Proceedings of Machine Learning Research vol 120:1–11, 2020

Robust Guarantees for Perception-Based Control

Sarah Dean DEAN_SARAH@BERKELEY.EDU
Nikolai Matni NMATNI@SEAS.UPENN.EDU
Benjamin Recht BRECHT@BERKELEY.EDU
Vickie Ye VYE@BERKELEY.EDU

Abstract

Motivated by vision-based control of autonomous vehicles, we consider the problem of controlling a known linear dynamical system for which partial state information, such as vehicle position, is extracted from complex and nonlinear data, such as a camera image. Our approach is to use a learned perception map that predicts some linear function of the state and to design a corresponding safe set and robust controller for the closed loop system with this sensing scheme. We show that under suitable smoothness assumptions on both the perception map and the generative model relating state to complex and nonlinear data, parameters of the safe set can be learned via appropriately dense sampling of the state space. We then prove that the resulting perception-control loop has favorable generalization properties. We illustrate the usefulness of our approach on a synthetic example and on the self-driving car simulation platform CARLA.

Keywords: Robust control, learning theory, generalization, perception, robotics.

1. Introduction

Incorporating insights from rich, perceptual sensing modalities such as cameras remains a major challenge in controlling complex autonomous systems. While such sensing systems clearly have the potential to convey more information than simple, single output sensor devices, interpreting and robustly acting upon the high-dimensional data streams remains difficult. Recent end-to-end approaches tackle the problem of image based control by learning an optimized map from pixel values directly to low level control inputs. Though there has been tremendous success in accomplishing sophisticated tasks, critical gaps in understanding robustness and safety still remain (Levine et al., 2016). On the other hand, methods rooted in classical state estimation and robust control explicitly characterize a model of the underlying system and its environment in order to design a feedback controller. These approaches have provided strong and rigorous guarantees of robustness and safety in domains such as aerospace and process control, but the level of specification of the underlying system required has thus far limited their impact for complex sensor inputs.

In this paper, we aim to leverage contemporary techniques from machine learning and robust control to understand the conditions under which safe and reliable behavior can be achieved in controlling a system with uncertain and high-dimensional sensors. Whereas much recent work has been devoted to proving safety and performance guarantees for learning-based controllers applied to systems with unknown dynamics (Wabersich and Zeilinger, 2018; Akametalu et al., 2014; Ostafew et al., 2014; Hewing and Zeilinger, 2017; Dean et al., 2017, 2018; Abbasi-Yadkori et al., 2019; Abbasi-Yadkori and Szepesvári, 2011; Cheng et al., 2019; Taylor et al., 2019; Williams et al., 2018; Agarwal et al., 2019; Berkenkamp et al., 2017), we focus on the practical scenario where the underlying dynamics of a system are well understood, and it is instead the integration of perceptual sensor data into the control loop that must be learned.

Specifically, we consider controlling a known linear dynamical system for which partial state information can only be extracted from complex observations. Our approach is to design a *virtual sensor* by learning both

a perception map (i.e., a map from observations to a linear function of the state) and a bound on its estimation errors. We show that under suitable smoothness assumptions, we can guarantee bounded errors within a neighborhood of the training data. This model of uncertainty allows us to synthesize a robust controller that ensures that the system does not deviate too far from states visited during training. Our main result shows that the perception and robust control loop is able to robustly generalize under adversarial noise models. To the best of our knowledge, this is the first such guarantee for a vision-based control system.

Related Work We refer to our full technical report (Dean et al., 2019) for an extensive treatment of related work, and give only a brief overview here. There is a rich body of work that integrates cameras into estimation, planning, and control loops. Techniques that focus on estimation, integrating camera measurements with inertial odometry via an Extended Kalman Filter and through Simultaneous Localization and Mapping (Jones and Soatto, 2011; Lynen et al., 2013), can be leveraged to enable aggressive control maneuvers, for example in unmanned aerial vehicles (Tang et al., 2018). The machine learning community has taken a more data-driven approach. The earliest such example is likely Pomerleau (1989), in which a 3-layer neural-network is trained to infer road direction from images. More recent work tackles low level vision-based control via imitation learning (Codevilla et al., 2018), resulting in policies that map pixels directly to low-level control inputs. Inspired by these works, our theoretical contributions are similar in spirit to those of the online learning community, in that we provide generalization guarantees under adversarial noise models (Hassibi and Kaliath, 2001; Yasini and Pelckmans, 2018).

