

Certifiable Robot Design Optimization using Differentiable Programming

Charles Dawson

Dept. of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, USA
cbd@mit.edu

Chuchu Fan

Dept. of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, USA
chuchu@mit.edu

Abstract—There is a growing need for computational tools to automatically design and verify autonomous systems, especially complex robotic systems involving perception, planning, control, and hardware in the autonomy stack. Differentiable programming has recently emerged as powerful tool for modeling and optimization. However, very few studies have been done to understand how differentiable programming can be used for robust, certifiable end-to-end design optimization. In this paper, we fill this gap by combining differentiable programming for robot design optimization with a novel statistical framework for certifying the robustness of optimized designs. Our framework can conduct end-to-end optimization and robustness certification for robotics systems, enabling simultaneous optimization of navigation, perception, planning, control, and hardware subsystems.

Using simulation and hardware experiments, we show how our tool can be used to solve practical problems in robotics. First, we optimize sensor placements for robot navigation (a design with 5 subsystems and 6 tunable parameters) in under 5 minutes to achieve an 8.4x performance improvement compared to the initial design. Second, we solve a multi-agent collaborative manipulation task (3 subsystems and 454 parameters) in under an hour to achieve a 44% performance improvement over the initial design. We find that differentiable programming enables much faster (32% and 20x, respectively for each example) optimization than approximate gradient methods. We certify the robustness of each design and successfully deploy the optimized designs in hardware. An open-source implementation is available at <https://github.com/MIT-REALM/architect>.

I. INTRODUCTION

To design complex systems, engineers in many fields use computer-aided tools to boost their productivity. Mechanical engineers can use a suite of 3D CAD (computer-aided design) and FEA (finite-element analysis) tools to design structures and understand their performance. Likewise, electrical engineers use electronic design automation tools, including hardware description languages like Verilog, to design and analyze large-scale, reliable, and yet highly complex integrated circuits. Sadly, when it comes to designing autonomous systems and robots, engineers often take an ad-hoc approach, relying heavily on experience and tedious parameter tuning.

C. Dawson is supported by the NSF GRFP under Grant No. 1745302. The Defense Science and Technology Agency in Singapore and IBM provided funds to assist the authors with their research, but this article solely reflects the opinions and conclusions of its authors and not DSTA Singapore, the Singapore Government, or IBM.

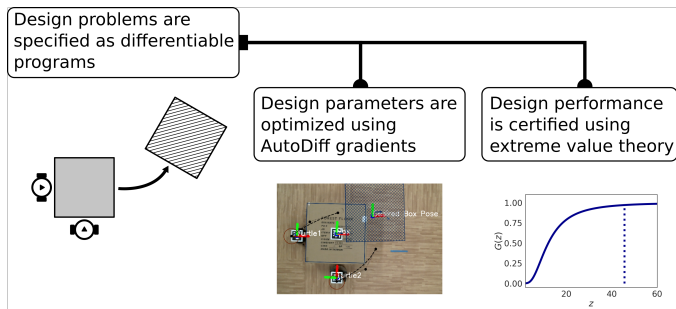


Fig. 1: An overview of our framework for robot design optimization and certification. Differentiable programming allows the user to flexibly specify a robot design problem, which can be efficiently optimized using exact gradients and verified using an extreme value statistical analysis.

Two factors have made it difficult to develop automated design tools for robotics. The first is complexity: most robots are composed of many interacting subsystems. Although some tools may aid in designing certain subsystems (e.g. Simulink for controllers, SolidWorks or CATIA for hardware, custom software for training perception systems), these tools cover only a small part of the overall robotics design problem, which includes sensing, actuation, perception, navigation, control, and decision-making subsystems. In addition to being interconnected, these subsystems often have a large number of parameters that require tuning to achieve good performance (neural network-based perception is an extreme example of this trend). Moreover, since few robotic systems are exactly alike, an effective design tool must allow the user to select an appropriate level of abstraction for the problem at hand. As a result, there is a need for flexible computational tools that can help designers optimize complex robotic systems.

The second difficulty is uncertainty. Robots operate in dynamic environments that cannot be fully specified *a priori*, and nonlinear interactions between the robot and its environment can make this uncertainty difficult to quantify. Nevertheless, we must account for this uncertainty during the design process and ensure that our designs perform robustly. The nature of this uncertainty can vary from problem to problem, reiterating the requirement that an automated design tool must be flexible

enough to adapt to different robot design problems.

To be successful, an automated robot design tool must address these two challenges (complexity and uncertainty). In addition, just as mechanical and electrical engineers use automated tools to both *design* and *verify* their designs, a robot design tool must enable its user to both design autonomous systems and certify the robustness of those designs. In this paper, we address these challenges by combining differentiable programming for design optimization with a novel statistical approach to design certification. In short:

- 1) We present a robot design optimization framework that is flexible (using differentiable programming to model complex systems) and robust (avoiding “brittle” optima).
- 2) We develop a novel statistical approach to certifying a design’s robustness to environmental uncertainty.
- 3) We validate our approach with experiments in simulation and hardware to show how our methods can be used to solve practical robot design problems.

Our goal is to develop a general-purpose robot design optimization tool that can be applied to a range of robot design problems with multiple subsystems. This goal is in contrast with other approaches that are restricted either to specific applications [1, 2, 3, 4, 5, 6] or subsystems [7]. To accomplish this goal, we make two novel contributions. The first is algorithmic: our approach builds on recent developments in programming languages (i.e. automatic differentiation) to provide the flexibility to model complex systems while still allowing fast gradient-based optimization. The second concerns certification: to ensure that our optimized designs are robust in the face of uncertainty, we pair design optimization with a novel statistical approach to robustness analysis.

Our experiments show that our methods can (in our first case study) optimize a robotic system with five subsystems and six design variables in under five minutes, achieving an 8.4x performance improvement over the initial design. In our second case study, we optimize a system with three subsystems and 454 design variables in under an hour, achieving a 44% performance improvement over the initial design. Our use of differentiable programming allows us to complete this optimization 32% and 20x faster, respectively in each example, compared to approximate gradient methods. Both of these designs are certified using a statistical robustness analysis and successfully deployed in hardware. An open-source implementation of our framework, including repeatable code examples, is available at <https://github.com/MIT-REALM/architect>. Our hope is that this prototype implementation will provide the foundation for a fully-featured, easy-to-use design tool for practicing robotics engineers.

II. RELATED WORK

1) *Design optimization for robotics*: Most existing works on design optimization for robotics focus on a particular application, such as simple walking robots [1], quadrotors [2], and soft robots [3, 4, 5]. Other works employ optimization to design specific subsystems, such as controllers [7] or motion plans [8]. In contrast, the purpose of this work is to develop

a general-purpose robot design optimization tool that can be applied not only to a range of robot design problems but also to optimize the design of multiple subsystems simultaneously. This goal is related to that of a large family of multi-disciplinary design optimization (MDO) methods in aerospace engineering [9]. As discussed above, our approach differs from MDO in its use of differentiable programming as a flexible modeling tool and our novel statistical approach to robustness analysis. We review the related work for automatic differentiation and robustness analysis in the next two sections.

2) *Programming languages for design optimization*: When it comes to managing complexity in a general-purpose design framework, programming languages are a natural tool. They allow users (i.e. programmers) to define precisely which abstractions are appropriate for any given application (e.g. by defining appropriate class hierarchies and function interfaces) without sacrificing generality. To take advantage of this expressivity, we can view engineering designs as programs that define the behavior of the system given suitable choices for design structure and parameters. We can then use automatic differentiation to derive gradients connecting these parameters to the system’s behavior and optimize accordingly. This view is inspired by recent work in 3D design optimization [10], aircraft design [11], and machine learning [12, 13].

In recent years, the robotics community has also developed special-purpose differentiable simulators for robotic systems, particularly those involving rigid body contact dynamics [14, 15, 16, 17]. These simulators have been used to solve system identification and controller design tasks, but they do not represent a general-purpose framework, as gradients are often derived by hand and the simulators are not expressive enough to model full-stack robotic systems (e.g. with perception and navigation capabilities). We take inspiration from these methods in our case studies, where we implement a simple differentiable contact simulator in our second case study.

