Dissertation summary: ***Learning and Optimization for Mixed Autonomy Systems - A Mobility Context***
By Cathy Wu

How will automated vehicles change mobility?

A running example throughout this thesis concerns the complex integration of automated vehicles into existing mobility systems, which we term *mixed autonomy mobility*. Mobility is a natural and important instance of mixed autonomy systems because it is an existing system in which many human agents interact closely and regularly with many automated agents, such as traffic lights. An upcoming and long anticipated event is the introduction of automated vehicles into the mobility system. This example helps identify and highlight a number of new scaling challenges which push the limits of our optimization frameworks, as exemplified in this thesis.

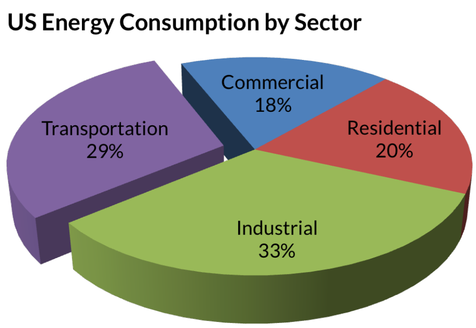

(a) Energy consumption in the United States by sector (April 2018). Source: US Energy Information Administration (2018, Table 2.1).

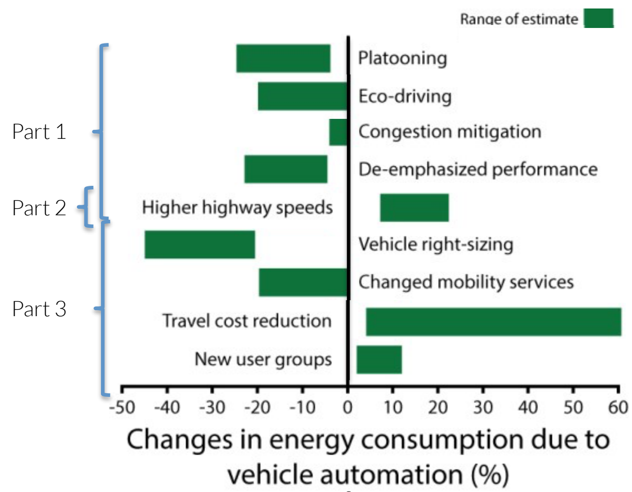

(b) Dissertation summary, by type of automated vehicle impacts addressed. Source: Wadud et al. (2016).

Figure 1: Transportation accounts for nearly a third of energy consumption in the United States. Automating transportation will not necessarily make transportation more efficient. For example, the reduced cost of traveling may induce more people to travel. Everything from future adoption rates of automated vehicles to government regulations will influence the net impact of automated transportation on transit-related energy consumption.

Transportation systems today form a literal physical backbone to civilization. Daily, they touch the lives of 3.9 billion people who reside in urban areas or 54% of the world population (United Nations, 2014) and support 1.32 billion motor vehicles worldwide[1] (WardsAuto, 2017). At the same time, the transportation sector accounting for 29% of energy consumption in the US (see Figure 1a) (US Energy Information Administration, 2018, Table 2.1) and more than 20% of energy-related global GHG emissions world-wide.

Self-driving vehicles are slated to bring about dramatic changes in terms of energy consumption, safety, access and time savings. They can affect energy consumption in a variety of ways, including though vehicle platooning, eco-driving (Gense, 2000), and many others, as summarized in Figure 1b. In particular, the reduction in travel cost could result in a 5-60% increase in energy consumption. Depending on weighted likelihoods of each of these factors, studies determined that with full adoption of self-driving cars, the US mobility system could see anywhere from a minus 40% to a doubling of energy consumption (Wadud et al., 2016). In short, we have a great amount of uncertainty with regards to the impact of automation, even for a single metric (energy, in this case), let alone a holistic measure. How do we even start to reason about all these factors? And then shape the outcome?

