Traffic Optimization through Collective Intelligence

Vikram Prashant Gajendragadkar, Si Min Qiang, Gabriel Reynoso
AAI 595-B Applied Machine Learning

Acknowledgements
Dr. Hao Wang
Dr. Joseph Helsing

Introduction

In the United States, the average driver loses over 40 hours every year to traffic delays, which caused an estimated $74 billion loss in productivity in 2024[1] and greatly affects quality of life.

Can we design routing strategies that reduce overall congestion by coordinating drivers more intelligently?

In this project, we explore whether collective intelligence principles and imitation learning can improve traffic routing efficiency.

Specifically:
- Compare Shortest Path Assignment (SPA) vs. Ford- Fulkerson FFA.
- Train a reinforcement-learning policy (imitating FFA) to imitate an expert
- Evaluate how routing choices affect system travel time, flow distribution, and network congestion.

Methodology

Data & Network Construction

Extract Manhattan roadway graph from OpenStreetMaps shapefile.
Convert street geometry into a directed traffic network.
Assign stochastic free-flow speeds and edge capacities based on link length, spacing, and lane assumptions.

Routing Algorithms

SPA (Shortest Path Assignment): selfish routing using free-flow travel times.
FFA: Ford–Fulkerson–style marginal-cost routing
COIN - FF: Collective Intelligence trained Q-table policy trained via imitation learning of FFA behavior.

Evaluation Metrics

Total system travel time
Flow heatmaps
Fraction of agents reaching destination
Average path length and travel time

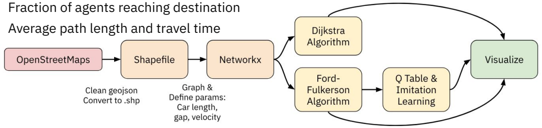

Results

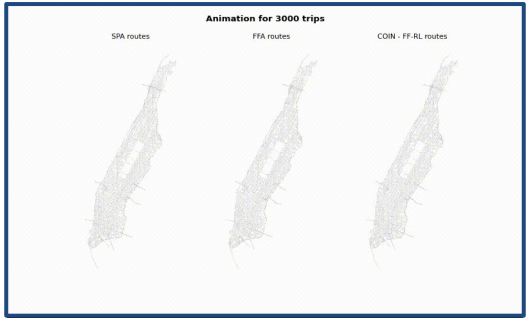

	SPA	FFA	Coin-FF
G = Vehicle Travel time (hrs)	1,325.08	2,508.73	2,097.26
Comparison vs SPA		-89.33%	-58.27%
Succes rate (Agents reaching destination)	100%	100%	100%
Average travel time per agent (min)	26.5	50.17	41.95
Total distance traveled (km)	57665.32	94359.46	59527.08
Average distance per agent (km)	19.22	31.45	19.84

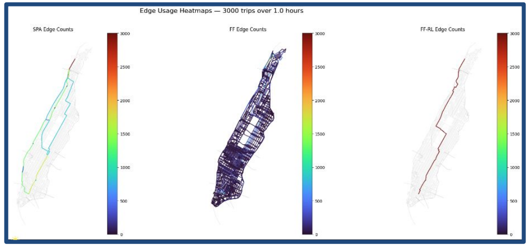

Conclusion & Discussion

Collective intelligence (**COIN**) combined with reinforcement learning is a **feasible** approach for modeling and optimizing **traffic routing** in a **fixed urban** environment. By integrating difference rewards with a Q-learning policy initialized from Ford–Fulkerson flow solutions, agents learn routing behaviors that align local decisions with a global congestion objective.

The COIN-based policy **improves average distance traveled per agent compared to shortest-path assignment**, while maintaining a 100% success rate in reaching destinations. However, under higher congestion levels, the learned policy shows **limited ability** to **redistribute** traffic across multiple routes, tending to reinforce dominant corridors even when alternative paths exist.

Overall, the results suggest that COIN-based reinforcement learning can reduce inefficiencies of selfish routing, but further improvements are needed to enhance adaptability under heavy congestion.

Finally, we believe that once more autonomous vehicles start to transit our streets, this kind of algorithm will become more relevant in order to try to find a good global optimum and reduce heavy traffic.

Future Work

Currently, we are only simulating a single point-to-point route with the assumption that all agents leave the starting point at the same time. This can be expanded to include **multi-route simulations**, and also incorporate **dynamic time frames**.

We could learn from **real-world trajectory data** (e.g., GPS traces) to see the realistic effects of road closures, congestion pricing, and other extraneous factors.

Implement a model of **heterogeneous driver behavior** and compliance rates, as driving style varies between both vehicles types and individual drivers.

For more **scalable routing**, replace Q-table with Graph Neural Networks. This way, with enough time or computing power, simulation of larger road networks should be possible.

Integrate travel time prediction (LSTM) to **anticipate future congestion** and use this as a tool for informing drivers.

STEVENS
INSTITUTE OF TECHNOLOGY

1. https://inrix.com/press-releases/2024-global-traffic-scorecard-us/
2. Wolpert, D. H., & Tumer, K. (1999). *An introduction to Collective Intelligence.*
3. Brackstone, M., & McDonald, M. (2000). *Car-following: a historical review.* Transportation Research Part F: Traffic Psychology and Behaviour.