

Learning to Reach, Swim, Walk and Fly in One Trial: Data-Driven Control with Scarce Data and Side Information

Franck Djeumou

FDJEUMOU@UTEXAS.EDU

Ufuk Topcu

UTOPCU@UTEXAS.EDU

The University of Texas at Austin, United States

Editors: R. Firoozi, N. Mehr, E. Yel, R. Antonova, J. Bohg, M. Schwager, M. Kochenderfer

Abstract

We develop a learning-based control algorithm for unknown dynamical systems under very severe data limitations. Specifically, the algorithm has access to streaming and noisy data only from a single and ongoing trial. It accomplishes such performance by effectively leveraging various forms of side information on the dynamics to reduce the sample complexity. Such side information typically comes from elementary laws of physics and qualitative properties of the system. More precisely, the algorithm approximately solves an optimal control problem encoding the system's desired behavior. To this end, it constructs and iteratively refines a data-driven differential inclusion that contains the unknown vector field of the dynamics. The differential inclusion, used in an interval Taylor-based method, enables to over-approximate the set of states the system may reach. Theoretically, we establish a bound on the suboptimality of the approximate solution with respect to the optimal control with known dynamics. We show that the longer the trial or the more side information is available, the tighter the bound. Empirically, experiments in a high-fidelity F-16 aircraft simulator and MuJoCo's environments illustrate that, despite the scarcity of data, the algorithm can provide performance comparable to reinforcement learning algorithms trained over millions of environment interactions. Besides, we show that the algorithm outperforms existing techniques combining system identification and model predictive control.

Keywords: Physics-informed learning; data-driven control; system identification; reachable sets.

1. Introduction

Learning how to achieve a complex task has found numerous applications ranging from robotics (Lillicrap et al., 2016; Schulman et al., 2015; Deisenroth et al., 2013) to fluid dynamics (Kutz, 2017). However, learning algorithms generally suffer from high sample complexity, often requiring millions of samples to achieve the desired performance (Nagabandi et al., 2018; Schulman et al., 2015). Such data requirements limit the practicability of learning algorithms in real-world scenarios where an excessive number of trials cannot be performed on a physical system. A rather extreme example of such a scenario is an aircraft trying to retain a certain degree of control after abrupt changes in its dynamics, e.g., due to the loss of an engine. In such a scenario, there is a need to learn the dynamics after the abrupt changes using data from only the current trajectory.

We develop a learning-based control algorithm that utilizes data from a single trial and leverages side information on the unknown dynamics to reduce the sample complexity. The data include finitely many *noisy* samples of the states, the states' derivatives, and the control signals applied. Under such a severe limitation on the amount of available data, learning can be performed efficiently only by incorporating already known invariant properties of the dynamical system. We refer to such




Figure 1: The developed learning-based algorithm can achieve near-optimal control of simulated robots and an F-16 aircraft using streaming data obtained from the systems’ ongoing trajectory and side information derived from laws of physics. From left to right, we have the Reacher, Swimmer, Cheetah, and F-16 aircraft simulator environments.

extra knowledge as *side information*. The side information, typically derived from elementary laws of physics, may be a priori knowledge of the regularity of the dynamics, monotonicity or bounds on the vector field, algebraic constraints on the states, or knowledge of parts of the vector field.

The developed algorithm, using the data and side information available to it, computes an over-approximation of the set of states the system may reach. Then, it incorporates the over-approximation into a constrained short-horizon optimal control problem, which is solved on the fly.

Specifically, it leverages a data-driven differential inclusion to compute over-approximations of the reachable sets of the system. It first constructs a differential inclusion that contains the unknown vector field. Next, it builds on set contractor programming (Chabert and Jaulin, 2009) to refine the differential inclusion as more data become available. Then, it computes over-approximations of the reachable sets of all dynamics described by the differential inclusion through an interval Taylor-based method (Berz and Makino, 1998; Nedialkov et al., 1999) that can enforce constraints from the side information to reduce the width of the over-approximations.

The obtained over-approximations enable to formulate the data-driven optimal control problem as a nonconvex and uncertain optimization problem. Specifically, we encode the control task as the sequential optimization of a cost function over a time horizon. Even for convex cost functions, the control problem is typically nonconvex. Besides, the predictions of the states’ values at future times cannot be computed due to the unknown dynamics. The developed algorithm leverages the obtained over-approximations to optimize the nonconvex problem under the uncertain states’ predictions.

The algorithm computes approximate solutions to the nonconvex optimization problem through convex relaxations. We develop a sequential convex optimization scheme (Mao et al., 2019) that uses the obtained over-approximation and iteratively linearizes its nonconvex constraint around the previous iteration solution. Thus, each iteration solves a *convex* optimization problem, and we leverage trust regions to account for the potential errors due to the linearization.

