ELECTRICAL & COMPUTER ENGINEERING
SUMMER RESEARCH INTERNSHIP PROGRAM

FACULTY MENTOR
Yang Zheng

PROJECT TITLE
Robust Predictive Control for Autonomous Vehicles in Mixed Traffic Systems

PROJECT DESCRIPTION
The emergence of connected and autonomous vehicles (CAVs) is expected to improve traffic efficiency and fuel economy. In practice, the mixed traffic system, where human-driven vehicles (HDVs) coexist with CAVs, is expected to be the predominant pattern in the near future. Thus, the control of CAVs needs to consider the behaviors of HDVs while designing their driving strategies. Although predictive control of CAVs has shown great performance from both small-scale real-world experiments (e.g., https://youtu.be/2mBjYZTeaTc) and large-scale numerical simulations (e.g., https://soc-ucsd.github.io/mixed-traffic/), its simplistic assumption of human-driven behaviors may pose safety concerns. Recent work has started to predict a series of future trajectories of HDVs with more accurate models and consider the worst-case scenario to improve safety.

In this project, the student will begin by reading some existing analyses and control methods of autonomous vehicles in mixed traffic, e.g., [1]-[3], and the references therein. Then, the student will need to explore new estimation methods (e.g., Gaussian Process Regression, machine learning-based methods) for predicting HDVs' potential future trajectory set and incorporate it with robust prediction control. Another major focus will be the implementation of the analysis and control methods in Python or Julia. A range of numerical demonstrations will be carried out in this project too.

Reference:
[1] Zheng, Y., Wang, J., & Li, K. (2020). Smoothing traffic flow via control of autonomous vehicles. IEEE Internet of Things Journal, 7(5), 3882-3896.
[2] Wang, J., Zheng, Y., Li, K., & Xu, Q. (2022). DeeP-LCC: Data-enabled predictive leading cruise control in mixed traffic flow. arXiv preprint arXiv:2203.10639.
[3] Shang, X., Wang, J., & Zheng, Y. (2023). Smoothing Mixed Traffic with Robust Data-driven Predictive Control for Connected and Autonomous Vehicles. arXiv preprint arXiv:2310.00509.

This project will be in person.

INTERNS NEEDED
- 2

PREREQUISITES
- A strong background in linear algebra and optimization and good programming skills (Python, Matlab, Julia, etc.) are required.
- Basic knowledge of connected vehicles is preferred.

FACULTY MENTOR
Yang Zheng

PROJECT TITLE
Structured Low-Rank Approximation with Machine Learning

PROJECT DESCRIPTION
In system identification and signal processing, it is often assumed that the underlying systems or noise-free signals have low-order properties and that the collected data satisfies some specific patterns. The objective is then to estimate the particular low-order structure from the noisy data. This type of problem is called a structured low-rank approximation (SLRA) problem. Although the well-known singular value decomposition can provide the optimal low-rank approximation to the noisy data matrix, it does not guarantee any optimality of estimating the original noise-free matrix. Another difficulty is to incorporate the structural constraints into optimization, and there is no closed-form solution or convex formulation in general. This project aims to explore whether machine learning techniques can contribute to the SLRA problem.

In this project, the student will begin by reading some existing algorithms and convergence analysis of SLRA, e.g. [1]-[3], and the references therein. Then, the student will need to explore new approximation methods for structured low-order matrices, especially for estimating the shrinkage law. Another major focus will be adjusting existing algorithms for different matrix structures and comparing their denoising performance.

Reference:
[1] Yin, M., & Smith, R. S. (2021). On low-rank hankel matrix denoising. IFAC-PapersOnLine, 54(7), 198-203.
[2] Wang, C., Zhu, Z., Gu, H., Wu, X., & Liu, S. (2018). Hankel low-rank approximation for seismic noise attenuation. IEEE Transactions on Geoscience and Remote Sensing, 57(1), 561-573.
[3] Alsalti, M., Barkey, M., Lopez, V. G., & Müller, M. A. (2023). Sample and computationally efficient data-driven predictive control. arXiv preprint arXiv:2309.11238.

This project will be in person.

INTERNS NEEDED
- 2

PREREQUISITES
- A strong background in linear algebra and optimization is required.
- Basic knowledge of the linear system is preferred.