3) *Formal methods for robustness analysis*: Safety and robustness are critical concerns for any robotic system. When it comes to low-level control, there is a rich history of reachability [18] and stability [19, 20] analysis tools that can be used to answer questions of safety and robustness for the control subsystem. Other works apply reachability analysis at the system level using black-box tools [21]. This work builds on this history by incorporating formal analysis into a design-optimize-analyze loop to provide rapid feedback on robustness as part of the design process. In particular, we develop a novel statistical method for quantifying the worst-case performance and sensitivity of an optimized design to external perturbations.

III. PRELIMINARIES AND ASSUMPTIONS

Key to the design of robotic systems is the tension between the factors a designer can control and those she cannot. For instance, a designer might be able to choose the locations of sensors and tune controller gains, but she cannot choose the sensor noise or disturbances (e.g. wind) encountered during operation. Robot design is therefore the process of choosing




Fig. 2: A glass-box model of a generic robotic system. Design optimization involves finding a set of design parameters so that the simulated cost is minimized, while robustness analysis involves quantifying how changes in the exogenous parameters affect the simulated cost.

feasible values for the controllable factors (here referred to as *design parameters*) that achieve good performance despite the influence of uncontrollable factors (*exogenous parameters*).

Of course, this is a deliberately narrow view of engineering design, since it focuses on parameter optimization and ignores important steps like problem formulation and system architecture selection. Our focus on parameter optimization is intentional, as it allows the designer to focus her creative abilities and engineering judgment on the architecture problem, using computational aids as interactive tools in a larger design process [10, 11]. This focus is common in design optimization (e.g. aircraft design in [11] and 3D CAD optimization in [10]).

To formalize the design optimization problem, we take a high-level view of the robot design problems (shown in Fig. 2), where a design problem has five components:

1) *Design parameters*: The system designer has the ability to tune certain continuous parameters $\theta \in \Theta \subseteq \mathbb{R}^n$; e.g., control gains or the positions of nodes in a sensor network.

2) *Exogenous parameters*: Some factors are beyond the designer’s control, such as wind speeds or sensor noise. We model these effects as random variables with some distribution $\phi \sim \Phi$ supported on a subset of \mathbb{R}^m . We assume no knowledge of Φ other than the ability to draw samples i.i.d..

3) *Simulator*: Given particular choices for θ and ϕ , the system’s state $s \in \mathcal{S}$ evolves in discrete time according to a known simulator $S : \Theta \times \Phi \mapsto \mathcal{S}^T$. This simulator describes the system’s behavior over a finite horizon T as a trace of states s_1, \dots, s_T . S should be deterministic; randomness must be “imported” via the exogenous parameters.

4) *Cost*: We assume access to a function $J : \mathcal{S}^T \mapsto \mathbb{R}$ mapping system behaviors (i.e. a trace of states) to a scalar performance metric that we seek to minimize.

5) *Constraints*: The choice of design parameters is governed by a set of constraints $c_i : \Theta \mapsto \mathbb{R}$ with index set $i \in \mathcal{I}_c$. Design parameters θ are feasible if $c_i(\theta) \geq 0 \forall i \in \mathcal{I}_c$. Here, we consider constraints as functions of θ only; we leave the extension to robust constraints involving ϕ to future work.

We can make this discussion concrete with an example: consider the autonomous ground vehicle (AGV) design problem illustrated in Fig. 3. In this problem, our goal is to design a localization and navigation system that will allow the AGV to safely navigate between two obstacles. The AGV can estimate its position using an extended Kalman filter (EKF) with noisy




Fig. 3: A design optimization problem for an AGV localization and navigation system. The goal is to find placements for two range sensors along with parameters for the navigation system that allow the robot to safely pass through the narrow doorway.

measurements of its range from two nearby beacons and its heading from an IMU. The robot uses this estimate with a navigation function [22] and feedback controller to track a collision-free path between the obstacles.

In this problem, the design parameters θ include the (x, y) locations of the two range beacons $b_1, b_2 \in \mathbb{R}^2$ and the feedback controller gains $k \in \mathbb{R}^2$. The exogenous parameters ϕ are the actuation and sensor noises at each timestep $w_t \in \mathbb{R}^3$ and $v_t \in \mathbb{R}^3$, drawn i.i.d. from Gaussian distributions $\mathcal{N}(0, Q)$ and $\mathcal{N}(0, R)$, respectively, as well as the initial state (also Gaussian). The simulator ξ integrates the AGV’s dynamics using a fixed timestep, updating the EKF and evaluating the navigation controller at each step. The cost function J assigns a penalty to collisions with the environment, estimation errors, and deviations from the goal location. We will return to this example in more detail in Section VI-A; first, we discuss our approach to design optimization and robustness analysis in Sections IV and V, respectively.

IV. DESIGN OPTIMIZATION

Given the notation from Section III, we can formally pose the robot design optimization problem. In formulating the optimization objective, it is important to consider the variance introduced by the exogenous parameters ϕ . Simply minimizing the expected value of the cost $\mathbb{E}_{\phi \sim \Phi} [J \circ S(\theta, \phi)]$ (where \circ denotes composition) can lead to myopic behavior where exceptional performance for some values of ϕ compensates for poor performance on other values; this is related to the phenomenon of “reward hacking” in reinforcement learning [23].

Ideally, we would like our designs to be robust to variations in exogenous parameters: changing ϕ should not cause the performance to change much. We can include this requirement as a heuristic by penalizing the variance of J . Intuitively, this heuristic “smooths” the cost function with respect to the exogenous parameters: regions of high variance (containing sharp local minima) are penalized, while regions of low variance are rewarded. We return to justify this connection to robustness in Section V-C. This heuristic leads us to the

variance-regularized robust design optimization problem:

$$\min_{\theta \in \Theta} \mathbb{E}_{\phi \sim \Phi} [J \circ S(\theta, \phi)] + \lambda \text{Var}_{\phi \sim \Phi} [J \circ S(\theta, \phi)] \quad (1a)$$

$$\text{s.t. } c_i(\theta) \geq 0 \quad \forall i \in \mathcal{I}_c \quad (1b)$$

Practically, we replace the expectation and variance with unbiased estimates over N samples $\phi_i \sim \Phi, i = 1, \dots, N$.

$$\min_{\theta \in \Theta} \frac{1}{N} \sum_{i=1}^N [J \circ S(\theta, \phi_i)] \quad (2a)$$

$$+ \lambda \left[\frac{\sum_{i=1}^N (J \circ S(\theta, \phi_i))^2}{N-1} - \frac{\left(\sum_{i=1}^N J \circ S(\theta, \phi_i)\right)^2}{(N-1)N} \right]$$

$$\text{s.t. } c_i(\theta) \geq 0 \quad \forall i \in \mathcal{I}_c \quad (2b)$$

Of course, these Monte-Carlo estimators will require multiple evaluations of $J \circ S$ to evaluate (2a). Since S might itself be expensive to evaluate, approximating the gradients of (2a) and (2b) using finite differences will impose a large computational cost ($2nN$ additional evaluations of $J \circ S$ and c_i at each step). Instead, we can turn to automatic differentiation (AD) to directly compute these gradients with respect to θ , which we can use with any off-the-shelf gradient-based optimization engine. The precise choice of optimization algorithm is driven by the constraints and is not central to our framework. If the constraints are hyper-rectangle bounds on θ , then algorithms like L-BFGS-B may be used, but if the constraints are more complex then sequential quadratic programming or interior-point methods may be used. Our implementation provides an interface to a range of optimization back-ends through SciPy [24], and we plan to add support for hybrid methods combining local gradient descent with gradient-free population methods in a future work.

In this framework, the user need only implement the simulator and cost function for their specific problem using a differentiable programming framework like the JAX library for Python [13], and this implementation can be used automatically for efficient gradient-based optimization. By implementing a library of additional building blocks in this AD paradigm (e.g. estimation algorithms like the EKF), we can provide an AD-based design optimization tool that strikes a productive balance between flexibility and ease of use. In the supplementary materials, we provide a prototype implementation of this tool, containing some of these AD building blocks. In future work, we hope to further expand this library to include more common robotics algorithms.

