Scalable Polynomial Optimization for Trustworthy Autonomy

Introduction: A self-driving car driving in any environment needs to have an accurate geometric model of the road in order to safely navigate in real time. The key challenges for this autonomous navigation task are to transform sensor data (in real time) into state estimates of the environment that update as the car maps a path and drives through the environment. Solving these fundamental problems of robotic planning, perception, and control can be naturally formulated as mathematical optimization problems, which often take the form of (or can be arbitrarily well approximated by) *polynomial optimization problems* (POPs): optimization problems of the form

$$p^{*} = \arg \min_{x \in \mathbb{R}^{n}} f(x),\ s.t.\ g_{i}(x) \leq 0,\ h_{j}(x) = 0,\ \forall\, i \in \mathcal{I},\ j \in \mathcal{J}, (1)$$

in which the functions $f, g_i, h_j \in \mathbb{R}[x]$ are polynomials. General (nonconvex) POPs are *at least NP-hard* to solve (including the problem classes that commonly arise in machine intelligence applications), making this class of problems impractical to solve *in general*.

As a result, current state of the art methods in robotics typically resort to *heuristic local search* due to the need for real-time operation. While this enables fast computation, it comes at the expense of reliability, as there are no guarantees on the quality of the solutions that local search will return. To highlight this issue, consider the (real-world) example trajectory estimation problem like the one shown in *Figure 1*. Image *1a* shows a garage in which a car is attempting to navigate. Images *1b, 1c,* and *1d* are all examples of actual state estimates (stationary points) recovered from local search applied to an instance of Problem 1; however, estimates *1c* and *1d* are so egregiously wrong that they cannot support reliable navigation. This brittleness and unreliability is a common characteristic of perception, planning, and control methods that are based on heuristic local search.

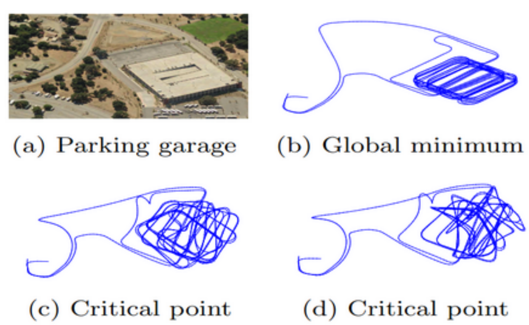

Figure 1: State estimate (b) is the only safe model of the parking garage.

As an alternative to local search, recent work has shown that the *specific instances* of (1) encountered in real-world robotics applications can often be solved much more efficiently than its worst-case NP-hard computational complexity would suggest [2]. In particular, *semidefinite relaxation* strategies have proven to be shockingly effective in many specific machine intelligence problems, enabling the efficient recovery of nearly (and often exactly) optimal solutions in practice [3]. Semidefinite relaxation methods allow us to approximate the original optimization problem with a convex surrogate (in the form of a semidefinite program (SDP)) that is easier to solve, and whose solution can be taken as an approximate solution for the original problem. Remarkably, recent work has shown that certain SDP relaxations will *provably* produce global minimizers for certain classes of generally-intractable estimation problems (such as robotic mapping) under realistic conditions [3].

While these initial results are promising, at present these SDP relaxation approaches are difficult to use because current semidefinite optimization methods struggle to scale to the large-scale SDP relaxations that arise in machine intelligence applications (which are often on the order of hundreds of thousands of dimensions). Prior work has proposed specialized, *problem-specific* SDP optimizers that are tailored to individual applications; at present there is no *general* algorithmic framework for solving the large-scale, but sparse, POP relaxations that appear in real-world machine intelligence tasks.

Research Plan: ***The goal of this project is to develop efficient algorithms (and associated software libraries) that ingest a description of a large-scale but sparse POP, and construct and solve its associated SDP relaxation.*** We propose to pursue two independent but complementary research arms. The first arm focuses on developing faster optimization techniques for the large-scale semidefinite relaxations arising in machine intelligence applications, while the second arm investigates developing techniques to create a simpler SDP relaxation of POPs to solve.