This dissertation introduces several new machine learning and optimization techniques that are needed to help guide the evolution of urban mobility in light of the adoption of automated vehicles. We charac-

1 This measure, from 2016, excludes motorbikes.

terize the study of the integration of automated vehicles into existing mobility systems as *mixed autonomy mobility*.

A note on vehicle automation. In the spirit of near-term practicality, this thesis aims to be compatible with a variety of forms and levels of vehicle automation. Vehicles may be automated in a variety of ways, including a dedicated driver such as in a taxi, an instructed or incentivized driver, or an electro-mechanical system such as those designed for a self-driving vehicle. Levels of automation vary as well, from partial automation in the forms of cruise control, braking or parking assistance, routing guidance, or trip management to full automation of the driving or trip related tasks. Although there are many other interesting forms of automation in mobility systems, such as traffic lights, ramp meters, variable speed limits and other electronic roadway signage, mobility usage fees, congestion pricing, and road directionality, the study of their integration into the mobility system are out of scope for the present manuscript. However, the presented approaches may extend readily to many of these forms of automation.

Thesis overview and contributions

This thesis introduces, develops scalable computational tools (including algorithms and systems) for, and addresses concrete challenges in mixed autonomy systems. We begin with a review of the common and powerful optimization frameworks of reinforcement learning, convex optimization, and combinatorial optimization. We briefly review the history, research, challenges, and opportunities of automation in mobility systems. An early version of the historical overview was published as Wu (2016). Then, what follows is the structure and primary contributions of the thesis, as diagrammed in Figures 2 and 1b.

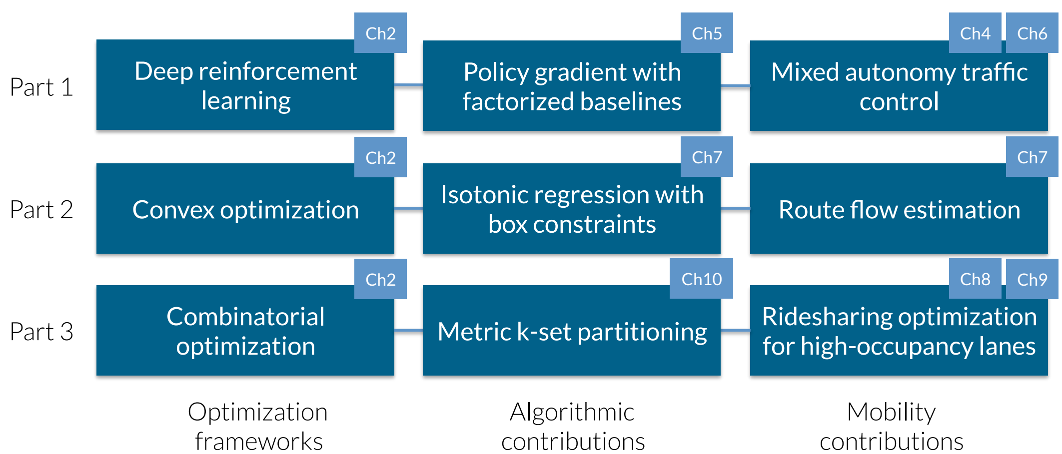

Figure 2: Summary of contributions of the thesis, in terms of optimization frameworks, specific optimization methods, and mobility challenges.