Notation We use letters such as x and A to denote vectors and matrices, and boldface letters such as $\boldsymbol{x}$ and $\boldsymbol{\Phi}$ to denote infinite horizon signals and linear convolution operators. Thus, for $\boldsymbol{y} = \boldsymbol{\Phi}\boldsymbol{x}$, we have by definition that $y_k = \sum_{t=0}^{k} \Phi_t x_{k-t}$. We write $x_{0:t} = \{x_0, x_1, \ldots, x_t\}$ for the history of signal $\boldsymbol{x}$ up to time t. For a function $x_k \mapsto f_k(x_k)$, we write $\boldsymbol{f}(\boldsymbol{x})$ to denote the signal $\{f_k(x_k)\}_{k=0}^{\infty}$. We overload the norm $\|\cdot\|$ so that it applies equally to elements x_k, signals $\boldsymbol{x}$, and linear operators $\boldsymbol{\Phi}$. For any element norm, we define the signal norm as $\|\boldsymbol{x}\| = \sup_k \|x_k\|$ and the linear operator norm as $\|\boldsymbol{\Phi}\| = \sup_{\|\boldsymbol{w}\| \le 1} \|\boldsymbol{\Phi}\boldsymbol{w}\|$. We primarily focus on the triple $(\|x_k\|_\infty, \|\boldsymbol{x}\|_\infty, \|\boldsymbol{\Phi}\|_{\mathcal{L}_1})$. Note that as $\|\boldsymbol{\Phi}\|$ is an induced norm, it satisfies the sub-multiplicative property $\|\boldsymbol{\Phi}\boldsymbol{\Psi}\| \le \|\boldsymbol{\Phi}\|\|\boldsymbol{\Psi}\|$. We define the ball $B_r(x_0) = \{x \ : \ \|x - x_0\| \le r\}$. We say that a fuction f is locally S-slope bounded for a radius r around x_0 if for $x \in B_r(x_0)$, $\|f(x) - f(x_0)\| \le S\|x - x_0\|$, or similarly locally L-Lipschitz if for $x, y \in B_r(x_0)$, $\|f(x) - f(y)\| \le L\|x - y\|$.

2. Problem setting

Consider the LTI dynamical system

$$x_{k+1} = Ax_k + Bu_k + Hw_k \,, \tag{1}$$
$$z_k = q(x_k) \tag{2}$$

with system state $x \in \mathbb{R}^n$, control input $u \in \mathbb{R}^m$, disturbance $w \in \mathbb{R}^w$, and observation $z \in \mathbb{R}^M$. We take the dynamics matrices (A, B, H) to be known. The observation process is determined by the unknown *generative model* q, which is nonlinear and potentially quite high-dimensional. As an example, consider a camera affixed to the dashboard of a car tasked with driving along a road. Here, the observations $\{z_k\}$ are the captured images and the map q generates these images as a function of position and velocity.

Motivated by such vision based control systems, our goal is to solve the optimal control problem

$$\text{minimize}_{\{\gamma_k\}} \quad c(\boldsymbol{x}, \boldsymbol{u}) \ \text{ subject to } \ \text{dynamics (1) and measurement (2)}, \ u_k = \gamma_k(z_{0:k}), \tag{3}$$

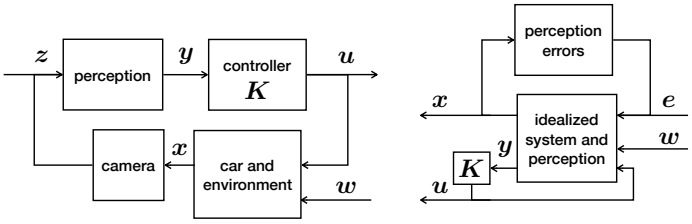

Figure 1: (Left) A diagram of the proposed perception-based control pipeline. (Right) The conceptual rearrangement of the closed-loop system permitted through our perception error characterization.

where $c(\boldsymbol{x}, \boldsymbol{u})$ is a suitably chosen cost function and γ_k is a measurable function of the image history $z_{0:k}$. This problem is made challenging by the nonlinear, high-dimensional, and unknown generative model.[1]

Suppose instead that there exists a *perception map* p that imperfectly predicts partial state information; that is $p(z_k) = Cx_k + e_k$ for $C \in \mathbb{R}^{\ell \times n}$ a known matrix, and error $e_k \in \mathbb{R}^{\ell}$. Such a matrix C might be specified to encode, for example, that camera images provide good signal on position, but not velocity or acceleration. We therefore define a new measurement model in which the map p plays the role of a noisy sensor:

$$y_k = p(z_k) = Cx_k + e_k. \tag{4}$$

This allows us to reformulate problem (3) as a *linear* optimal control problem, where the measurements are defined by (4) and the control law $u_k = \boldsymbol{K}(y_{0:k})$ is a *linear* function of the outputs of past measurements $y_{0:k}$. As illustrated in Figure 1, guarantees of performance, safety, and robustness require designing a controller which suitably responds to system disturbance $\boldsymbol{w}$ and sensor noise $\boldsymbol{e}$.

