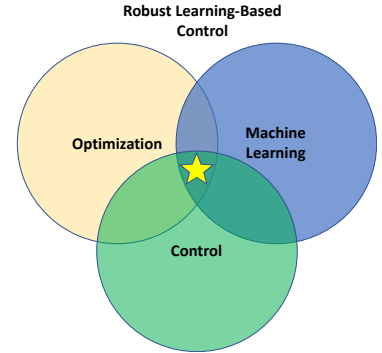


Research Statement

Nikolai Matni

Machine learning techniques – bolstered by successes in video games, sophisticated robotic simulations, and Go – are now being applied to plan and control the behavior of autonomous systems interacting with physical environments. Such systems, which include self-driving vehicles, distributed sensor networks, and agile robots, must interact with complex environments that are ever changing and difficult to model, strongly motivating the use of data-driven decision making and control. However, if machine learning techniques are to be applied in these new settings, it is critical that they be accompanied by guarantees of reliability, robustness, and safety, as failures could be catastrophic. To address these challenges, my research is focused on developing **learning-based control strategies for the design of safe and robust autonomous networked systems**.



In its broadest sense, my work seeks to understand how well a system (or the environment it operates in) needs to be learned in order to effectively and safely control it. This involves understanding and integrating the following three design steps: (i) system/environmental identification and modeling, (ii) policy and control design, and (iii) adaptation and refinement. My approach is to recognize that machine learning algorithms produce estimates or predictions that are inherently uncertain, and therefore this uncertainty must be explicitly quantified (e.g., using non-asymptotic guarantees of contemporary high-dimensional statistics) and accounted for (e.g., using robust control/optimization) when designing safety critical systems. I describe my approach and research vision for **robust learning-based control** in Section 1 below, with the ultimate goal of this line of research being to provide a rigorous and contemporary view on **reinforcement learning as applied to continuous control problems**.

Further, in addition to the challenges described above, many systems of interest today are large-scale, distributed, and networked (e.g., the smart-grid, software defined networks, automated transportation systems). Another major thrust of my work is extending methods from learning and control to this large-scale networked setting, wherein information exchange constraints between agents, as well as scalability of methods, play an important role. I describe my approach and research vision for **learning and control in large-scale networked systems** in Section 2 below, which builds on previous work that already has had significant impact in the controls community (with [4] winning the **IEEE CDC 2013 Best Student Paper Prize**, and [12] winning the **IEEE ACC 2017 Best Student Paper Prize**).

There is a broad and exciting research frontier at the intersection of learning and control, and my research background makes me exceptionally well suited to make significant impact in this interdisciplinary and timely area. Going forward, I aim to build a research program that both expands our theoretical understanding of the interplay between **machine learning and control**, and that uses these theoretical insights to design and implement **real-world algorithms that are safe, sample-efficient, and robust**. Specifically, I plan to develop methods that are applicable to richer system models, learning algorithms, and control approaches (e.g., nonlinear, model-predictive-control, vision based control, etc.), and to extend these approaches to be applicable to large-scale networked systems of autonomous agents (e.g., moving from a single self-driving car to a fleet of autonomous vehicles). I also aim to work with practitioners in domains such as self-driving vehicles, computer networking, and robotics to validate our methods experimentally. Concrete steps towards achieving these goals are outlined in more detail below.

1 Robust Learning-Based Control

Given the dramatic successes in machine learning and reinforcement learning (RL) over the past half decade, there has been a resurgence of interest in applying these techniques to continuous control problems in robotics, self-driving cars, and unmanned aerial vehicles. Though such applications appear to be straightforward generalizations of standard RL, few fundamental baselines have been established prescribing how well one must know a system in order to control it. Further, learning algorithms produce estimates and predictions that are inherently uncertain – when using such algorithms in safety and performance critical control loops, it becomes imperative to explicitly quantify and account for this uncertainty in a sample-efficient manner.

To that end, in [1], we describe a contemporary view that merges techniques from statistical learning theory and robust control to derive baselines for an optimal control problem – the linear quadratic regulator (LQR) – wherein the system model to be controlled is unknown. We propose a simple 3-step strategy: (i) use least-squares to estimate a model based on experiments, (ii) use bootstrapping techniques to build probabilistic guarantees about the distance between the estimated and true model, and (iii) solve an optimal control problem that robustly optimizes the performance of the system while guaranteeing stable and robust execution for any realization of the estimated model uncertainty. Leveraging our recently developed System Level Synthesis (SLS) framework, we provide the first known baselines delineating the possible control performance achievable given a fixed amount of data collected about the system, bringing rigor to reinforcement learning methods in continuous control. We show that (i) robustness is not only necessary from a practical perspective, but is also essential to making the analysis of our approach tractable, and (ii) that by explicitly identifying a model, we are able to rapidly (i.e., in a sample-efficient manner) synthesize stabilizing and high-performing controllers.