Part 1: Control. The first part of the thesis first demonstrates the existence of mixed autonomy systems in the open world, showing that a mixture of automated and human-driven vehicles may yield vastly different system characteristics, such as average velocity of all vehicles in the system, as compared to a system with only human-driven vehicles. Termed mixed autonomy traffic, this serves as a running example throughout the thesis for concretely demonstrating the problems and challenges of mixed autonomy systems. Simultaneously, we demonstrate the potential of machine learning methods for studying mixed autonomy systems, in contrast to classical techniques based on partial differential equations and manual controller design. In particular, by casting the problem into the framework of model-agnostic reinforcement learning, it is established that a small fraction of automated vehicles has the potential to dramatically improve overall road velocities for all vehicles, and a number of interesting driving behaviors emerge. Many advances are needed to enable mixed autonomy systems at the scale of thousands or even millions of agents, including high-speed distributed simulation and algorithmic developments. To this end, we present generic deep reinforcement learning techniques for scaling up to higher dimensional control problems, such as controlling many vehicles in a mobility system. Additionally, an open-source library is introduced which allows for integrated studies of reinforcement learning and traffic microsimulation, with scalable distributed simulation and cloud deployment. The work in this part was published in Wu et al. (2017b), Wu et al. (2018), and Wu et al. (2017a). The contributions of this part have implications for the environment and public policy concerning the regulation of automated vehicles, as well as scalable reinforcement learning.

Part 2: State estimation. The second part of the thesis explores the sensing challenges and requirements to enable mixed autonomy systems. In mixed autonomy systems, only parts of the system are automated and thus naturally observed, but information about the other parts of the system may be required in order to achieve optimal system performance. Thus being able to measure or estimate relevant quantities is critical to the performance of mixed autonomy systems. In mobility, owing to control theoretic analysis and advances in transportation engineering, vehicle throughput is such a critical quantity to estimate, but has classically been hindered by sparse sensing of the transportation infrastructure. By casting the problem into the framework of convex optimization, we determined that currently available transportation sensing infrastructure, augmented with also currently available aggregate data from cellular networks enables accurate throughput estimation. The structure of cellular network data and the large scale of urban systems motivates the design of a new algorithm for projected gradient descent with a block simplex constraint. The work in this part was published in Wu et al. (2015) and Wu et al. (2019). The contributions of this part have implications for near-term transportation management, as opposed to long-term planning such as land-use planning, and infrastructure design, as well as scalable convex optimization.

Part 3: System design. The third part of the thesis explores the design of the system itself, to mitigate the potential effects of the external process on the system. Mixed autonomy systems may be viewed as being embedded within another dynamical system, one which dictates the progression of the integration, adoption, and use of automation. This process, external to the mixed autonomy system, may induce substantial effects upon the system, both positive and negative. Mobility is embedded in an overall socioeconomic system, and one major anticipated long-term impact of automated vehicles is induced demand, in which more people travel in response to the newly available roadway capacity (enabled in Part 1). This additional demand on the mobility system may compromise the benefits in road velocity and throughput by resulting in elevated energy consumption. We start by empirically studying the dynamics of the overall socioeconomic system and in particular, its couplings with the mobility system. To this end, we build a model of human mobility preferences based on a user study of 300 employees at a major technology corporation. We identify ridesharing as a promising approach and give it treatment as a design paradigm for the mobility system itself, with the goal of mitigating the effects of induced demand by dramatically improving the throughput of the mobility system. We conclude therefore that, with lightly modified existing infrastructure, ridesharing has the potential to dramatically improve (nearly triple) the throughput of the mobility system, and provide combinatorial algorithms to solve the allocation. The structure of the ridesharing problem motivates the adaptation of clustering algorithms from machine learning for set partitioning in the combinatorial optimization framework. The work in this part was published in Wu et al. (2016b) and Wu et al. (2016a). The contributions of this part have implications for mobility system design, urban planning, and public policy, as well as scalable combinatorial optimization.

Discussion and closing remarks. This thesis is only a first step in enabling a science and engineering of mixed autonomy systems, and takes the approach of introducing the broad definition and challenges of mixed autonomy systems, developing scalable computational tools, and addressing concrete challenges, in particular, in mobility. We therefore close with a discussion of the road ahead. We discuss remaining research questions and system requirements for deployment. We present also directions of future research for mixed autonomy systems.