Theoretically, we establish a bound on the suboptimality of the approximate solution with respect to the optimal control solution in the case where the dynamics were known. The bound is proportional to the width of the obtained over-approximations. We show that the longer the trial or the more the side information available, the tighter the over-approximations. Thus, the algorithm achieves near-optimal control as more data streams or more side information is available.

Empirically, our preliminary results show that the proposed approach can provide performance comparable to reinforcement learning (RL) algorithms, such as D4PG (Barth-Maron et al., 2018) and SAC (Haarnoja et al., 2018), while outperforming a popular system identification technique, SINDYC, with model predictive control (Brunton et al., 2016; Kaiser et al., 2018). It is worth emphasizing that it is unfair to compare the proposed approach with RL algorithms—yet we still do it to exemplify the potential strengths—because the proposed approach relies on data only from a

single trajectory, whereas RL algorithms typically rely on excessively large interactions with the environment. Indeed, in several control tasks from MuJoCo (Todorov et al., 2012; Tassa et al., 2018), D4PG and SAC were trained over millions of environment interactions before comparing to the proposed approach. Besides, in a ground collision avoidance scenario of an F-16 aircraft (Heidlauf et al., 2018), we show that SINDYc is unable to avoid the ground collision while being at least two orders of magnitude computationally slower than the proposed approach.

Related Work. In our prior work (Djeumou et al., 2021, 2020), we described a data-driven algorithm similar to the algorithm developed in this paper. However, the algorithm (Djeumou et al., 2021, 2020) works only for control-affine dynamics. Further, most of the considered side information is not tailored for robotics systems, and only one-step optimal control problems were investigated. In contrast, the algorithm in this paper is applicable for a more general class of dynamics with polynomial dependency in control. We also evaluate the developed algorithm on highly-complex systems and consider a larger set of side information, e.g., algebraic constraints on states and unknown terms. Besides, we investigate short-horizon rather than one-step optimal control problems.

Several approaches for data-driven control combine model predictive control with system identification or data-driven reachable set estimation. These approaches achieve system identification through sparse regression over a library of nonlinear functions (Kaiser et al., 2018), regression over the set of polynomials of fixed degree with physics-based side information (Ahmadi and Khadir, 2020), spectral properties of the collected data (Proctor et al., 2016), Koopman theory (Korda and Mezić, 2018), or Gaussian processes (Krause and Ong, 2011; Gahlawat et al., 2020). The approaches (Devonport and Arca, 2020; Haesaert et al., 2017; Chakrabarty et al., 2018) achieve data-driven estimation of the reachable sets of partially unknown dynamics using either supervised learning or Gaussian processes. They provide only probabilistic guarantees of the correctness of the computed reachable sets while our algorithm computes correct over-approximations. Recent work (Berberich et al., 2020a,b; Markovsky and Dörfler, 2021; van Waarde et al., 2020; Van Waarde et al., 2020) and DeePC (Coulson et al., 2019) have proposed data-driven control techniques based on the behavioral systems theory foundation (Willems et al., 2004), which bypass the system identification step. These techniques mostly assume linear time-invariant dynamical systems and are extremely performant in such a setting. Except for Ahmadi and Khadir (2020) that consider limited side information and build on computationally expensive semidefinite programs solvers, none of the above approaches (in their current form) can exploit the side information in this paper. Besides, through extensive comparisons with SINDYc, DeePC, and Gaussian-based approaches, Djeumou et al. (2021, 2020) empirically demonstrate that: (a) Even for a simple system such as a unicycle, these techniques achieve significant lower performance (computation time and control suboptimality) than an approach that can exploit side information; (b) These techniques struggle to learn on high-dimensional and complex systems. Thus, this paper compares against RL techniques even though they work in a drastically different data regime and under different assumptions.

Model-free (Mnih et al., 2015; Oh et al., 2016; Lillicrap et al., 2016; Mnih et al., 2016; Schulman et al., 2016) and model-based (Nagabandi et al., 2018; Deisenroth and Rasmussen, 2011; Gu et al., 2016; Boedecker et al., 2014; Levine and Abbeel, 2014) RL algorithms have been widely used for data-driven control. Model-free algorithms can achieve high performance at the expense of high sample complexity (Schulman et al., 2016) while model-based algorithms are more data-efficient but generally achieve lower performance than model-free approaches. In contrast, our algorithm can work with data from only the current trajectory by exploiting side information on the dynamics.