For linear optimal control problems, a variety of cost functions and noise models have well-understood solutions; for further background, see the full technical report (Dean et al., 2019). Perhaps the most well known is the combination of *Kalman filtering* with static state feedback, which arises as the solution to the linear quadratic Guassian (LQG) problem. However, the perception errors $\boldsymbol{e}$ do not necessarily obey assumptions about measurement noise made in traditional optimal control, and must be handled carefully.

In light of this discussion, we approach our problem with the following data-driven procedure: First, we use supervised learning to fit a perception map and explicitly characterize its errors in terms of the data it was trained on. Second, we compute a robust controller that mitigates the effects of the measurement error $\boldsymbol{e}$. We will show that under suitable local smoothness assumptions on the generative model q and perception map p we can guarantee closed-loop robustness using a number of samples depending only on the state dimension.

3. Data-dependent perception error

In this section, we introduce a procedure to estimate regions for which a learned perception map can be used safely during operation. We first suppose access to initial training data $\mathcal{S}_0 = \{(x_i, z_i)\}_{i=1}^{N_0}$, used to learn a perception map via any of the wide variety of traditional supervised methods. We then estimate safe regions around the training data, potentially using parameters learned from a second dataset. The result is

1. We remark that a nondeterministic appearance map q may be of interest for modeling phenomena like noise and environmental uncertainty. While we focus our exposition on the deterministic case, we note that many of our results can be extended in a straightforward manner to any noise class for which the perception map has a bounded response. For example, many computer vision algorithms are robust to random Gaussian pixel noise, gamma corrections, or sparse scene occlusions.

a characterization of how quickly the learned perception map degrades as we move away from the initial training data.

We will describe the regions of the state space within which the sensing is reliable using a safe set which approximates sub-level sets of the error function $e(x) = p(q(x)) - Cx$. We make this precise in the following Lemma, defining a safe set which is valid under an assumption that the error function is locally slope bounded around training data. We defer the proof, and the proofs of all results to follow, to the full technical report (Dean et al., 2019).

Lemma 1 (Closeness implies generalization) *Suppose that $p \circ q - C$ is locally S-slope bounded with a radius of r around training datapoints. Define the safe set*

$$\mathcal{X}_\gamma = \bigcup_{(x_d, z_d) \in \mathcal{S}_0} \{x \in B_r(x_d) \ : \ \|p(z_d) - Cx_d\| + S\|x - x_d\| \leq \gamma\} \, . \tag{5}$$

Then for any (x, z) with $x \in \mathcal{X}_\gamma$, the perception error is bounded: $\|p(z) - Cx\| \leq \gamma$.

The validity of the safe set $\mathcal{X}_\gamma$ depends on bounding the slope of the error function locally around the training data. We remark that this notion of slope boundedness has connections to sector bounded nonlinearities, a classic setting for nonlinear system stability analysis (Desoer and Vidyasagar, 1975).

Deriving the slope boundedness of the error function relies on the learned perception map as well as the underlying generative model, so we propose a second learning step to estimate a bound on S. We use an additional dataset composed of samples around each training data point: $\mathcal{S}_N = \cup_{x_d \in \mathcal{S}_0} \mathcal{S}_{x_d}$, where each $\mathcal{S}_{x_d}$ contains N points densely sampled from $B_r(x_d)$. For each x_d in the training set, we then fix a radius r and estimate S with the maximum observed slope:

$$\widehat{S}_{x_d} = \max_{x_i \in \mathcal{S}_{x_d}} s(x_i, x_d) \, , \quad \text{where } \ s(x, x') = \|e(x) - e(x')\| / \|x - x'\|. \tag{6}$$

By taking the maximum over training datapoints x_d, the slope bound S can be estimated with $|\mathcal{S}_0|N$ samples.

Proposition 2 *Assume that the error function $p \circ q - C$ is locally L-Lipschitz for a radius r around a training point x_d. Further assume that the true maximum slope $S_{x_d} := \max_{x \in B_r(x_d)} s(x, x_d)$ is achieved by a point with distance from x_d greater than $0 < \tau \leq r$. Then for $\widehat{S}_{x_d}$ estimated with an ε-covering[2] of $B_r(x_d)$,*

$$S_{x_d} \leq \widehat{S}_{x_d}(1 + \varepsilon/\tau) + L\varepsilon/\tau \, .$$

This result shows that adding a margin to the observed slope bound $\widehat{S}$ provides an upper bound on S. This margin decreases as the density of samples increases and depends on an assumption about the local smoothness of the system. This estimation problem is similar to Lipschitz estimation, which has been widely studied with an eye towards optimization (Sergeyev and Kvasov, 2010). We note that Lipschitz smoothness is a stronger condition than our bounded slope condition, and requires several additional smoothness assumptions to estimate (Wood and Zhang, 1996).