BIBLIOGRAPHY

Gense, N. (2000). *Driving style, fuel consumption and emissions–final report*. Tech. rep. TNO Automotive Technical Report Number 00. OR. VM. 021.1/NG, TNO Automotive (cit. on p. 1).

United Nations (2014). "World urbanization prospects: The 2014 revision, highlights. department of economic and social affairs." In: *Population Division, United Nations*. URL: http://www.un.org/en/development/desa/news/population/world-urbanization-prospects-2014.html (cit. on p. 1).

US Energy Information Administration (2018). *Monthly energy review*. Table 2.1 (cit. on p. 1).

Wadud, Z., D. MacKenzie, and P. Leiby (2016). "Help or hindrance? The travel, energy and carbon impacts of highly automated vehicles." In: *Transportation Research Part A: Policy and Practice* 86, pp. 1–18 (cit. on p. 1).

WardsAuto (2017). *World Vehicle Population Rose 4.6% in 2016*. URL: http://subscribers.wardsintelligence.com/analysis/world-vehicle-population-rose-46-2016 (cit. on p. 1).

Wu, C. (2016). "Traffic Jammin': Making automated transportation a reality." In: *Berkeley Science Review* (cit. on p. 2).

Wu, C., E. Kamar, and E. Horvitz (2016a). "Clustering for set partitioning with a case study in ridesharing." In: *Intelligent Transportation Systems (ITSC), 2016 IEEE 19th International Conference on*. IEEE, pp. 1384–1388 (cit. on p. 3).

Wu, C., A. Kreidieh, K. Parvate, E. Vinitsky, and A. M. Bayen (2017a). "Flow: Architecture and Benchmarking for Reinforcement Learning in Traffic Control." In: *arXiv preprint arXiv:1710.05465*. URL: https://arxiv.org/abs/1710.05465 (cit. on p. 3).

Wu, C., A. Kreidieh, E. Vinitsky, and A. M. Bayen (2017b). "Emergent Behaviors in Mixed-Autonomy Traffic." In: *Proceedings of the 1st Annual Conference on Robot Learning*. Ed. by S. Levine, V. Vanhoucke, and K. Goldberg. Vol. 78. Proceedings of Machine Learning Research. PMLR, pp. 398–407. URL: http://proceedings.mlr.press/v78/wu17a.html (cit. on p. 3).

Wu, C., A. Pozdnukhov, and A. M. Bayen (2019). "Block Simplex Signal Recovery: Methods, Trade-Offs, and an Application to Routing." In: *IEEE Transactions on Intelligent Transportation Systems* (cit. on p. 3).

Wu, C., A. Rajeswaran, Y. Duan, V. Kumar, A. M. Bayen, S. Kakade, I. Mordatch, and P. Abbeel (2018). "Variance Reduction for Policy Gradient with Action-Dependent Factorized Baselines." In: *International Conference on Learning Representations*. URL: https://openreview.net/forum?id=H1tSsb-AW (cit. on p. 3).

Wu, C., K. Shankari, E. Kamar, R. Katz, D. Culler, C. Papadimitriou, E. Horvitz, and A. M. Bayen (2016b). "Optimizing the diamond lane: A more tractable carpool problem and algorithms." In: *Intelligent Transportation Systems (ITSC), 2016 IEEE 19th International Conference on*. IEEE, pp. 1389–1396 (cit. on p. 3).

Wu, C., J. Thai, S. Yadlowsky, A. Pozdnoukhov, and A. Bayen (2015). "Cellpath: Fusion of cellular and traffic sensor data for route flow estimation via convex optimization." In: *Transportation Research Part C: Emerging Technologies* 59. Special Issue on International Symposium on Transportation and Traffic Theory, pp. 111–128. ISSN: 0968-090X. DOI: https://doi.org/10.1016/j.trc.2015.05.004. URL: http://www.sciencedirect.com/science/article/pii/S0968090X15001758 (cit. on p. 3).