2. Background

Notation. We denote an interval by $[a, b] = \{x \in \mathbb{R} | a \leq x \leq b\}$ for some $a, b \in \mathbb{R}$ such that $a \leq b$, the set $\{i, \dots, j\}$ by $\mathbb{N}_{[i,j]}$ for $i, j \in \mathbb{N}$ with $i \leq j$, the k^{th} component of a vector x and the (k, j) component of a matrix X by x_k and $X_{k,j}$, respectively, the weighted norm of a vector $x \in \mathbb{R}^n$ by $\|x\|_w = \sqrt{\sum_{i=1}^n (w_i x_i)^2}$ for some $w \in \mathbb{R}_+^n$, and the Lipschitz constant of $f : \mathcal{X} \rightarrow \mathbb{R}$ by $L_f^w = \sup\{L \in \mathbb{R} \mid |f(x) - f(y)| \leq L\|x - y\|_w, x, y \in \mathcal{X}, x \neq y\}$ for $\mathcal{X} \subseteq \mathbb{R}^n$.

Interval Analysis. We denote the set of intervals on \mathbb{R} by $\mathbb{IR} = \{\mathcal{A} = [\underline{\mathcal{A}}, \overline{\mathcal{A}}] \mid \underline{\mathcal{A}}, \overline{\mathcal{A}} \in \mathbb{R}, \underline{\mathcal{A}} \leq \overline{\mathcal{A}}\}$, the set of n -dimensional interval vectors by \mathbb{IR}^n , and the set of $n \times m$ -dimensional interval matrices by $\mathbb{IR}^{n \times m}$. We carry forward the definitions (Moore, 1966) of arithmetic operations, set inclusion, and intersections of intervals to interval vectors and matrices by applying them componentwise. We use the term interval to specify an interval vector or interval matrix when it is clear from the context. Given $f : \mathcal{X} \mapsto \mathcal{Y}$ with $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$, we define an *interval extension of f* as $\mathbf{f} : \mathbb{IR}^n \mapsto \mathbb{IR}^m$ satisfying $\mathbf{f}(\mathcal{A}) \supseteq \mathcal{R}(f, \mathcal{A}) = \{f(x) \mid x \in \mathcal{A}\}$, $\forall \mathcal{A} \subseteq \mathcal{X}$. Thus, given an interval \mathcal{A} , $\mathbf{f}(\mathcal{A})$ is an interval that *over-approximates* the range of values taken by f over \mathcal{A} .

Interval-Based Contractor. Interval-based contractor programming is a mathematical framework to solve constraints involving interval variables. Given an initial over-estimation of the constraint's solutions, a contractor filters such variable domains, i.e., reduces the interval of each variable, without loss of solutions of the constraints. Consider the constraint $h(\cdot) \leq 0$. Assume that $\mathcal{A} = [\mathcal{A}_1, \dots, \mathcal{A}_n] \in \mathbb{IR}^n$ is a set containing the solutions. Then, the contractor operator computes $C_{\mathcal{A}}^h = [C_{\mathcal{A}_1}^h, \dots, C_{\mathcal{A}_n}^h] \in \mathbb{IR}^n$ such that $C_{\mathcal{A}_i}^h \subseteq \mathcal{A}_i, \forall i \in \mathbb{N}_{[1,n]}$ and $h(x) > 0$ for all $x \in \mathcal{A} \setminus C_{\mathcal{A}}^h$.

Several polynomial-time algorithms (Benhamou et al., 1999; Van Hentenryck et al., 1997; Trombettoni et al., 2010) have been developed to compute contractors associated with a given constraint. For example, HC4-Revise (Benhamou et al., 1999) is a linear-time algorithm that provides optimal contractors when each variable appears only once in the constraint. In the following, we use $C_{\mathcal{A}}^h$ to refer to the contracted interval resulting from any of these algorithms.

3. Problem Formulation

This paper considers nonlinear dynamics with polynomial dependency in the control inputs as

$$\dot{x} = f(x) + \sum_{p=1}^d g_p(x) u[\alpha^p], \quad (1)$$

where $d \in \mathbb{N}$, $\alpha^p \in \mathbb{N}^m$ is known, the state $x : \mathbb{R}_+ \mapsto \mathcal{X}$ is a continuous-time signal evolving in $\mathcal{X} \in \mathbb{IR}^n$, $u[\alpha^p] = u_1^{\alpha_1^p} \dots u_m^{\alpha_m^p}$ is a monomial with variables from the control signal $u : \mathbb{R}^+ \mapsto \mathcal{U}$ where $\mathcal{U} \in \mathbb{IR}^m$. The vector-valued functions $f = [f_k] : \mathbb{R}^n \mapsto \mathbb{R}^n$ and $g_p = [g_{p,k}] : \mathbb{R}^n \mapsto \mathbb{R}^n$ are considered to be *nonlinear and unknown*. Note that even if the dynamics are not in the class above, Taylor expansion provides a tight approximation of the dynamics that lies in such a class.