In follow up work, we extend this approach to a model based reinforcement learning algorithm with provable sublinear regret bounds (the first such algorithm that is polytime computable, and that provides guarantees of robust stability and performance throughout) [2], and to allow for safety constraints to be satisfied during the learning process [3]. Our emphasis on verifiable robustness and stability has made this work perfectly suited for **DARPA’s Assured Autonomy** program, which has been funding this research for the past year.

Future Work

These results represent a modern take on robust, adaptive, and reinforcement learning based control, and to the best of our knowledge, establish the first end-to-end baselines for learning in optimal control problems that do not require restrictive or unrealistic assumptions. A key enabling technical tool in deriving these results is System Level Synthesis, which for the first time allows for clean and interpretable bounds on the performance of a robust controller to be obtained in terms of the size of the uncertainties of the system estimates. As we describe next, although these results are exciting, they represent the beginning of a much broader line of research.

Richer Models and Modalities: An important next step is to consider richer (and more realistic) ways in which learning enabled components are integrated into control and optimization loops. Indeed, in many applications of interest, the dynamics of the system are well understood, and it is rather the sensing modalities that require learning-enabled components. As a concrete example, consider a self-driving car: under normal operating conditions, simple and well understood models are sufficient for control. However, translating sensor data (e.g., pixels from a camera or point clouds from a lidar) to useful measurements for control (e.g., cross-track error from the center of the lane) requires a complex nonlinear map that must be learned. As the output of this uncertain map will be fed directly into a safety critical control loop, it is essential to be able to provide guarantees about the accuracy of its output, as well as to understand under what conditions those guarantees do and do not hold.

A natural question that then arises from the previous discussion is how can we learn such a map (e.g., from pixel space to position) in such a way that is both data-efficient and amenable to control. We conjecture that key to this procedure will be exploiting and enforcing the fact that the high-dimensional data (e.g., camera images) is generated by a process constrained to a lower-dimensional manifold defined by the dynamics of the underlying system (e.g., a subspace in the case of a linear system). This also raises the broader question of how should learning algorithms be tuned to the specific control task at hand, which we believe will be an important and fruitful direction for future work.

Beyond Linear Time Invariant Learning and Control: In [1], we consider the simplest possible control paradigm: an unconstrained linear-time-invariant system being controlled by an unconstrained linear-time-invariant controller. Ultimately, the systems that we care about, as well as the environments that they interact with, are nonlinear and constrained. Future work will look to bridge this gap in several ways. First, the results derived above, and in particular the SLS framework, can be naturally incorporated into a model predictive control (MPC) framework, allowing us to naturally incorporate constraints on the state and control input, allowing our approach to accommodate to a much richer notion of safety (see [3] for a flavor of such results). Further, when integrated with sequential linearization, these methods have proven to be effective heuristics in controlling nonlinear systems. Finally, recent lifting techniques (e.g., those based on Koopman theory) suggest that certain

nonlinear systems can indeed be analyzed using linear systems tools – this represents an interesting direction for future work.

Self-Driving Car Simulation and Experimental Validation: I am currently exploring applications of our safe learning and control algorithms using the Unity3d based self-driving car simulator that Udacity has released as open-source code (see Figure 1). This simulator is ideally suited to rapid prototyping and validation of learning-based control algorithms as it presents many of the challenges inherent to real-world hardware implementations (latency and multithreading issues, nonlinear and unspecified vehicle and environment dynamics, etc.) without the difficulties and overhead related to developing and maintaining a hardware testbed. It also allows for easily switching between different sensing modalities (ground truth, noisy sensors, lidar, vision based navigation), meaning that increasingly sophisticated and rich models/learning algorithms can be validated using this simulation testbed. I am also beginning a collaboration with Francesco Borelli’s group, and ultimately aim to move algorithms validated in this simulation environment to their 1/10 scale self-driving car platform over the coming months.



Figure 1: Unity3d based self-driving car simulator released by Udacity.

2 Learning and Control in Large-Scale Networked Systems

As large-scale networked systems become increasingly dynamic and decentralized, feedback control systems will be essential in guaranteeing their reliable and robust behavior – indeed robust and predictable behavior at the component level of a system allows higher system-level tasks, such as resource allocation, learning and exploration, and optimizing social/economic welfare, to be made simpler and more efficient.

