platoon Spacing
- +21.8%
- Convoy energy savings via 0.6m spacing
- String stability preserved across 10 vehicles

Edge Latency
- 6.2 ms
- Inference runtime on Jetson Orin Nano
- Comfortably within 20ms real-time control limit

Core Takeaways

- Decompose autonomy into 3 layers: strategic, tactical, and execution
- Control Barrier Functions guarantee mathematical safety over pure RL
- Multi-Agent Rollout cuts trajectory search complexity from $O(|U|^N)$ to $O(N \cdot |U|)$
- Spatial XR transforms raw sensor telemetry into actionable human-swarm oversight

Saptarashmi Bandyopadhyay

Questions?

Saptarashmi Bandyopadhyay