V. DESIGN CERTIFICATION VIA ROBUSTNESS ANALYSIS

Once we have found an optimal choice of design parameters, we need to verify that the design will be robust to uncertainty in the exogenous parameters. Similarly to 3D CAD and FEA packages for mechanical engineers, a successful design tool not only helps an engineer refine her design (i.e. using the design optimization framework in Section IV) but

also helps her analyze and predict its performance. To certify the performance of an optimized design, we are interested in two distinct questions. First, what is the maximum cost we can expect given variation in the exogenous parameters? Second, how sensitive is the cost to external disturbances: by how much can a change in the exogenous parameters increase the cost?

Answering these questions is difficult because we must extrapolate from a finite number of simulations to predict worst-case performance. To address this difficulty, we develop a probabilistic approach based on extreme value theory in statistics [25, 26, 27]. We begin by stating a relevant result:

Theorem V.1 (Extremal Types Theorem; 3.1.1 in [27]). *Let X_1, \dots, X_N be random variables drawn i.i.d. from an unknown distribution and $M_N = \max_i \{X_i\}$ be the sample maximum. If there exist sequences of normalizing constants $\{a_N > 0\}$ and b_N such that the limiting distribution of $(M_N - b_N)/a_N$ as $N \rightarrow \infty$ is non-degenerate, then*

$$\lim_{N \rightarrow \infty} \Pr[(M_N - b_N)/a_N \leq z] = G(z) \quad (3)$$

where $G(z)$ is a Generalized Extreme Value Distribution (GEVD) with location μ , scale σ , and shape ξ ,

$$G(z) = \exp \left\{ - \left[1 + \xi \left(\frac{z - \mu}{\sigma} \right) \right]^{-1/\xi} \right\}, \quad (4)$$

supported on $\{z : 1 + \xi(z - \mu)/\sigma > 0\}$.

In the special case $\xi = 0$, this distribution has a slightly different form (known as a Gumbel distribution), but the result holds. In practice, a_n and b_n are not estimated directly (this merely changes the fit values of μ and σ) and the GEVD is fit directly to M_N by either minimizing the log likelihood [27] or estimating the posterior distribution of (μ, σ, ξ) using Markov Chain Monte Carlo sampling [28]. A useful feature of the GEVD is that if our data suggest that $\xi < 0$, then the support of $G(z)$ is bounded above and we can estimate an upper bound on the maximum M_∞ . If $\xi \geq 0$, then we cannot estimate a strict upper bound, but we can provide for a confidence interval for M_∞ instead. In the following sections, we apply this theorem to analyze the robustness of an optimized design.

A. Estimating the worst-case performance

Our first robustness question concerns the worst-case performance of our design: given variation in ϕ , what is the maximum cost¹ we can expect for our choice of design parameters θ ? Our insight is that the variation $\phi \sim \Phi$ induces an (unknown) distribution in $J \circ S(\theta, \phi)$, so $J \circ S(\theta, \phi)$ a random variable to which the extremal types theorem applies. Algorithm 1 provides a means for estimating the maximum of $J \circ S(\theta, \phi)$ by fitting a GEVD to observed maximums M_N . Generally speaking, the block size N and sample size M should be chosen to be as large as computationally feasible to reduce the variance of the GEVD estimate [27].

¹Any function of the simulation trace can be substituted for cost without changing the framework.

Algorithm 1 An algorithm for estimating the parameters of a GEVD governing the expected maximum cost $J \circ S$

Require: Block size $N > 0$ and sample size $M > 0$

$X_j^i \leftarrow J \circ S(\theta, \phi_{ij})$; with $\phi_{ij} \sim \Phi$, $1 \leq j \leq N$, $1 \leq i \leq M$
 $M_N^i \leftarrow \max\{X_1^i, \dots, X_N^i\}$ for $i = 1, \dots, M$
 $(\mu, \sigma, \xi) \leftarrow$ posterior GEVD estimate given $\{M_N^i\}$

In practice, we use the automatic parallelization features of JAX to efficiently compute X_j^i and obtain the posterior distribution of μ , σ , and ξ using Markov Chain Monte Carlo sampling with the PyMC3 library [28]. From this posterior distribution, we take the 97% confidence level for each parameter (μ^*, σ^*, ξ^*) . If $\xi^* < 0$, we have confidence that the corresponding GEVD has bounded support on the right and estimate the maximum cost $J_{max} \leq \mu - \sigma/\xi$. Otherwise, we can estimate the 97% confidence level for J_{max} using the GEVD described by (μ^*, σ^*, ξ^*) .

B. Estimating sensitivity

In addition to the expected worst-case performance, it is also useful to know the sensitivity of that performance. That is, if the design performs well in one situation (i.e. for some value of ϕ), then how much can we expect its performance to degrade if ϕ changes? Formally, we define the sensitivity L as the least constant such that for any two $\phi_1, \phi_2 \sim \Phi$,

$$|J \circ S(\theta, \phi_1) - J \circ S(\theta, \phi_2)| \leq L \|\phi_1 - \phi_2\|$$

If $J \circ S$ is Lipschitz then L will be finite and equal the Lipschitz constant of $J \circ S$, but we do not require this assumption; if $J \circ S$ is not Lipschitz, then we can estimate a high-confidence upper bound on L .

In both cases, we can exploit the fact that L is an extreme value of the slope $|J \circ S(\theta, \phi_1) - J \circ S(\theta, \phi_2)|/\|\phi_1 - \phi_2\|$ and apply the extremal types theorem. Let $X = |J \circ S(\theta, \phi_1) - J \circ S(\theta, \phi_2)|/\|\phi_1 - \phi_2\|$ be a random variable with $\phi_1, \phi_2 \sim \Phi$. The distribution of X is unknown, but the extremal types theorem lets us characterize the sample maximum $L_N = \max\{X_1, \dots, X_N\}$ using a GEVD. Algorithm 2 provides our method for fitting this distribution, and a concrete Python implementation is provided in the supplementary materials. This approach is similar to that in [25, 29] but removes the assumption that L is bounded by fitting a GEVD instead of a reverse Weibull distribution, allowing our approach to apply when $J \circ S$ is not Lipschitz.

Algorithm 2 An algorithm for estimating the parameters of a GEVD governing the sensitivity of $J \circ S$

Require: Block size $N > 0$ and sample size $M > 0$

$X_j^i \leftarrow |J \circ S(\theta, \phi_{ij,1}) - J \circ S(\theta, \phi_{ij,2})|/\|\phi_{ij,1} - \phi_{ij,2}\|$,
with $\phi_{ij,1}, \phi_{ij,2} \sim \Phi$, $j = 1, \dots, N$, $i = 1, \dots, M$
 $L_N^i \leftarrow \max\{X_1^i, \dots, X_N^i\}$ for $i = 1, \dots, M$
 $(\mu, \sigma, \xi) \leftarrow$ posterior GEVD estimate given $\{L_N^i\}$

Algorithm 2 is similar to Algorithm 1, but the interpretation of the results differs in that the fit parameters from Algorithm 2

allow us to understand the sensitivity of a design. In particular, if the 97% confidence level for the shape parameter ξ^* is negative, then $J \circ S$ is likely Lipschitz continuous with Lipschitz constant $L \leq \mu - \frac{\sigma}{\xi}$. If $\xi > 0$, then $J \circ S$ is likely not Lipschitz but we can estimate the 97% confidence level for L . As a result, this statistical approach allows us to avoid making prior assumptions about the continuity of our system.