Assumption 1 (LIPSCHITZ SYSTEMS) *Given a set $\mathcal{A} \subseteq \mathbb{R}^n$, f_k and $g_{p,k}$ admit local Lipschitz constants $L_{f_k}^w, L_{g_{p,k}}^w > 0$ on \mathcal{A} , for some $w \in \mathbb{R}_+^n$ and for all $k \in \mathbb{N}_{[1,n]}, p \in \mathbb{N}_{[1,d]}$.*

Assumption 1 is common in the framework of optimal control. We emphasize that even though we use the weighted norm to define the Lipschitz constants, the results of this paper can be straightforwardly extended to general modulus of continuity assumption on f and g_p . The weighted norm has the advantage of providing information on the relative importance of each variable in the function.

Besides, the domain $\mathcal{X} \in \mathbb{R}^n$ is bounded. Thus, by Assumption 1, f_k and $g_{p,k}$ admit global Lipschitz constants on \mathcal{X} . We exploit such a knowledge by assuming *known upper bounds* on the Lipschitz constants. That is, we have access to $\bar{f}_k \in \mathbb{R}_+$ and $\bar{g}_{p,k} \in \mathbb{R}_+$ as known upper bounds on the Lipschitz constants $L_{f_k}^w$ and $L_{g_{p,k}}^w$, respectively, for $k \in \mathbb{N}_{[1,n]}$ and $p \in \mathbb{N}_{[1,d]}$. We emphasize that the Lipschitz bounds can be directly estimated from data at the expense of weakening some of the guarantees in this paper. Our numerical experiments use Lipschitz bounds estimated from data.

In a discrete-time setting, we denote the initial time by $t_1 \geq 0$ and the current time by $t_j > t_1$ for some $j > 1$. Let $\mathcal{T}_j = \{(\tilde{x}^i, \dot{\tilde{x}}^i, u^i)\}_{i=1}^{j-1}$ be the finite-length set of observations obtained between t_1 and t_j . The dataset \mathcal{T}_j contains $j - 1$ noisy samples of the exact state $x^i = x(t_i)$, the derivative $\dot{x}^i = \dot{x}(t_i)$ of the state, and the applied input $u^i = u(t_i)$. We build on the widely-used *bounded* noise assumption and consider that $|x(t) - \tilde{x}(t)| \leq \eta$, $|\dot{x}(t) - \dot{\tilde{x}}(t)| \leq \bar{\eta}$ for all $t \in \mathbb{R}_+$ and for some vector values $\eta, \bar{\eta} \in \mathbb{R}_+^n$. Here, the absolute value and the comparison are conducted elementwise.

We seek to control the unknown dynamical system (1) by finding $u^j, \dots, u^{j+N} \in \mathcal{U}$ that are solutions of the *N-step optimal control problem*

$$\underset{u^j, \dots, u^{j+N} \in \mathcal{U}}{\text{minimize}} \quad \sum_{q=j}^{j+N} c(x^q, u^q, x^{q+1} = x(t_{q+1}; x^q, u^q)), \quad (2)$$

where N is the planning horizon, c is a known cost function, $x^j = x(t_j)$ is the known current state of the system, $t_q = t_j + (q - j)\Delta t$, Δt is a constant time step, and $x^{q+1} = x(t_{q+1}; x^q, u^q)$ is the state at t_{q+1} , i.e., a solution of the differential equation (1) at t_{q+1} when x^q is the initial state and u^q is the constant control applied between $[t_q, t_{q+1}]$. The optimization problem (2) is generally nonconvex since the state at t_{q+1} is nonconvex due to the nonlinear dynamics. Besides, x^{q+1} cannot be computed due to the unknown dynamics. Thus, we seek for approximate solutions to (2).

Problem 1 *Given the dataset \mathcal{T}_j , the current state \tilde{x}^j , compute an approximate solution to the N-step optimal control problem (2) and characterize the suboptimality of such approximation.*

4. Reachable Set Over-Approximation via Data-Based Differential Inclusions

In this section, we first construct a differential inclusion $\dot{x} \in \mathbf{f}(x) + \sum_{p=1}^d \mathbf{g}_p(x)u[\alpha^p]$ that contains the unknown vector field. Then, we adapt an interval Taylor-based method to over-approximate the reachable set of dynamics described by the constructed differential inclusion. