CRII:CPS: Emerging Markets and Myopic Decision-Making in Multi-Modal Transportation Systems: Models and Validation

Lillian J. Ratliff (University of Washington)
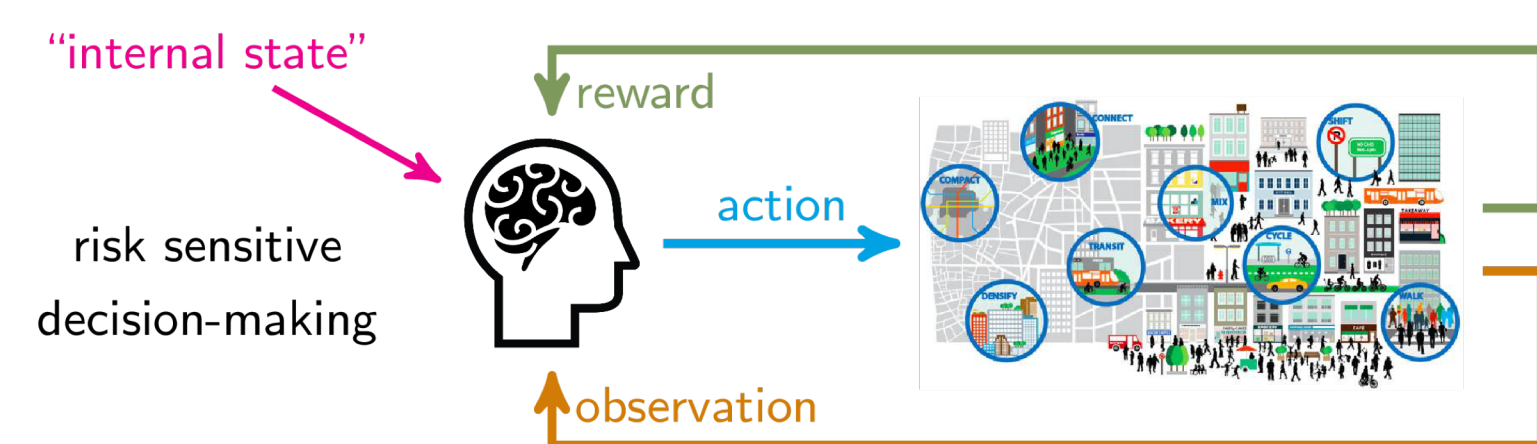
Award ID#: CNS-1656873

W

UNIVERSITY of WASHINGTON

Description

We tackle problems in urban mobility by creating high-fidelity models, validated with real-world data, of mixed-mode travel decisions and emerging mobility markets. An efficient transportation system is a crucial component of any well-functioning city. As cities expand, existing infrastructures are increasingly stressed and a **service gap** between what is needed and available is forming. At the same time, CPS and IoT technologies have enable intelligent infrastructure systems resulting in **rich data streams**, **constant connectivity**, and **emergence of innovative markets.**

Societal-Scale CPS

Recent Results

High-probability, finite time convergence guarantees for a class of gradient-based multi-agent learning algorithms (including policy gradient RL, GANs, gradient-based MABs, and online optimization) [1,2].

Regret guarantees for online incentive design algorithms under bandit feedback for matching budget constrained incentives to users with testing on bike-share supply and demand matching [3,4].