C. Connections to Design Optimization

Here, we will attempt to justify the variance regularization heuristic introduced in Section IV with reference to the worst-case performance J_{max} and sensitivity L computed by Algorithms 1 and 2. First, let's examine the connection with expected worst-case performance J_{max} . If we take the probability of observing a cost $J = J \circ S(\theta, \phi)$ within α of J_{max} ($0 < \alpha < 1$) and apply Cantelli's inequality [30], we see that

$$\Pr_{\phi \sim \Phi}(J \geq \alpha J_{max}) \leq \frac{\text{Var}_{\phi \sim \Phi}[J]}{\text{Var}_{\phi \sim \Phi}[J] + (\alpha J_{max} - \mathbb{E}_{\phi \sim \Phi}[J])^2}$$

Minimizing $\text{Var}_{\phi \sim \Phi}[J]$ in addition to $\mathbb{E}_{\phi \sim \Phi}[J]$ will correlate with decreasing this upper bound. As a result, we expect variance regularization to correlate with decreased probability of encountering near-worst-case performance.

We can also justify the connection between variance regularization and reducing sensitivity L by looking at the special case where $J \circ S$ is Lipschitz and the elements of ϕ are independent. The Bobkov-Houdré variance bound for Lipschitz functions [31] holds that $\text{Var}_{\phi \sim \Phi}[J] \leq L^2 \sigma_{\Sigma \phi}^2$, where $\sigma_{\Sigma \phi}^2$ is the variance of the sum of elements in ϕ . This bound does not explicitly show that minimizing $\text{Var}_{\phi \sim \Phi}[J]$ decreases L , but it suggests a correlation that we hope to revisit in future work.

VI. EXPERIMENTAL RESULTS

So far, we have developed the theoretical and algorithmic basis for our robot design framework. It remains for us to empirically answer two questions: first, is our framework useful for solving practical robot design problems? Second, is our statistical method for robustness analysis sound?

In this section, we answer these questions through the lens of two case studies. The first involves finding optimal sensor placements for robot navigation, and the second involves optimizing a pushing strategy for multi-agent manipulation. We demonstrate the success of our optimization and robustness analysis framework on each example, and we provide results from hardware testing in both cases. Next, we include an ablation study justifying our use of automatic differentiation and variance regularization. We conclude by verifying the soundness of our statistical robustness analysis.

A. Case study: optimal sensor placement for navigation

First, we return to the AGV localization and navigation example introduced in Fig. 3. This design problem requires finding an optimal placement for two ranging beacons to minimize estimation error and allow the robot to safely navigate between two obstacles. Range measurements from these beacons are integrated with IMU data via an EKF, and

the resulting state estimate is used as input to a navigation function and tracking feedback controller to guide the robot to its goal. This design problem has two important features. First, it involves interactions between multiple subsystems: the output from the EKF is used by the navigation function, which feeds input to the controller, which in turn influences future EKF predictions. Second, the effect of uncertainty on the robot’s performance is relatively strong.

The design parameters are the (x, y) locations of two range beacons and two feedback controller gains (6 total design parameters). The exogenous parameters include uncertainty in the robot’s initial state along with actuation and sensing noise at each of T timesteps ($3 + 6T$ total exogenous parameters). The cost function has three components: one penalizing large estimation errors, one penalizing deviations from the goal, and one penalizing collisions with the environment. A formal definition of the design and exogenous parameters, simulator, cost, and constraints is given in Table I in the appendix. We also include code in the supplementary materials for defining this design problem in our framework and running our design optimization and sensitivity analysis methods. The simulator and cost functions are implemented in Python using the JAX framework for automatic differentiation.

Fig. 4 compares simulated trajectories for the initial and optimized beacon placements and feedback gains, clearly showing the impact of design optimization. Initially, poor beacon placement causes the robot to accumulate estimation error and drift away from its goal. The optimized design moves the beacons off to the side to eliminate this drift. Optimization ($N = 512$, $\lambda = 0.1$, L-BFGS-B back-end) took 3 minutes 34 s on a laptop computer (8 GB RAM, 1.8 GHz 8-core processor).

We tested the initial and optimized design in hardware using the Turtlebot 3 platform. To emulate range beacon measurements in our lab, odometry and laser scan data were fused into a full state estimate from which range measurements were derived (the full state estimate was hidden from the robot, which only received the emulated range measurements). The control frequency was increased from 2 Hz in simulation to 10 Hz in hardware, and the obstacles were recreated in our laboratory. The hardware results, shown in Figs. 5 and 6, confirm our simulation results: the initial design suffers from drift and ends approximately 10 cm from its target position, while the optimized design does not drift and ends within 5 cm of the goal. This difference can be seen most clearly in the posterior error covariance from the EKF; Fig. 6 shows how the optimized design greatly reduces uncertainty in the state estimate compared to the initial design. No parameter estimation or tuning was required.

Finally, we apply the robustness analysis from Section V to certify the maximum absolute estimation error $\|x_t - \hat{x}_t\|$ in the optimized design (in meters, projected into the xy plane). Note that this error is different from the cost used during optimization, but we can still apply Algorithm 1 simply by changing the cost function for the duration of the analysis. Using block size $N = 1000$ and sample size $M = 1000$, we fit a GEVD using Algorithm 1 to the maximum estimation error

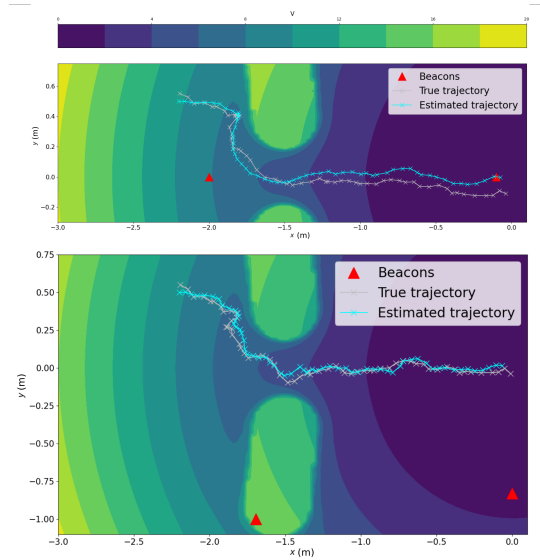


Fig. 4: Simulated trajectories for the initial (top) and optimized (bottom) AGV designs. Color indicates the value of the navigation function. Beacon positions are bounded within the area shown.



Fig. 5: Hardware performance of initial (left) and optimized (right) AGV designs. Square (green) shows the goal; triangles (red) show beacon locations. The optimized design eliminates drift relative to goal.

for both the initial and optimized designs. These distributions are shown in Fig. 7; the optimized design significantly reduces the expected maximum estimation error. We observe that the 97% confidence level for the shape parameter $\xi = 0.059$ is positive, so we cannot conclude that the worst-case estimation error is bounded, but we can derive a high-confidence bound of 0.21 m for our optimized design.




Fig. 6: Hardware results for EKF state estimates and posterior error covariance 3σ ellipse for initial and optimized designs.

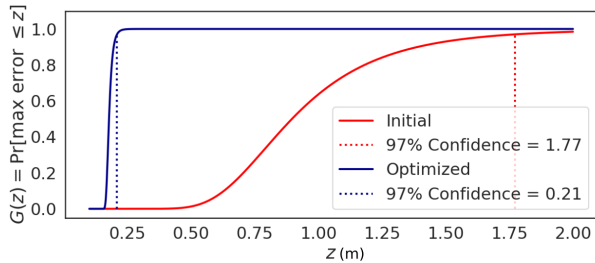


Fig. 7: GEVD CDF fit using Algorithm 1 for the maximum absolute estimation error in the xy -plane in both the initial and optimized designs, with 97% confidence levels.

B. Case study: collaborative multi-robot manipulation

Our second example involves finding a control strategy for multi-agent collaborative manipulation. In this setting, two ground robots must collaborate to push a box from its current location to a target pose (as in Fig. 8). Given the desired box pose and the current location of each robot, a neural network plans a trajectory for each robot, which the robots then track using a feedback controller (θ includes both the neural network parameters and the tracking controller gains, with a total of 454 design parameters). The exogenous parameters include the coefficient of friction for each contact pair, the mass of the box, the desired pose of the box, and the initial pose for each robot (a total of 13 exogenous parameters; we vary the desired box pose and initial robot poses to prevent over-fitting during optimization). The cost function is simply the squared error between the desired box pose (including position and orientation) and its true final pose after a 4s simulation. A full definition of this design problem and contact dynamics model is included in Table II in the appendix. We implement the contact dynamics simulator, trajectory planning neural network, and path tracking controller in Python using JAX.