Combining Lemma 1 with Proposition 2 results in estimated bounds on the perception errors within a neighborhood of the training data. In what follows, we use these local error bounds in robust control.

2. We can achieve an ε-cover by sampling from a grid with $N = \left(\frac{r}{\varepsilon}\right)^d$ points from $B_r(x_d)$. For a controllable system, it is possible to execute a trajectory to collect these dense samples.

4. Analysis and synthesis of perception-based controllers

Robust control for generalization Recall that for a state-observation pair (x, z), the perception error, defined as $e := p(z) - Cx$, acts as additive noise to the measurement model $y = p(z)$. While standard linear control techniques can handle uniformly bounded errors, more care is necessary to further ensure that the system remains with a safe region of the state space, as determined by the training data. Through a suitable convex reformulation of the safe region, this goal can be addressed through receding horizon strategies (e.g. Wan and Kothare (2002); Mayne et al. (2006)). While these methods are effective in practice, constructing terminal sets and ensuring a priori that feasible solutions exist is not an easy task. To make explicit connections between learning and control, we turn our analysis to a system level perspective on the closed-loop to characterize its sensitivity to noise.

Once the control input to dynamical system (1) is defined to be a linear function of the measurement (4), the closed-loop behavior is determined entirely by the process noise $\boldsymbol{w}$ and the measurement noise $\boldsymbol{e}$ (as in Figure 1). Therefore, we can write the system state and input directly as a linear function of the noise

$$\begin{bmatrix} \boldsymbol{x} \\ \boldsymbol{u} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Phi_{xw}} & \boldsymbol{\Phi_{xe}} \\ \boldsymbol{\Phi_{uw}} & \boldsymbol{\Phi_{ue}} \end{bmatrix} \begin{bmatrix} H\boldsymbol{w} \\ \boldsymbol{e} \end{bmatrix} . \tag{7}$$

In what follows, we will state results in terms of these system response variables. The connection between these maps and a feedback control law $\boldsymbol{u} = \boldsymbol{K}\boldsymbol{y}$ that achieves the response (7) is formalized in the *System Level Synthesis* (SLS) framework (Wang et al., 2016). SLS states that there exists a linear feedback controller $\boldsymbol{K}$ that achieves the response $\boldsymbol{\Phi} = \{\boldsymbol{\Phi_{xw}}, \boldsymbol{\Phi_{xe}}, \boldsymbol{\Phi_{uw}}, \boldsymbol{\Phi_{ue}}\}$ for any system response $\boldsymbol{\Phi}$ constrained to lie in an affine space defined by the system dynamics. The affine space $\mathcal{A}$ defined by (A, B, C) is

$$\mathcal{A} = \left\{ \boldsymbol{\Phi} \mid \begin{bmatrix} zI - A & -B \end{bmatrix} \boldsymbol{\Phi} = \begin{bmatrix} I & 0 \end{bmatrix}, \ \boldsymbol{\Phi} \begin{bmatrix} zI - A \\ -C \end{bmatrix} = \begin{bmatrix} I \\ 0 \end{bmatrix} \right\} . \tag{8}$$

A more comprehensive introduction to SLS, including computational concerns, can be found in the full technical report (Dean et al., 2019). We finish our brief overview by remarking that the linear optimal control problem can be written as

$$\text{minimize}_{\boldsymbol{\Phi}} \quad c(\boldsymbol{\Phi}) \ \text{ subject to } \ \boldsymbol{\Phi} \in \mathcal{A} , \tag{9}$$

where the cost function c is redefined to operate on system responses; many robust control costs can be written as system norms.

In what follows, we specialize controller design concerns to our perception-based setting, and develop further conditions on the closed-loop response $\boldsymbol{\Phi}$ to incorporate into the synthesis problem.

Lemma 3 (Generalization implies closeness) *For a perception map p with errors $\boldsymbol{e} = p(\boldsymbol{z}) - C\boldsymbol{x}$, let the system responses $\{\boldsymbol{\Phi_{xw}}, \boldsymbol{\Phi_{xe}}, \boldsymbol{\Phi_{uw}}, \boldsymbol{\Phi_{ue}}\}$ lie in the affine space defined by dynamics (A, B, C), and let $\boldsymbol{K}$ be the associated controller. Then the state trajectory $\boldsymbol{x}$ achieved by the control law $\boldsymbol{u} = \boldsymbol{K}p(\boldsymbol{z})$ and driven by noise process $\boldsymbol{w}$, satisfies, for any target trajectory $\boldsymbol{x}_d$,*

$$\|\boldsymbol{x} - \boldsymbol{x}_d\| \leq \|\widehat{\boldsymbol{x}} - \boldsymbol{x}_d\| + \|\boldsymbol{\Phi_{xe}}\| \|\boldsymbol{e}\|. \tag{10}$$

where we define the nominal closeness $\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_d\| = \|\boldsymbol{\Phi_{xw}} H\boldsymbol{w} - \boldsymbol{x}_d\|$ *to be the deviation from the target trajectory in the absence of measurement errors.*

Proof Notice that over the course of a trajectory, we have system outputs $\boldsymbol{y} = p(\boldsymbol{z}) = C\boldsymbol{x} + \boldsymbol{e}$. Then recalling that the system responses are defined such that $\boldsymbol{x} = \boldsymbol{\Phi}_{\boldsymbol{xw}} H\boldsymbol{w} + \boldsymbol{\Phi}_{\boldsymbol{xe}}\boldsymbol{e}$, we have that

$$\|\boldsymbol{x} - \boldsymbol{x}_d\| = \|\boldsymbol{\Phi}_{\boldsymbol{xw}} H\boldsymbol{w} + \boldsymbol{\Phi}_{\boldsymbol{xe}}\boldsymbol{e} - \boldsymbol{x}_d\| \leq \|\boldsymbol{\Phi}_{\boldsymbol{xw}} H\boldsymbol{w} - \boldsymbol{x}_d\| + \|\boldsymbol{\Phi}_{\boldsymbol{xe}}\|\|\boldsymbol{e}\| \,.$$

■

The terms in bound (10) capture different generalization properties. The first is small if we plan to visit states during operation that are similar to those seen during training. The second term is a measure of the robustness of the closed-loop system to the error $\boldsymbol{e}$.

We are now in a position to state the main result of the paper, which shows that under an additional robustness condition, Lemmas 1 and 3 combine to define a control invariant set around the training data within which the perception errors, and consequently the performance, are bounded.

Theorem 4 *Let the assumptions of Lemmas 1 and 3 hold and, for simplicity of presentation, suppose that the training error is bounded:* $\|p(z_d) - Cx_d\| \leq R_0$ *for all* $(x_d, z_d) \in \mathcal{S}_0$. *Then as long as*

$$\|\boldsymbol{\Phi}_{\boldsymbol{xe}}\| \leq \frac{1 - \frac{1}{r}\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_d\|}{S + \frac{R_0}{r}} \,, \tag{11}$$

the perception errors remain bounded

$$\|p(\boldsymbol{z}) - C\boldsymbol{x}\| \leq \frac{\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_d\| + R_0}{1 - S\|\boldsymbol{\Phi}_{\boldsymbol{xe}}\|} := \gamma, \tag{12}$$

and the closed-loop trajectory lies within $\mathcal{X}_\gamma$.

Theorem 4 shows that the bound (11) should be used during controller synthesis to ensure generalization. Feasibility of the synthesis problem depends on the controllability and observability of the system (A, B, C), which impose limits on how small $\|\boldsymbol{\Phi}_{\boldsymbol{xe}}\|$ can be made to be, and on the planned deviation from training data as captured by the quantity $\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_d\|$.

Robust synthesis for waypoint tracking We now specialize to the task of waypoint tracking to simplify the term $\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_d\|$ and propose a robust control synthesis problem. As discussed further in the extended technical report (Dean et al., 2019), waypoint tracking can be encoded by defining the state $x_k := [\xi_k - r_k; r_k]$ as the concatenation of tracking error and waypoint and the disturbance as the change in reference, $w_k := r_{k+1} - r_k$.

Proposition 5 *For a reference tracking problem with* $\|r_{k+1} - r_k\| \leq \Delta_{\text{ref}}$ *and a reference trajectory within a ball of radius* r_{ref} *from the training data,*

$$\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_d\| \leq \Delta_{\text{ref}}\| \begin{bmatrix} I & 0 \end{bmatrix} \boldsymbol{\Phi}_{\boldsymbol{xw}} H\| + r_{\text{ref}} \,.$$

Using this setting, we add the robustness condition in (11) to a control synthesis problem as

$$\begin{array}{ll} \text{minimize}_{\boldsymbol{\Phi}} & c(\boldsymbol{\Phi}) \\ \text{subject to} & \boldsymbol{\Phi} \in \mathcal{A}, \;\; (S + \frac{R_0}{r})\|\boldsymbol{\Phi}_{\boldsymbol{xe}}\| + \frac{\Delta_{\text{ref}}}{r}\| \begin{bmatrix} I & 0 \end{bmatrix} \boldsymbol{\Phi}_{\boldsymbol{xw}} H\| + \frac{r_{\text{ref}}}{r} \leq 1 \,. \end{array}$$

Notice the trade-off between the size of different parts of the system response. Because the responses must lie in an affine space, they cannot both become arbitrarily small. Therefore, the synthesis problem must trade off between sensitivity to measurement errors and tracking fidelity. These trade-offs are mediated by the quality of the perception map, through its slope-boundedness and training error, and the ambition of the control task, through its deviation from training data and distances between waypoints.

Necessity of robustness Robust control is notoriously conservative, and our main result relies heavily on small gain-like arguments. Can the conservatism inherent in this approach be generally reduced? In this section, we answer in the negative by describing an example for which the robustness condition in Theorem 4 is necessary. For simplicity, we specialize to the goal of regulating a system to the origin, and assume that $(0, z_0)$ is in the training set, with perfect perception at that point: $p(z_0) = 0$.

Consider the double integrator

$$x_{k+1} = \begin{bmatrix} 1 & dt \\ 0 & 1 \end{bmatrix} x_k + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u_k, \quad y_k = p(z_k) = x_k + e(x_k)\,, \tag{13}$$

and the control law resulting from the $\mathcal{H}_2$ optimal control problem,

$$\min_{\mathbf{\Phi}} \left\| \begin{bmatrix} \mathbf{\Phi}_{\boldsymbol{xw}} & \mathbf{\Phi}_{\boldsymbol{xe}} \\ \mathbf{\Phi}_{\boldsymbol{uw}} & \mathbf{\Phi}_{\boldsymbol{ue}} \end{bmatrix} \right\|_{\mathcal{H}_2} \text{subject to } \mathbf{\Phi} \in \mathcal{A},\ \ \|\mathbf{\Phi}_{\boldsymbol{xe}}\| \leq \alpha\,.$$

Notice that the solution to the unconstrained problem would be the familiar LQG combination of Kalman filtering and static feedback on the estimated state.

Proposition 6 *For system* (13) *with* $dt = 0.5$*, there exist error functions* $e(x) = p(q(x)) - Cx$ *with slope globally bounded about the origin by* $S \approx 0.276$ *such that the closed-loop system is asymptotically stable if and only if* $\alpha < \frac{1}{S}$*. Thus the robustness condition* (11) *is necessary as well as sufficient.*

5. Experiments

We demonstrate our theoretical results with examples of control from pixels, using both simple synthetic images and complex graphics simulation. The synthetic example uses generated 64×64 pixel images of a moving blurry white circle on a black background; the complex example uses 800×600 pixel dashboard camera images of a vehicle in the CARLA simulator platform (Dosovitskiy et al., 2017).

For both visual settings, we demonstrate our results in controlling a 2D double integrator to track a periodic reference. The 2D double integrator state is given by $x_k^\top = [(x_k^{(1)})^\top\, (x_k^{(2)})^\top]$, and each component evolves independently according to the dynamics in (13) with $dt = 0.1$. For all examples, the sensing matrix C extracts the position of the system, i.e., $Cx_k = [x_{1,k}^{(1)}, x_{1,k}^{(2)}]$. Training and validation trajectories are generated by driving the system with an optimal state feedback controller (i.e. where measurement $\boldsymbol{y} = \boldsymbol{x}$) to track a desired reference trajectory $\boldsymbol{w} = \boldsymbol{r} + \boldsymbol{v}$, where $\boldsymbol{r}$ is a nominal reference, and $\boldsymbol{v}$ is a random norm bounded random perturbation satisfying $\|v_k\|_\infty \leq 0.1$.

We consider three different perception maps: a linear map for the synthetic example and both visual odometry and a convolutional neural net for the CARLA example. For the CNN, we collect a dense set of training examples around the reference to train the model. We use the approach proposed by Coates and Ng (2012) to learn a convolutional representation of each image: each resized and scaled image is passed through a single convolutional, ReLU activation, and max pooling layer. We then fit a linear map of these learned image features to position and heading of the camera pose. We require approximately $30,000$ training points. During operation, pixel-data z is passed through the CNN, and the position estimates y are used by the controller. We note that other more sophisticated architectures for feature extraction would also be reasonable to use in our control framework; we find that this one is conceptually simple and sufficient for tracking around our fixed reference.