Fig. Reduction in rejections at each station

Aim 1

Learn plausible models of human decision-making and preferences with theoretical foundations by drawing on intelligent infrastructure data.
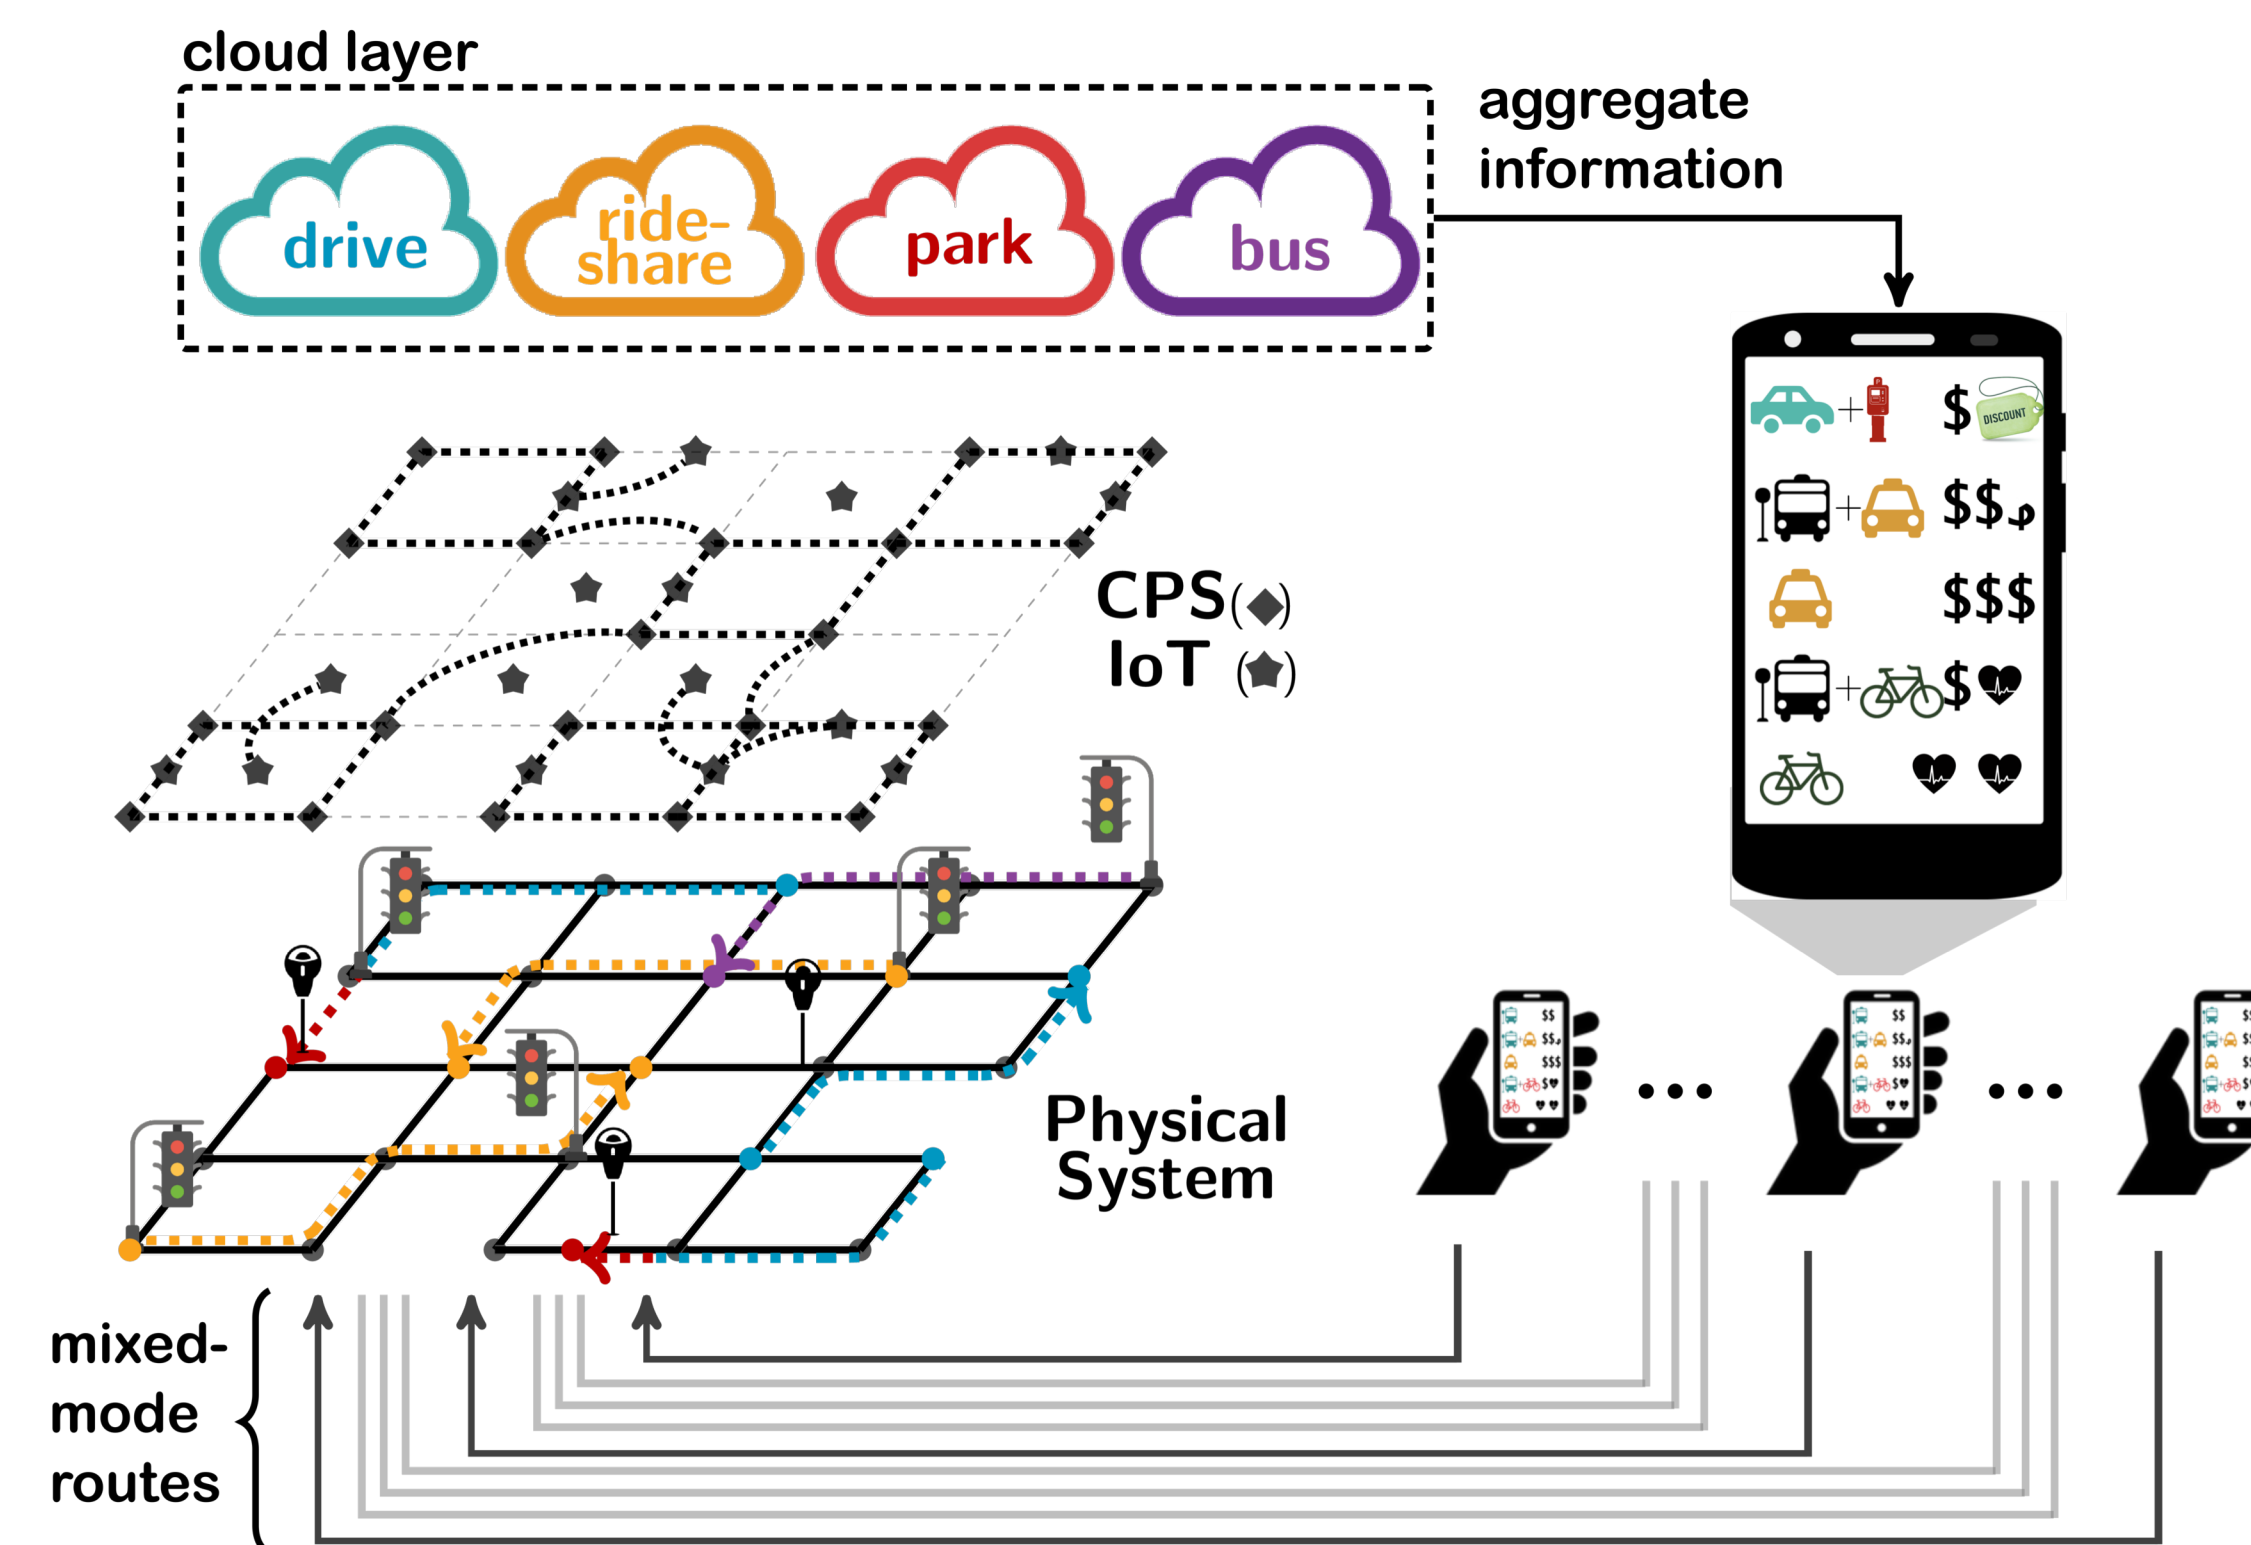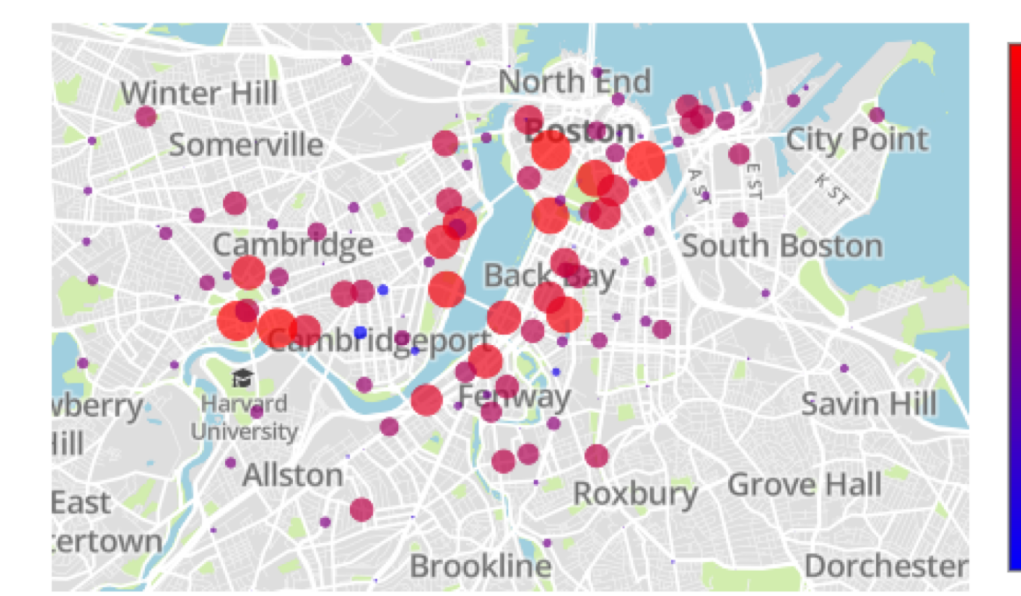
Challenges: Humans are not completely rational and face both exogenous and endogenous uncertainties. Intelligently augmented humans (via system-level CPS/IoT infrastructure and personal devices) can make and revise decisions in near real-time.