Compared to the design problem in our first case study, this system has a simpler architecture (fewer subsystems) but more complicated dynamics and a much higher-dimensional design space. This example also showcases a different interpretation of the exogenous parameters: instead of representing true sources of randomness, these parameters represent quantities that are simply unknown at design-time. For example, the




Fig. 8: Multi-agent manipulation design optimization problem. The goal is to find parameters for robot controllers and a neural network planner that push the box from an initial position (solid) to a desired position (striped).




Fig. 9: GEVD CDF fit using Algorithm 2 for the maximum sensitivity of the optimized collaborative manipulation strategy to variation in friction coefficient. z has units of meters per unit change in friction coefficient.

target position for the box is not random in the same way as sensor noise in the previous example, but since we cannot choose this value at design-time it must be included in ϕ . As a result, minimizing the expected cost with respect to variation in ϕ yields a solution that achieves good performance for many different target poses, enabling the user to select one at run-time and be confident that the design will perform well.

To solve this design problem, the neural network parameters are initialized i.i.d. according to a Gaussian distribution, and the tracking controller gains are set to nominal values. We then optimize the parameters using $N = 512$, $\lambda = 0.1$, and L-BFGS-B back-end. This optimization took 45 minutes 32s on a laptop computer (8 GB of RAM and a 1.8 GHz 8-core processor). Fig. 10 shows a comparison between the initial and optimized strategies, and Fig. 13 in the appendix shows additional examples of the optimized behavior. The target pose is drawn uniformly $[x, y, \theta] \in [0, 0.5]^2 \times [-\pi/4, \pi/4]$, and the optimized design achieves a mean squared error of 0.0964.

We tested the optimized design in hardware, again using the Turtlebot 3 platform. An overhead camera and AprilTag [32] markers were used to obtain the location of the box and each robot. At execution, each robot first moves to a designated starting location near the box, plans a trajectory using the neural network policy, and tracks that trajectory at 100 Hz until the box reaches its desired location or a time limit is reached. Results from this hardware experiment are shown in Fig. 10, and a video is included in the supplementary materials. Again, no parameter tuning or estimation was needed.

After successfully testing the optimized design in the laboratory, it is natural to ask how its performance might change as conditions (particularly the coefficients of friction) change. Using $M = 500$ blocks of size $N = 1000$ each, we use Algorithm 2 to fit a GEVD for the sensitivity constant L with respect to the coefficients of friction between each contact pair. We do this by allowing these coefficients to vary and freezing other elements of ϕ at nominal values (box mass 1 kg and target pose $[0.3, 0.3, 0.3]$). The fit distribution is shown in Fig. 9. The 97% confidence level for the shape parameter is $\xi = 0.118 > 0$, so we cannot conclude that the performance of our design is Lipschitz with respect to the friction coefficients, but we can estimate the 97% confidence level for L as 0.63.

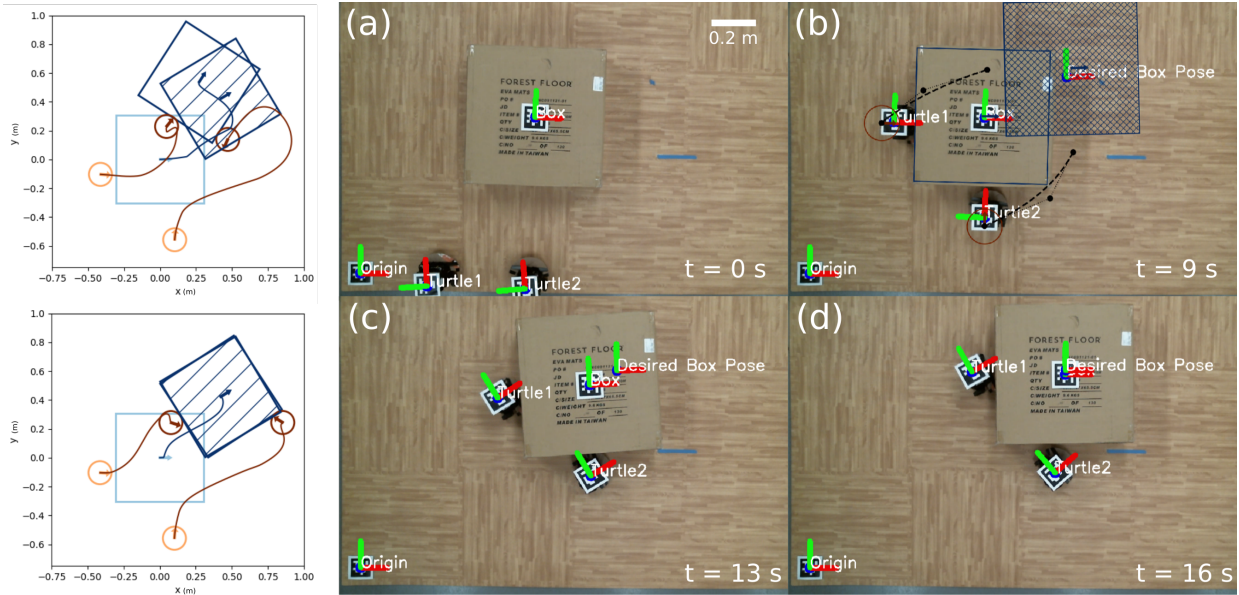


Fig. 10: Left: Initial (top) and optimized (bottom) manipulation strategies in simulation (light/dark colors indicate initial/final positions, stripes indicate desired position). Right: Optimized manipulation strategy deployed in hardware (video included in the supplementary materials). (a) The robots first move to positions around the box. (b) Using the optimized neural network, the robots plan a cubic spline trajectory pushing the box to its desired location. (c-d) The robots execute the plan by tracking that trajectory.

C. Design optimization ablation study

Our case studies in Sections VI-A and VI-B help demonstrate the utility of our framework for solving realistic robotics problems. However, it remains to justify the choices we made in designing this framework. For instance, how does automatic differentiation compare with other methods for estimating the gradient (e.g. finite differences)? What benefit does variance regularization in problem (2) bring? We answer these questions here using an ablation study where we attempt to isolate the impact of each of these features.

First, why use automatic differentiation? On the one hand, AD allows us to estimate the gradient with only a single evaluation of the objective function, while other methods (such as finite differences, or FD) require multiple evaluations. On the other hand, AD necessarily incurs some overhead at runtime, making each AD function call more expensive than those used in an FD scheme. Additionally, some arguments [17] suggest that exact gradients may be less useful than finite-difference or stochastic approximations when the objective is stiff or discontinuous. We compare AD with a 3-point finite-difference method by re-solving problem (2) for both case studies, keeping all parameters constant ($N = 512$, $\lambda = 0.1$, same random seed) and substituting the gradients obtained using AD for those computed using finite differences. Fig. 11 shows the results of this comparison. In the sensor placement example, AD achieves a lower expected cost and cost variance, and it runs in 32% less time. In the collaborative manipulation example, both methods achieve similar expected cost and variance, but the AD version runs nearly 19x faster. These


results lead us to conclude that AD enables more effective optimization than finite differences and is an appropriate choice for our framework. An exciting extension of our framework involves combining AD with stochastic population methods, but we leave this to future work.

The next question is whether variance regularization brings any benefit to the design optimization problem. To answer this question, we compare the results of re-solving both case studies with variance weight $\lambda = 0.1$ and $\lambda = 0$. These results are shown in Fig. 11; surprisingly, in the sensor placement example we see that the variance-regularized problem results in a lower expected cost, contrary to the intuition that regularization requires a trade off with increased expected cost. We expect that this lower expected cost may be a result of the regularization term smoothing the objective with respect to the exogenous parameters. However, these benefits are less pronounced than the benefits from automatic differentiation, and we do not see a distinct benefit in our second case study.

D. Accuracy of robustness analysis

To verify the soundness of our statistical robustness analysis methods, we need to determine whether the fit GEVD is likely to either under- or overestimate the worst-case performance of a design. Put simply, is our approach falsely optimistic (underestimating the worst-case) or conservative (overestimating)?

To answer these questions, we compare the cumulative distribution function (CDF) of the fit GEVD with an empirical CDF observed from data. Algorithms 1 and 2 both estimate a posterior distribution for μ , σ , and ξ , allowing us to construct an upper-bound and lower-bound GEVD using the 97% and



(a) AD vs. FD; sensor placement (b) AD vs. FD; manipulation (c) Effect of VR; sensor placement (d) Effect of VR; manipulation
 Fig. 11: (a)-(b) Improvement of automatic differentiation (AD) over finite differences (FD) in both case studies. (c)-(d) Effect of variance regularization (VR) in both case studies.

3% confidence level parameter estimates. Using these distributions, we can measure false optimism and conservatism using a one-sided Kolmogorov-Smirnov (KS) test [33].

Fig. 12 compares the estimated GEVDs and empirical data for worst-case performance in the sensor placement example (fit using Algorithm 1) and sensitivity in the manipulation example (fit using Algorithm 2). In the former case, we see that the empirical CDF lies between the upper- and lower-confidence limits for the fit distribution, indicating that the fit is neither falsely optimistic at the 97% level nor conservative at the 3% level (these conclusions are confirmed by the KS statistics provided in Table III in the appendix). In the latter case, even though the empirical CDF extends slightly beyond the estimated bounds in some regions, the statistical analysis in Table IV indicates that the estimated GEVD is neither falsely optimistic at the 97% level nor conservative at the 3% level. In addition, we see that the gap between the 3% and 97% distributions is relatively small in both examples in Fig. 12.

VII. DISCUSSION AND CONCLUSION

In this paper, we develop an automated design tool to improve the productivity of robot designers by a) enabling efficient optimization of robot designs and b) allowing users to certify the robustness of those designs. In developing this framework, we make two main algorithmic and theoretical contributions. First, we use differentiable programming for end-to-end optimization of robotic systems, creating a flexible software framework for design optimization. Second, we develop a novel statistical framework for certifying the worst-case performance and sensitivity of optimized designs.

To validate this framework and demonstrate the usefulness of our contributions, we present two case studies to highlight how our framework can be used for design optimization in practical robotics problems. Moreover, we show that our optimized designs are robust enough to deploy in hardware, and data from these hardware experiments validate our optimization approach. Finally, we provide an ablation study to justify the architecture of our optimization framework and a statistical analysis showing the soundness of our robustness analysis techniques. We hope that by combining flexible design optimization with robustness certification in our framework we can increase the productivity of robotics engineers, shorten the design cycle, and help bring more complex robotic systems to life.

There are a number of interesting directions for future work. First, since our approach relies on sampling from Φ without any further information, it will require a large number of samples to accurately capture rare events. We can close this gap when more information about Φ is available, perhaps using adversarial testing or importance sampling. Second, our framework is currently focused on tuning continuous parameters; we hope to incorporate stochastic search over discrete parameters in a future work. Finally, we hope to expand the software implementation of our framework to include a richer library of autonomy building blocks and demonstrate a wider range of applications in designing autonomous systems, including robotic arms, autonomous air and spacecraft, and networked autonomous systems.

REFERENCES

- [1] A. Schulz, C. Sung, A. Spielberg, W. Zhao, R. Cheng, E. Grinspun, D. Rus, and W. Matusik, “Interactive robogami: An end-to-end system for design of robots with ground locomotion,” *The International Journal of Robotics Research*, vol. 36, no. 10, pp. 1131–1147, 2017.
- [2] T. Du, A. Schulz, B. Zhu, B. Bickel, and W. Matusik, “Computational multicopter design,” *ACM Transactions on Graphics (TOG)*, vol. 35, no. 6, p. 227, 2016.
- [3] F. Chen and M. Y. Wang, “Design optimization of soft robots: A review of the state of the art,” *IEEE Robotics Automation Magazine*, vol. 27, no. 4, pp. 27–43, 2020.
- [4] T. Du, J. Hughes, S. Wah, W. Matusik, and D. Rus, “Underwater soft robot modeling and control with differentiable simulation,” *IEEE Robotics and Automation Letters*, 2021.
- [5] P. Ma, T. Du, J. Z. Zhang, K. Wu, A. Spielberg, R. K. Katzschmann, and W. Matusik, “Diffaqua: A differentiable computational design pipeline for soft underwater swimmers with shape interpolation,” *ACM Transactions on Graphics (TOG)*, vol. 40, no. 4, p. 132, 2021.
- [6] J. Zhang, J. Liu, C. Wang, Y. Song, and B. Li, “Study on multidisciplinary design optimization of a 2-degree-of-freedom robot based on sensitivity analysis and structural analysis,” *Advances in Mechanical Engineering*, vol. 9, no. 4, p. 1687814017696656, 2017. [Online]. Available: <https://doi.org/10.1177/1687814017696656>
- [7] J. Xu, T. Du, M. Foshey, B. Li, B. Zhu, A. Schulz, and W. Matusik, “Learning to fly: Computational controller

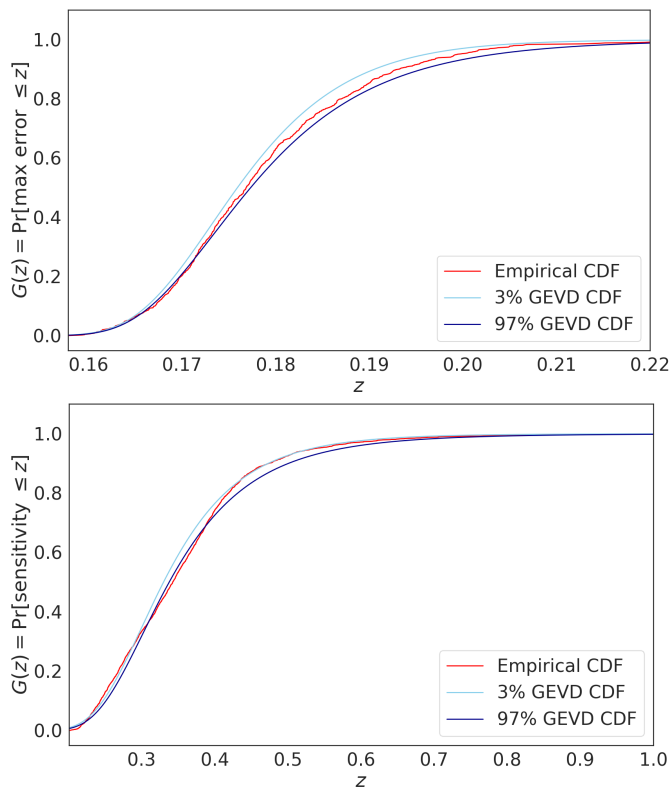


Fig. 12: Comparison of fit GEVD CDFs and empirical CDF for worst-case estimation error in the sensor placement example (top) and sensitivity in the manipulation example (bottom).

design for hybrid uavs with reinforcement learning,” *ACM Trans. Graph.*, vol. 38, no. 4, jul 2019. [Online]. Available: <https://doi.org/10.1145/3306346.3322940>

