Lagrangian Methods for Learning-Based Constrained Optimal Control

Abstract:

This thesis explores the development of learning-based algorithms that approximately solve the optimal control problems for uncertain systems under user-defined safety constraints. As autonomous technologies, ranging from self-driving vehicles to robot-assisted surgeries become increasingly integrated into our society, the need for learning-based controllers under epistemic uncertainties becomes increasingly apparent. While traditional optimal control methods frequently depend on the knowledge of the underlying system dynamics which is often not possible in practice, data-driven approaches on the other hand may have safety violations during the training phase. Additionally, ensuring safety often yields sub-optimal performance. This thesis bridges these gaps by leveraging Lagrangian dual formulations and the online estimation of state-dependent Lagrange multipliers to ensure that safety is treated as a fundamental component of the control design. The first major contribution of this work is a safe Q-learning framework for continuous-time linear time-invariant systems with partially unknown dynamics. By integrating reciprocal control barrier functions (RCBFs) into an actor-critic architecture, the proposed method formulates a safety-augmented cost functional that ensures the forward invariance of a user-defined safe set. This approach allows for the online estimation of a constrained optimal control law while providing mathematical proof that the system states never leave the predefined safety boundaries during the learning process.

Subsequently, we extend these concepts to more complex, uncertain nonlinear systems through the introduction of the Actor-Critic-Identifier-Lagrangian (ACIL) algorithm. Recognizing that static Lagrange multipliers can lead to overly conservative system behavior and sub-optimal performance, this thesis proposes a dynamic adaptation law for multipliers using a smooth approximation of the ReLU function. This online Lagrange multiplier estimation approach enables the controller to approximately solve the constrained optimal control problem in real-time for systems with unknown drift dynamics. Through a rigorous Lyapunov-based analysis, the algorithm is shown to maintain closed-loop stability and constraint satisfaction, outperforming existing online RL approaches in terms of the cost accrued.

Finally, we generalize the notion of safety by moving away from traditional state-space representations to propose a data-driven framework for constrained linear-quadratic control based on behavioral systems theory. By treating system dynamics and constraints as fundamental trajectory-level objects, the work utilizes the theory of quadratic differential forms and latent-variable image representations to solve optimization problems directly from measured data. This approach characterizes the feasibility and optimality of behavioral constraints and demonstrates how a finite set of persistently exciting data trajectories can be used to synthesize optimal policies without requiring the explicit knowledge of the system dynamics. Collectively, these contributions demonstrate that the concurrent pursuit of optimality and safety through adaptive Lagrangian methods provides the theoretical foundations necessary for the reliable deployment of learning-based controllers in safety-critical environments.