To perform visual odometry, we first collect images from known poses around the desired reference trajectory. We then use ORB SLAM (Mur-Artal and Tardós, 2016) to build a global map of visual features

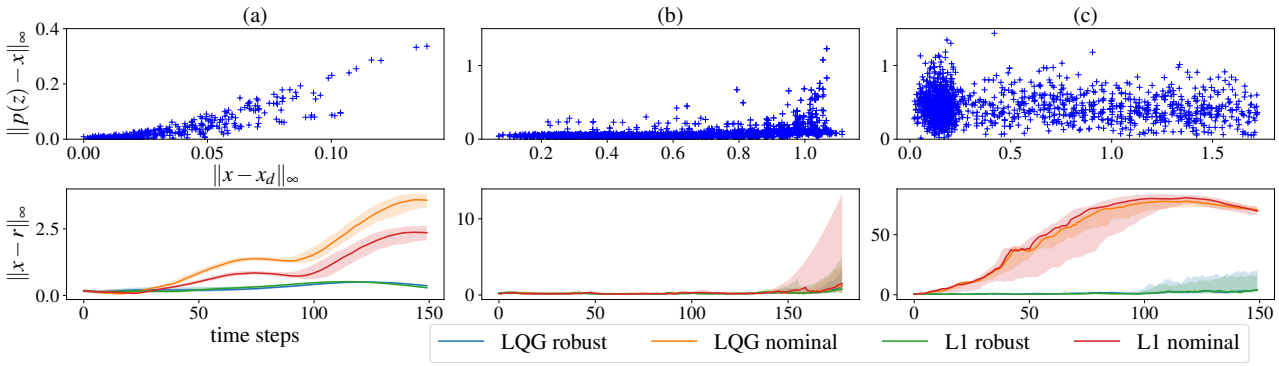

Figure 2: (top) Test perception error $\|p(z)-Cx\|_\infty$ vs. distance to the nearest training point $\|x-x_d\|_\infty$; and (bottom) median ℓ_∞ tracking error for 200 rollouts using the (a) linear map on synthetic images, (b) SLAM and (c) CNN on CARLA dashboard images. Error bars show upper and lower quartiles.

and a database of reference images with known poses. This is the "training" phase. We use one trajectory of 200 points; the reference database is approximately this size. During operation, an image z is matched with an image z_d in the database. The reprojection error between the matched features in z_d with known pose x_d and their corresponding features in z is then minimized to generate a pose estimate. For more details on standard visual odometry methods, see Scaramuzza and Fraundorfer (2011). We highlight that modern visual SLAM algorithms incorporate sophisticated filtering and optimization techniques for localization in previously unseen environments with complex dynamics; we use a simplified algorithm under this training and testing paradigm in order to better isolate the data dependence.

We then estimate safe regions for all three maps, as described in Section 3. In the top panels of Figure 2 we show the error profiles as a function of distance to the nearest training point for the linear (left), SLAM (middle), and CNN (right) maps. We see that these data-dependent localization schemes exhibit the local slope bounded property posited in the previous section.

We compare robust synthesis to the behavior of *nominal controllers* that do not take into account the nonlinearity in the measurement model. In particular, we compare the performance of naively synthesized LQG and $\mathcal{L}_1$ optimal controllers with controllers designed with the robustness condition (11). LQG is a standard control scheme that explicitly separates state estimation (Kalman Filtering) from control (LQR control), and is emblematic of much of standard control practice. $\mathcal{L}_1$ optimal control minimizes worst case state deviation and control effort by modeling process and sensor errors as ℓ_∞ bounded adversarial processes. For further details on control synthesis, refer to our full technical report (Dean et al., 2019). The bottom panels of Figure 2 show that the robustly synthesized controllers remain within a bounded neighborhood around the training data. On the other hand, the nominal controllers drive the system away from the training data, leading to a failure of the perception and control loop. We note that although visual odometry may not satisfy smoothness assumptions when the feature detection and matching fails, we nevertheless observe safe system behavior, suggesting that using our robust controller, no such failures occur.

Acknowledgments

This work is generously supported in part by ONR awards N00014-17-1-2191, N00014-17-1-2401, and N00014-18-1-2833, the DARPA Assured Autonomy (FA8750-18-C-0101) and Lagrange (W911NF-16-1-0552) programs, a Siemens Futuremakers Fellowship, and an Amazon AWS AI Research Award. SD and VY are supported by an NSF Graduate Research Fellowship under Grant No. DGE 1752814.

References

Yasin Abbasi-Yadkori and Csaba Szepesvári. Regret bounds for the adaptive control of linear quadratic systems. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 1–26, 2011.

Yasin Abbasi-Yadkori, Nevena Lazic, and Csaba Szepesvári. Model-free linear quadratic control via reduction to expert prediction. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3108–3117, 2019.

Naman Agarwal, Brian Bullins, Elad Hazan, Sham M Kakade, and Karan Singh. Online control with adversarial disturbances. *arXiv preprint arXiv:1902.08721*, 2019.

Anayo K Akametalu, Jaime F Fisac, Jeremy H Gillula, Shahab Kaynama, Melanie N Zeilinger, and Claire J Tomlin. Reachability-based safe learning with gaussian processes. In *53rd IEEE Conference on Decision and Control*, pages 1424–1431. IEEE, 2014.

Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In *Advances in neural information processing systems*, pages 908–918, 2017.

Richard Cheng, Gábor Orosz, Richard M Murray, and Joel W Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. *arXiv preprint arXiv:1903.08792*, 2019.

Adam Coates and Andrew Y Ng. Learning feature representations with k-means. In *Neural networks: Tricks of the trade*, pages 561–580. Springer, 2012.

Felipe Codevilla, Matthias Miiller, Antonio López, Vladlen Koltun, and Alexey Dosovitskiy. End-to-end driving via conditional imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–9. IEEE, 2018.

Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. On the sample complexity of the linear quadratic regulator. *arXiv preprint arXiv:1710.01688*, 2017.

Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu. Regret bounds for robust adaptive control of the linear quadratic regulator. In *Advances in Neural Information Processing Systems*, pages 4192–4201, 2018.

Sarah Dean, Nikolai Matni, Benjamin Recht, and Vickie Ye. Robust guarantees for perception-based control. *arXiv preprint arXiv:1907.03680*, 2019.

Charles A Desoer and Mathukumalli Vidyasagar. *Feedback systems: input-output properties*, volume 55. Siam, 1975.

Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. *arXiv preprint arXiv:1711.03938*, 2017.

Babak Hassibi and Thomas Kaliath. H-∞ bounds for least-squares estimators. *IEEE Transactions on Automatic Control*, 46(2):309–314, 2001.

Lukas Hewing and Melanie N Zeilinger. Cautious model predictive control using gaussian process regression. *arXiv preprint arXiv:1705.10702*, 2017.

Eagle S Jones and Stefano Soatto. Visual-inertial navigation, mapping and localization: A scalable real-time causal approach. *The International Journal of Robotics Research*, 30(4):407–430, 2011.

Sergey Levine, Chelsea Finn, Trevor Darrell, and Pieter Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.

Simon Lynen, Markus W Achtelik, Stephan Weiss, Margarita Chli, and Roland Siegwart. A robust and modular multi-sensor fusion approach applied to mav navigation. In *2013 IEEE/RSJ international conference on intelligent robots and systems*, pages 3923–3929. IEEE, 2013.

David Q Mayne, SV Raković, Rolf Findeisen, and Frank Allgöwer. Robust output feedback model predictive control of constrained linear systems. *Automatica*, 42(7):1217–1222, 2006.

Raul Mur-Artal and Juan D. Tardós. ORB-SLAM2: an open-source SLAM system for monocular, stereo and RGB-D cameras. *CoRR*, abs/1610.06475, 2016. URL http://arxiv.org/abs/1610.06475.

C. J. Ostafew, A. P. Schoellig, and T. D. Barfoot. Learning-based nonlinear model predictive control to improve vision-based mobile robot path-tracking in challenging outdoor environments. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4029–4036, May 2014. doi: 10.1109/ICRA.2014.6907444.

Dean A Pomerleau. Alvinn: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pages 305–313, 1989.

Davide Scaramuzza and Friedrich Fraundorfer. Visual odometry [tutorial]. *IEEE robotics & automation magazine*, 18(4):80–92, 2011.

Yaroslav D Sergeyev and Dmitri E Kvasov. L ipschitz global optimization. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.

Sarah Tang, Valentin Wüest, and Vijay Kumar. Aggressive flight with suspended payloads using vision-based control. *IEEE Robotics and Automation Letters*, 3(2):1152–1159, 2018.

Andrew J Taylor, Victor D Dorobantu, Hoang M Le, Yisong Yue, and Aaron D Ames. Episodic learning with control lyapunov functions for uncertain robotic systems. *arXiv preprint arXiv:1903.01577*, 2019.

Kim P Wabersich and Melanie N Zeilinger. Linear model predictive safety certification for learning-based control. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 7130–7135. IEEE, 2018.

Zhaoyang Wan and Mayuresh V Kothare. Robust output feedback model predictive control using off-line linear matrix inequalities. *Journal of Process Control*, 12(7):763–774, 2002.

Yuh-Shyang Wang, Nikolai Matni, and John C Doyle. A system level approach to controller synthesis. *arXiv preprint arXiv:1610.04815*, 2016.

Grady Williams, Paul Drews, Brian Goldfain, James M Rehg, and Evangelos A Theodorou. Information-theoretic model predictive control: Theory and applications to autonomous driving. *IEEE Transactions on Robotics*, 34(6):1603–1622, 2018.

GR Wood and BP Zhang. Estimation of the lipschitz constant of a function. *Journal of Global Optimization*, 8(1):91–103, 1996.

Sholeh Yasini and Kristiaan Pelckmans. Worst-case prediction performance analysis of the kalman filter. *IEEE Transactions on Automatic Control*, 63(6):1768–1775, 2018.