- [8] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, “Motion planning with sequential convex optimization and convex collision checking,” *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, Aug. 2014.
- [9] J. R. R. A. Martins and A. B. Lambe, “Multidisciplinary design optimization: A survey of architectures,” *AIAA Journal*, vol. 51, no. 9, pp. 2049–2075, Sept. 2013.
- [10] D. Cascaval, M. Shalah, P. Quinn, R. Bodik, M. Agrawala, and A. Schulz, “Differentiable 3d cad programs for bidirectional editing,” *arXiv*, vol. abs/2110.01182, 2021.
- [11] P. D. Sharpe, “Aerosandbox: A differentiable framework for aircraft design optimization,” Master’s thesis, MIT, 2021.
- [12] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035.
- [13] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, “JAX: composable transformations of Python+NumPy programs,” 2018. [Online]. Available: <http://github.com/google/jax>
- [14] E. Heiden, D. Millard, E. Coumans, Y. Sheng, and G. S. Sukhatme, “NeuralSim: Augmenting differentiable simulators with neural networks,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2021. [Online]. Available: <https://github.com/google-research/tiny-differentiable-simulator>
- [15] F. de Avila Belbute-Peres, K. Smith, K. Allen, J. Tenenbaum, and J. Z. Kolter, “End-to-end differentiable physics for learning and control,” in *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., vol. 31. Curran Associates, Inc., 2018.
- [16] R. Tedrake and the Drake Development Team, “Drake: Model-based design and verification for robotics,” 2019. [Online]. Available: <https://drake.mit.edu>
- [17] H. Suh, T. Pang, and R. Tedrake, “Bundled gradients through contact via randomized smoothing,” *ArXiv*, vol. abs/2109.05143, 2021.
- [18] M. Althoff, G. Frehse, and A. Girard, “Set propagation techniques for reachability analysis,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, no. 1, pp. 369–395, 2021. [Online]. Available: <https://doi.org/10.1146/annurev-control-071420-081941>
- [19] Y.-C. Chang, N. Roohi, and S. Gao, “Neural lyapunov control,” in *NeurIPS*, 2019.
- [20] C. Dawson, Z. Qin, S. Gao, and C. Fan, “Safe nonlinear control using robust neural lyapunov-barrier functions,” in *5th Annual Conference on Robot Learning*, 2021. [Online]. Available: https://openreview.net/forum?id=8K5kisAnb_p
- [21] C. Fan, B. Qi, S. Mitra, and M. Viswanathan, “Dryvr: Data-driven verification and compositional reasoning for automotive systems,” in *Computer Aided Verification*, R. Majumdar and V. Kuncak, Eds. Cham: Springer International Publishing, 2017, pp. 441–461.
- [22] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach (2nd Edition)*, December 2002. [Online]. Available: <http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/0137903952>
- [23] D. Amodei, C. Olah, J. Steinhardt, P. F. Christiano, J. Schulman, and D. Mané, “Concrete problems in ai safety,” *ArXiv*, vol. abs/1606.06565, 2016.
- [24] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett,

- J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [25] K. Sridhar, O. Sokolsky, I. Lee, and J. Weimer, “Improving neural network robustness via persistency of excitation,” *arXiv*, 2021. [Online]. Available: <https://arxiv.org/abs/2106.02078>
- [26] G. Wood and B. Zhang, “Estimation of the lipschitz constant of a function,” *Journal of Global Optimization*, vol. 8, no. 1, 1996.
- [27] S. Coles, *An introduction to statistical modeling of extreme values*. Springer, 2001.
- [28] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, “Probabilistic programming in python using pymc3,” *PeerJ Computer Science*, vol. 2, 2016.
- [29] C. Knuth, G. Chou, N. Ozay, and D. Berenson, “Planning with learned dynamics: Probabilistic guarantees on safety and reachability via lipschitz constants,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5129–5136, 2021.
- [30] S. Boucheron, G. Lugosi, and P. Massart, *Concentration inequalities: A nonasymptotic theory of Independence*. Oxford University Press, 2016.
- [31] S. G. Bobkov and C. Houdré, “Variance of lipschitz functions and an isoperimetric problem for a class of product measures,” *Bernoulli*, vol. 2, no. 3, pp. 249–255, 1996. [Online]. Available: <http://www.jstor.org/stable/3318522>
- [32] E. Olson, “AprilTag: A robust and flexible visual fiducial system,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2011, pp. 3400–3407.
- [33] “Kolmogorov-smirnov goodness-of-fit test.” [Online]. Available: <https://www.itl.nist.gov/div898/handbook/eda/section3/eda35g.htm>

APPENDIX

SENSOR PLACEMENT DESIGN PROBLEM STATEMENT

We model the robot with discrete-time Dubins dynamics with three state variables ($q = [x, y, \theta]$), two control inputs for linear and angular velocity ($u = [v, \omega]$), and noisy transition model

$$\begin{bmatrix} x \\ y \\ \theta \end{bmatrix}_{t+1} = \begin{bmatrix} x \\ y \\ \theta \end{bmatrix}_t + \begin{bmatrix} \Delta t v \cos(\theta + \Delta t \omega / 2) \\ \Delta t v \sin(\theta + \Delta t \omega / 2) \\ \Delta t \omega \end{bmatrix} + w_t$$

where $\Delta t = 0.5$ and $w_t \in \mathbb{R}^3$ is the actuation noise ($w_t \sim \mathcal{N}(0, Q)$ with covariance $Q \in \mathbb{R}^{3 \times 3}$). The measurement model

TABLE I: Formal statement of the sensor placement design problem with T discrete timesteps.

Design parameters	$\theta = [b_1, b_2, k] \in \mathbb{R}^6$ Beacon locations: $b_i = (x_{b_i}, y_{b_i}) \in \mathbb{R}^2$ for $i = 1, 2$ Feedback gains: $k \in \mathbb{R}^2$
Exogenous parameters	$\phi = [q_0, w_0, \dots, w_{T-1}, v_0, \dots, v_{T-1}] \in \mathbb{R}^{3+6T}$ Initial state: $q_0 \in \mathbb{R}^3$, $q_0 \sim \mathcal{N}(\bar{q}_0, P_0)$; $P_0 = 0.001 I_{3 \times 3}$ Actuation noise: $w_t \in \mathbb{R}^3$, $w_t \sim \mathcal{N}(0, Q)$; $Q = (\Delta t)^2 \text{diag}([0.001, 0.001, 0.01])$ Measurement noise: $v_t \in \mathbb{R}^3$, $v_t \sim \mathcal{N}(0, R)$ $R = \text{diag}([0.1, 0.01, 0.01])$
Simulator	S initializes the robot with state q_0 and EKF state estimate \bar{q}_0 and error covariance P_0 , then steps forward with interval $\Delta t = 0.5$ for $T = 60$ total steps. At each step, the simulator <ol style="list-style-type: none"> 1) Evaluates the navigation function to find a collision-free path to the goal, 2) Uses a feedback controller to track that path, 3) Updates the state using forward Euler integration, 4) Performs an EKF prediction, obtains a measurement z_t, and performs an EKF update. S returns a trace $s_t = [q, \hat{q}, P_t, V_t]$ containing true states, estimated states, estimated posterior error covariance, and the value of the navigation function at each time step.
Cost	J has three components. The first ($\ q_t - \hat{q}_t\ ^2$) minimizes the estimation error of the EKF, the second ($\ q_t\ $) guides the robot towards the goal, and the third (both V_t terms) avoids collision with the environment: $J = \frac{1}{T} \sum_{t=1}^T (100 \ q_t - \hat{q}_t\ ^2 + \ q_t\ ^2 + 0.1 V_t) + 0.1 \max_t V_t$
Constraints	$(x_{b_i}, y_{b_i}) \in [-3, 0] \times [-1, 1]$ for $i = 1, 2$

is

$$z_t = \begin{bmatrix} (x_t - x_{b_1})^2 + (y_t - y_{b_1})^2 \\ (x_t - x_{b_2})^2 + (y_t - y_{b_2})^2 \\ \theta \end{bmatrix} + v_t$$

where v_t is the measurement noise ($v_t \sim (0, R)$ and covariance $R \in \mathbb{R}^{3 \times 3}$), modeling range measurements from radio or acoustic beacons b_1 and b_2 and inertial or magnetic measurements of θ . The initial state of the robot is normally distributed $q_0 \sim \mathcal{N}(\bar{q}_0, P_0)$ for mean initial state $\bar{q}_0 \in \mathbb{R}^3$ and initial covariance $P_0 \in \mathbb{R}^{3 \times 3}$. The navigation function (shown in Fig. 4) is $V_t(x_t, y_t) = 2(x_t^2 + y_t^2) + 0.05/d_t$ (d_t is the distance from the robot to the nearest obstacle at step t). Formally, we define this problem in the language of our framework in Table I.