FACULTY MENTOR
Yang Zheng

PROJECT TITLE
A Unified Analysis for First-Order Algorithms in Convex Optimization

PROJECT DESCRIPTION
First-order algorithms are iterative methods that only use first-order information such as (sub)gradients and first-order derivatives in optimization problems. This class of algorithms has low computational complexity per iteration, scaling very favorably to large-scale instances and model-free setups, which has applications in a wide range of areas. Indeed, first-order algorithms are the main workhorse for modern machine learning and reinforcement learning, which have achieved many successful successes (such as AlphaGo and ChatGPT). Under favorable conditions, many first-order algorithms are known to converge at a linear rate for convex problems, even without strong convexity. These conditions include quadratic growth, subdifferential error bound, proximal mapping error, and others.

However, it is a crucial question to choose the correct condition, and not all the proof techniques are intuitive. Some are opaque, at least at first glance. Therefore, this project is interested in developing a unified and intuitive framework that covers many first-order algorithm analyses in both smooth and nonsmooth optimization. Students will start by reading existing results about different analysis approaches, i.e., [1]-[4], and understanding the connections and differences among them. The goal is to develop a unified framework for algorithm analyses and develop new algorithms with faster convergence guarantees and lower computational complexity.

Reference:
[1] Atenas, F., Sagastizábal, C., Silva, P. J., & Solodov, M. (2023). A unified analysis of descent sequences in weakly convex optimization, including convergence rates for bundle methods. SIAM Journal on Optimization, 33(1), 89-115.
[2] Drusvyatskiy, D., & Lewis, A. S. (2018). Error bounds, quadratic growth, and linear convergence of proximal methods. Mathematics of Operations Research, 43(3), 919-948.
[3] Luo, Z. Q., & Tseng, P. (1993). Error bounds and convergence analysis of feasible descent methods: a general approach. Annals of Operations Research, 46(1), 157-178.
[4] Necoara, I., Nesterov, Y., & Glineur, F. (2019). Linear convergence of first-order methods for non-strongly convex optimization. Mathematical Programming, 175, 69-107.

This project will be in person.

INTERNS NEEDED
- 1

PREREQUISITES
- A strong background in linear algebra is required.
- Basic knowledge of coding (Matlab or Python) is preferred.

FACULTY MENTOR
Yang Zheng

PROJECT TITLE
Continuous Control from Optimization and Reinforcement Learning Perspectives

PROJECT DESCRIPTION
Motivated by the exceptional achievements of reinforcement learning in domains such as mastering video games and Go, there has been a renewed interest in exploring the theoretical underpinnings of policy optimization when applied to the control of dynamical systems [1]. Significant progress has been made in various benchmark control problems, including the classical Kalman filter, linear quadratic stochastic control [2][3], and robust control [4]. These control problems pose challenges as they are nonconvex in the policy space. Consequently, analyzing the optimization landscape or further ensuring the global optimality of policy search methods is highly nontrivial.

The focus of this project lies in the development of model-free learning algorithms for continuous control problems with provable performance guarantees. This necessitates insights from mathematical optimization, statistics, control theory, and reinforcement learning.

Reference:
[1] Hu, Bin, et al. "Toward a Theoretical Foundation of Policy Optimization for Learning Control Policies." Annual Review of Control, Robotics, and Autonomous Systems 6 (2023): 123-158.
[2] Fazel, Maryam, et al. "Global convergence of policy gradient methods for the linear quadratic regulator." International conference on machine learning. PMLR, 2018.
[3] Zheng, Yang, Yujie Tang, and Na Li. "Analysis of the optimization landscape of linear quadratic Gaussian (LQG) control." arXiv preprint arXiv:2102.04393 (2021).
[4] Tang, Yujie, and Yang Zheng. "On the Global Optimality of Direct Policy Search for Nonsmooth H infty Output-Feedback Control." arXiv preprint arXiv:2304.00753 (2023).

This project will be in person.

INTERNS NEEDED
- 2

PREREQUISITES
- A strong background in linear algebra and optimization is essential.
- A keen interest or experience in the theoretical/practical aspect of control/reinforcement learning would be advantageous.
- Basic knowledge of coding (Matlab or Python) is preferred.

FACULTY MENTOR
Yang Zheng

PROJECT TITLE
Semidefinite Relaxation for Machine Learning Problems

PROJECT DESCRIPTION
Neural networks (NNs) have shown promising results in many applications, including image recognition, natural language processing, and robotic control. Some recent work has also incorporated Deep NNs (DNNs) into safety-critical systems, like autonomous driving. Although DNNs obtain unprecedented empirical success, their opacity poses significant challenges, especially in safety-critical systems. Therefore, it is desirable to provide rigorous analysis and guarantees about the behaviors of DNNs, either as a standalone component or as an integrated component in a feedback loop.