Arm One: Scalable, structure-exploiting SDP optimization: Currently, general-purpose interior-point SDP solvers (such as SeDuMi [1]) are limited to matrices on the order of a few thousand dimensions, due to the computational cost of forming and solving the associated Newton systems for dense matrices [4]. However, many SDPS exhibit useful structure (such as sparsity and/or low-rank solutions) that can be algorithmically exploited. The goal of this arm is to identify algorithmically useful structures that are commonly present in the SDP relaxations arising in machine intelligence tasks, and develop novel SDP optimization algorithms that exploit these structures to efficiently scale to real-world use cases (hundreds of thousands of dimensions).
Arm Two: More tractable POP Relaxations for Machine Intelligence: Standard SDP relaxations of POPs often impose many constraints to ensure tightness between the original problem and the relaxation. However, recent work has shown that many of these constraints can be discarded at *negligible* cost in solution quality [2]. In this arm, we will explore the design of more parsimonious SDP relaxations for machine intelligence problems that use simpler constraints while preserving solution quality. One possible approach might be to start with a "simpler" SDP relaxation, and then iteratively tighten it by adding new constraints if and when they are needed (essentially, a form of dynamic cut generation for SDP relaxations). Another potential approach is to devise concise but effective constraint sets for SDP relaxations that are specialized to the types of problems commonly encountered in robotics (such as those dealing with, for example, states that model 3D orientations).
These two research arms are important steps to developing a general purpose POP solver for large-scale, real world problems.
Intellectual Merit: Our investigation into *large-scale, structured* SDPs isolates a specific but broadly-applicable problem class with substantial real-world impact. The SDP problem structure that our algorithms will exploit creates a meaningful mathematical separation between problems that are easy to solve and problems that are fundamentally intractable. This is analogous to the development of polynomial time interior point methods for linear programming compared to the nonpolynomial time of the simplex method on the Klee-Minty polytope. With this separation into a distinct problem subclass, we can beat the speed limit seemingly imposed by NP-hardness for a large class of problems. Our algorithms will demonstrate that mathematically developing scalable algorithms grants the ability to develop mathematically rich and realistic algorithms that can be deployed in real systems.
Broader Impacts: Optimization is a foundational tool across a wide variety of scientific and engineering fields. While our interest in solving large-scale but structured SDP relaxations of POPs is motivated by applications in machine intelligence, this technology is immediately useful for an extremely broad array of applications across science and engineering, such as robotics, computer vision, and mathematics [1].
Beyond its intrinsic scientific interest, another key goal of our work is to democratize access to this powerful class of optimization technologies, making them easy for practitioners to deploy in their own work. In the specific case of robotics, this work will serve as a strong step towards a future with safer, more trustworthy robotic autonomous systems. More generally, we expect that this will also serve as a useful motivating example, computational tool, and learning resource for applied and computational mathematics instruction.
References: [1] Lasserre, J.-B. (2010). Moments, Positive Polynomials and Their Applications. Imperial College Press. [2] Rosen, D. M., Doherty, K. J., Terán Espinonza, A., & Leonard, J. J. (2021). Advances in Inference and Representation for Simultaneous Localization and Mapping. Annual Review of Control, Robotics, and Autonomous Systems, 4, 215-242. [3] Rosen, D. M., Carlone, L., Bandeira, A. S., & Leonard, J. J. (2019, March). SE-Sync: A Certifiably Correct Algorithm for Synchronization over the Special Euclidean Group. The International Journal of Robotics Research, 38, 95-125. [4] Vandenberghe, L., Balakrishnan, V.R., Wallin, R., Hansson, A., Roh, T. Interior-Point Algorithms for Semidefinite Programming Problems Derived from the KYP Lemma. In: Henrion, D., Garulli, A. (eds) Positive Polynomials in Control. Lecture Notes in Control and Information Science, vol 312. Springer, Berlin, Heidelberg.