MULTI-AGENT MANIPULATION DESIGN PROBLEM STATEMENT

We model each ground robot as a double integrator with states $[p_x, p_y, \theta, v_x, v_y, \omega]$. Given control inputs representing desired linear velocity v_d in the $[\cos \theta, \sin \theta]$ direction and desired angular velocity ω_d , the robot tracks those desired velocities by applying forces and torques subject to a friction cone constraint. The box is modeled as a rigid body with friction against the ground. Contact forces between the

box and each robot are modeled using a penalty method described in [17], where the normal force is given by $f_n = k_c \min(\phi, 0) - k_d \dot{\phi} \mathbb{1}_{\phi < 0}$ (ϕ is the signed distance between the robot and the box, $k_c = 300 \text{ N/m}$ is the contact stiffness, k_d is a damping coefficient chosen to ensure critical damping, and $\mathbb{1}_{\phi < 0}$ is the indicator function equal to 1 when the box and robot are in contact and 0 otherwise). Friction in the box/ground and box/robot contacts was modeled as Coulomb friction, resulting in a tangential force $f_t = \mu f_n$ with $\mu = c\psi$ if $\psi < \psi_s$ and $\mu = \mu_d$ otherwise, where μ_d is the coefficient of dynamic friction (μ_d varies for each contact pair), ψ is the tangential velocity at the point of contact, $\psi_s = 0.3 \text{ m/s}$ is the tangential velocity where slipping begins, and $c = \mu_d/\psi_s$ was chosen to ensure a continuous friction model.

Each ground robot uses a proportional controller (with tunable gains) to find v_d and ω_d to track a cubic spline reference trajectory. The start point of each spline is set to match the robot’s current position, the end point is set based a known offset from the desired box location, and the central control point of the spline is set using a neural network (with tunable parameters). The neural network is given inputs including the current position of each robot and the desired box pose, all referenced against the current box pose, and it predicts (x, y) locations for the control point for each robot. The network uses tanh activations on each hidden layer.

Formally, we define this problem in the language of our framework (design parameters, exogenous parameters, etc.) in Table II. The design parameters include the trajectory tracking control gains and network parameters, while the exogenous parameters include the desired box pose, coefficients of friction, box mass, and initial robot poses.

KOLMOGOROV-SMIRNOV TEST RESULTS

Table III provides results from one-sided KS tests for the GEVD estimated using Algorithm 1 in the sensor placement case study, while Table IV provides similar results for Algorithm 2 in the collaborative manipulation case study.

TABLE II: Formal statement of the collaborative manipulation design problem using a planning network with n_p total parameters (weights and biases).

Design parameters	$\theta = [k_v, k_\omega, w_i, b_i] \in \mathbb{R}^{2+n_p}$ Trajectory tracking gains: $[k_v, k_\omega] \in \mathbb{R}^2$ Network weights and biases: (w_i, b_i) for $i = 1, \dots, n_p$
Exogenous parameters	$\phi = [\mu_{rg}, \mu_{bg}, \mu_{br}, m_b, p_{bd}, p_{r1}, p_{r2}] \in \mathbb{R}^{13}$ Robot/ground, box/ground, box/robot coefficients of friction: $[\mu_{rg}, \mu_{bg}, \mu_{br}] \in [0.6, 0.8] \times [0.4, 0.6] \times [0.1, 0.3]$ Box mass: $m_b \in [0.9, 1.1]$ Desired box pose: $p_{bd} = [x_d, y_d, \theta_d] \in [0, 0.5]^2 \times [-\pi/4, \pi/4]$ (Above parameters are uniformly distributed) Initial robot pose: $p_{ri} = [x_0, y_0, \theta_0] \sim \mathcal{N}(\bar{p}_{ri}, \Sigma)$ $\Sigma = 0.01 I_{3 \times 3}, i = 1, 2.$
Simulator	S initializes the robots at the initial states in ϕ relative to the box. Since these initial states may be in contact, we simulate 0.5 s of settling time at a 0.01 s timestep, then re-index the robot positions and desired box pose relative to the settled box pose. We then evaluate the planning network and track the planned path for 4 s at a 0.01 s timestep. At each timestep, 1) evaluate the spline tracking controller, 2) evaluate contact dynamics between the box, robots, and ground, and 3) integrate forces and torques to obtain box and robot states at the next timestep. S returns a trace $s_t = [q_{r1}, q_{r2}, q_b]$ containing the states of each robot and the box over time (relative to the initial pose of the box after the settling period).
Cost	J is simply the squared distance between the final box pose and the desired box position $(x-x_d)^2 + (y-y_d)^2 + (\theta-\theta_d)^2$
Constraints	Network parameters were not constrained. k_v and k_ω were constrained to be less than 10.

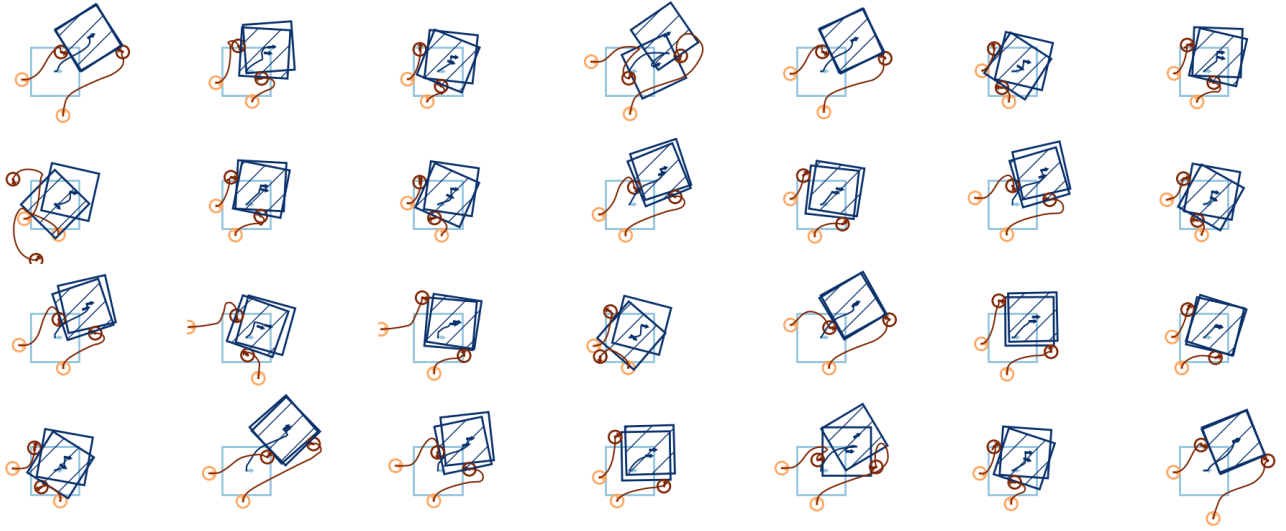


Fig. 13: Additional examples of optimized multi-agent manipulation behavior in simulation, showing that the optimized strategy reaches the goal in most cases. Each example shows the results of executing the optimized pushing strategy for 4s with a randomly selected set of friction coefficients, random target pose, and random initial robot poses. Light/dark colors indicate initial/final positions, respectively, and the striped box indicates the target pose.

	Null Hypothesis	KS Statistic	p-value	Conclusion ($p < 0.05$)
False Optimism	97% GEVD under-estimates worst-case performance	0.0410	0.0337	Reject; 97% GEVD <i>does not</i> under-estimate worst-case performance
Conservatism	3% GEVD over-estimates worst-case performance	0.0529	0.00354	Reject; 3% GEVD <i>does not</i> over-estimate worst-case performance

TABLE III: Results of one-sided KS tests for the sensor placement case study. These results indicate that Algorithm 1 is sound in this case.

	Null Hypothesis	KS Statistic	p-value	Conclusion ($p < 0.05$)
False Optimism	97% GEVD under-estimates sensitivity	0.0399	6.75×10^{-5}	Reject; 97% GEVD <i>does not</i> under-estimate sensitivity
Conservatism	3% GEVD over-estimates sensitivity	0.0618	1.03×10^{-10}	Reject; 3% GEVD <i>does not</i> over-estimate sensitivity

TABLE IV: Results of one-sided KS tests for the collaborative manipulation case study. These results indicate that Algorithm 2 is sound in this case.