One interesting problem is the robustness verification of NNs, which aims to find adversarial examples that reveal ways NNs can fail or certify no adversarial examples exist. The input/output behavior of NNs is nonlinear and non-convex, making it hard to solve verification problems exactly. A variety of convex relaxation techniques, such as interval bound propagation, linear programs (LP) [1]-[3], and semidefinite programs [4], have been proposed to derive an approximated solution. It is shown empirically that semidefinite relaxation can offer much tighter bounds than LP relaxations [4], but semidefinite relaxation, in general, suffers from the scalability issue. Currently, the techniques in [4] can only deal with small NNs.

This project aims to develop efficient semidefinite relaxation for NN verification. The starting point is to exploit the inherent neural network structure (e.g., cascading structure) to develop layer-wise relaxations; a related topic is chordal decomposition (see [5][6]). Several possible objectives can be targeted, including:

- The development of efficient methods for verifying NNs via semidefinite relaxation (layer-wise decomposition, and tightening by LP cuts etc.)
- The development of strategies to balance relaxation quality and solution efficiency for semidefinite relaxation
- The implementation of a verification toolkit for the verification of NNs
- Numerical comparison with the existing toolkits.

References

[1] Salman, H., Yang, G., Zhang, H., Hsieh, C. J., & Zhang, P. (2019). A convex relaxation barrier to tight robustness verification of neural networks. In Advances in Neural Information Processing Systems (pp. 9835-9846).
[2] Singh, G., Ganvir, R., Püschel, M., & Vechev, M. (2019). Beyond the single neuron convex barrier for neural network certification. In Advances in Neural Information Processing Systems (pp. 15098-15109).
[3] Tjandraatmadja, C., Anderson, R., Huchette, J., Ma, W., Patel, K., & Vielma, J. P. (2020). The convex relaxation barrier, revisited: Tightened single-neuron relaxations for neural network verification. arXiv preprint arXiv:2006.14076.
[4] A. Raghunathan, J. Steinhardt, and P. S. Liang, "Semidefinite relaxations for certifying robustness to adversarial examples," inAdvances in Neural Information Processing Systems, pp. 10877–10887,2018.
[5] Vandenberghe, Lieven, and Martin S. Andersen. "Chordal graphs and semidefinite optimization." Foundations and Trends in Optimization 1.4 (2015): 241-433.
[6] Zheng, Y., Fantuzzi, G., Papachristodoulou, A., Goulart, P., & Wynn, A. (2020). Chordal decomposition in operator-splitting methods for sparse semidefinite programs. Mathematical Programming, 180(1), 489-532.

This project will be in person.

INTERNS NEEDED

➢ 2

PREREQUISITES

➢ A strong background in linear algebra and optimization and good programming skills (Python, Matlab, Julia, etc.) are required.

FACULTY MENTOR
Yang Zheng

PROJECT TITLE
Analysis of Performance Estimation Problems for Nonsmooth and Nonconvex Optimization

PROJECT DESCRIPTION
This project is interested in how computers can assist in designing (and verifying) convergence proofs for first-order optimization methods. In particular, we consider the performance estimation problems (PEPs), introduced by Yoel Drori and Marc Teboulle [1]. This framework has been successfully used for the analysis and design of many first-order algorithms [2]-[4]. In this project, we will investigate the analysis of PEPs for nonsmooth and nonconvex optimization, particularly for the class of weakly convex functions. The first step will be developing new interpolation conditions for weakly convex functions under appropriate setups, and then we will apply the interpolation conditions to derive QCQP and SDP formulations for the analysis and design of first-order algorithms.

References:

[1] Yoel Drori, Marc Teboulle. Performance of first-order methods for smooth convex minimization: a novel approach. Mathematical Programming, 145(1-2), 451-482, 2014.
[2] Adrien Taylor, Julien Hendrickx, François Glineur. Exact worst-case convergence rates of the proximal gradient method for composite convex minimization. Journal of Optimization Theory and Applications, vol. 178, no 2, p. 455-476, 2018.
[3] Adrien Taylor, Julien Hendrickx, François Glineur. Exact worst-case performance of first-order methods for composite convex optimization. SIAM Journal on Optimization, vol. 27, no 3, p. 1283-1313, 2017.
[4] Gupta, Shuvomoy Das, Bart PG Van Parys, and Ernest K. Ryu. "Branch-and-bound performance estimation programming: A unified methodology for constructing optimal optimization methods." arXiv preprint arXiv:2203.07305 (2022).

This project will be in person.

INTERNS NEEDED
- 1

PREREQUISITES
- A strong background in linear algebra and optimization and good programming skills (Python, Matlab, Julia, etc.) are required.