Therefore, in order to move from the design of a single autonomous system to a large network of interconnected and interacting autonomous agents, several technical challenges must be overcome. Even without learning enabled components, the problem of computing distributed optimal controllers at scale is a difficult one. Indeed, even seemingly simple decentralized optimal control problems are intractable (e.g., Witsenhausen’s counterexample), making it a challenge to extend foundational results of the field (e.g., the Youla parameterization and DGKF state-space solutions) to the distributed setting. A recent breakthrough was the identification of a broad class of tractable (convex) distributed optimal control problems by Rotkowitz and Lall in 2006. These systems satisfy a quadratic invariance (QI) property, and allow classical (Youla) based synthesis methods to be used to compute structured controllers. Although QI was an important step forward, the resulting controllers were not scalable to synthesize or implement. This lack of scalability has prevented the adoption of these methods in practice, limiting their scope to systems with at most a few hundred states. We address this issue by using our recently developed System Level Synthesis framework [14, 13], in which we present an alternative to the Youla parameterization that is naturally suited to distributed (structured) controller synthesis. By rethinking and generalizing this foundational pillar of modern control theory, we are able to identify the broadest known class of convex constrained linear optimal control problems, of which QI distributed optimal control problems are a special instance.

Unlike its predecessors, our approach allows for the synthesis of localized system responses and corresponding controllers, in which local control policies depend only on local subsets of the global system model and state [10, 11, 9, 8]. The importance of this contribution was recognized by the controls community, with the paper [12] being awarded the **Best Student Paper Award at the 2017 IEEE American Control Conference**. The notion of locality, which could not be captured by existing theory, allows us to scale structured optimal controller synthesis methods to systems of arbitrary size (assuming sufficient parallel computing power): an example illustrating the benefits of our approach on controller synthesis computation time are shown in Figure 2. In addition to the aforementioned locality constraints, the system designer can simultaneously impose multiple performance objectives, structural constraints on the controller, robustness to quantization in the communication and internal computation of the controller, robustness to modeling error, and limits on the actuation, sensing and communication complexity of the resulting controller [11, 6, 4, 5, 7] (with [4] winning the **CDC 2013 Best Student Paper Prize**). Thus, by generalizing the foundational results of QI and the Youla parameterization,

System Level Synthesis provides, for the first time, a unified framework for customized controller synthesis that is applicable to large-scale distributed systems.

Future Work

Robustness and Safety for Networked Systems: I believe it is essential to develop a theory of robustness, akin to that from the 90s centered around H_∞ optimal control and μ -synthesis, that is applicable to large-scale distributed systems. In addition to these methods being poised to have significant impact on emerging areas such as smart-grids and intelligent transportation systems, they are key to the principled integration of learning enabled components. The key in this setting will be identifying appropriate structural constraints on model uncertainty (to capture the spatially distributed nature of the system) to impose in order to ensure robust, but not overly conservative, behavior. SLS is ideally suited for tackling this problem as it provides a transparent mapping between system uncertainty and its effect on system performance. Similarly, SLS provides a straightforward avenue for extending MPC (and its ability to include state and input constraints) to the distributed large-scale setting – this is an avenue that we are actively pursuing and expect to have preliminary results for shortly.

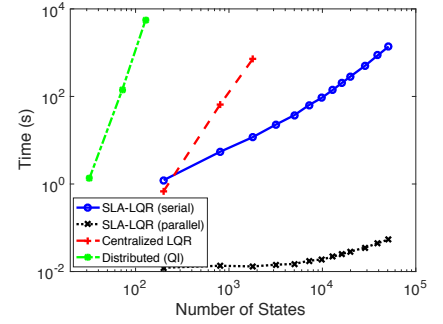


Figure 2: Computational time needed to synthesize centralized, distributed (QI), and localized (SLS) controllers as a function of system state size.

Distributed Learning and Control: Consider a fleet of similar (but not quite identical) self-driving vehicles navigating through a city and collecting camera, lidar, and other telemetry data. How can the data collected by the *entire* fleet be leveraged to improve the safety and performance of each individual vehicle, despite differences in individual vehicle dynamics, operating conditions, etc. Answering this question will involve drawing on tools from distributed learning and robust control in order to develop learning and control algorithms that leverage the collective data of the fleet to improve the behavior of an individual vehicle. As a concrete first step in this direction, I am currently collaboration with Somayeh Sojoudi’s group at UC Berkeley to integrate our localized optimal control paradigm with structured inference (namely LASSO) to extend the results of [1] to the large-scale distributed setting.