Approach: Develop novel inverse reinforcement learning algorithms that support estimating behavioral models [5].

Aim 2

Develop models of emerging market structures that capture user behavior.

Challenges: Developing computationally tractable models is challenging due to the fact that, facilitated by smart devices, market exchanges occur on multiple timescales and are subject to large amounts of uncertainty due to exogenous factors and behavioral aspects of humans.

Approach: Develop algorithms for online incentive design and models of multi-sided markets leveraging dynamical systems tools 3,4].

Aim 3

Testing and validation via computational platform supported by data.

Challenges: Intelligent infrastructure data is often sparse or provides only an indirect measure of phenomena of interest.

Approach: Work with domain experts (municipal and industry partners) to parse, organize and standardize data. Integrate data into simulation platform.

Broader Impact on Society

This project aims to advance our understanding of how urban mobility resources are being consumed and expose areas where municipalities can adjust management strategies to supplement emerging solutions such as platformed-supported resource sharing. We are in a unique position to **translate the results to practice** due to an established relationship with municipal and industry partners.

Education & Outreach

Grad and undergrad students from under-represented groups are heavily involved in the the development of the computational platform and testbed. Students engage with the community via our municipal partners and through this engagement, they are being exposed to the jobs of tomorrow's smart cities. An central goal of the project is to **train the next generation of engineers to incorporate socio-economic considerations into engineering design**.

Products:
[1] E. Mazumdar, L. Ratliff. On the Convergence of Competitive, Multi-Agent Gradient-Based Learning. *arXiv:1804.05464*, 2018.
[2] L. Ratliff, B. Chasnov, Calderone, E. Mazumdar, S. Burden. Convergence Guarantees for Gradient-Based Learning in Continuous Games, under review, 2018

[3] T. Fiez, S. Sekar, L.Zheng, Ratliff. Combinatorial Bandits for Incentivizing Agents with Dynamic Preferences. UAI, 2018.
[4] T. Fiez, S. Sekar, L. Ratliff, Multi-Armed Bandits for Correlated Markovian Environments, under review, 2018
[5] L. Ratliff, E. Mazumdar. Inverse Risk-Sensitive Reinforcement Learning, under review TAC, 2017.