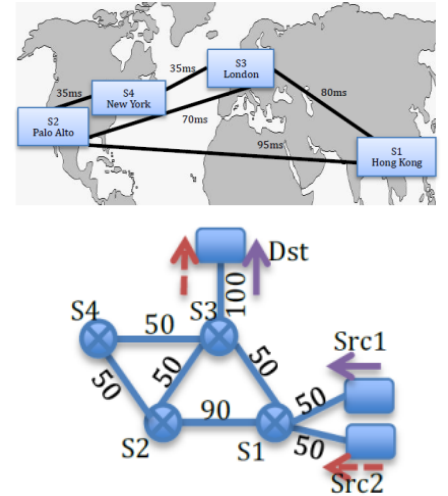


Figure 3: Production Wide Area Network (WAN) used to validate our High Frequency Traffic Control algorithm.

Learning and Control in Computer Networks: An exciting development in the world of computer networking is the introduction of Software Defined Networking (SDN), which provides an abstraction between the traditional forwarding (data) plane and the control plane of a network. This abstraction allows for an explicit separation between data forwarding and data control, and provides an interface through which network applications can programmatically control the network. This in turn allows for diverse, distributed application software to be run using diverse, distributed hardware in a seamless way. With this increased design freedom comes the opportunity to design more sophisticated, dynamic, and reactive learning-based network control algorithms.

As a concrete example of progress made in this area, I have developed and implemented (in collaboration with Ao (Kevin) Tang and his group at Cornell University) a novel approach to in-network congestion management, which we call High Frequency Traffic Control (HFTraC) [15], that operates at the network layer at the timescale of round-trip time. HFTraC’s objective is to minimize a weighted sum of queue length and flow rate fluctuation by utilizing available buffer space in routers network-wide, allowing for the principled exploration and optimization of the tradeoff between packet loss % and queueing length. Another key component of this work is quantifying how the achievable performance of HFTraC is determined by the network architecture used to implement it (e.g., whether router service rate decisions are computed in a decentralized, distributed or centralized manner). Finally,

in order to validate the effectiveness of HFTraC, we implement and evaluate its performance on a custom designed experimental testbed, a Mininet emulator, and a production wide area network (WAN) (see Figure 3).

Leveraging my expertise across learning-based control, distributed optimal control, and SDN, I am currently collaborating with Ion Stoica at UC Berkeley to develop reinforcement learning based congestion control algorithms.

References

- [1] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. On the sample complexity of the linear quadratic regulator. *Journal of Foundations of Computational Mathematics, Accepted with minor revisions*, 2018.
- [2] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *32nd Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- [3] S. Dean, S. Tu, N. Matni, and B. Recht. Safely learning to control the constrained linear quadratic regulator. *arXiv preprint arXiv:1809.10121*, 2018.
- [4] N. Matni. Communication delay co-design in \mathcal{H}_2 decentralized control using atomic norm minimization. In *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*, pages 6522–6529, Dec 2013.
- [5] N. Matni. Communication delay co-design in \mathcal{H}_2 -distributed control using atomic norm minimization. *IEEE Transactions on Control of Network Systems*, 4(2):267–278, June 2017.
- [6] N. Matni and V. Chandrasekaran. Regularization for design. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 1111–1118, Dec 2014.
- [7] N. Matni and V. Chandrasekaran. Regularization for design. *IEEE Transactions on Automatic Control*, 61(12):3991–4006, Dec 2016.
- [8] N. Matni, Y. Wang, and J. Anderson. Scalable system level synthesis for virtually localizable systems. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 3473–3480, Dec 2017.
- [9] Y. S. Wang and N. Matni. Localized LQG optimal control for large-scale systems. In *2016 American Control Conference (ACC)*, pages 1954–1961, July 2016.
- [10] Y.-S. Wang, N. Matni, and J. C. Doyle. Localized LQR optimal control. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 1661–1668, Dec 2014.
- [11] Y. S. Wang, N. Matni, and J. C. Doyle. Localized LQR control with actuator regularization. In *2016 American Control Conference (ACC)*, pages 5205–5212, July 2016.
- [12] Y. S. Wang, N. Matni, and J. C. Doyle. System level parameterizations, constraints and synthesis. In *2017 American Control Conference (ACC)*, pages 1308–1315, May 2017.
- [13] Y.-S. Wang, N. Matni, and J. C. Doyle. Separable and localized system level synthesis for large-scale systems. *IEEE Transactions on Automatic Control*, 2018.
- [14] Y.-S. Wang, N. Matni, and J. C. Doyle. A system level approach to controller synthesis. *IEEE Transactions on Automatic Control, Accepted*, 2018.
- [15] N. Wu, Y. Bi, N. Michael, A. Tang, J. C. Doyle, and N. Matni. A control-theoretic approach to in-network congestion management. *IEEE/ACM Transactions on Networking*, pages 1–14, 2018.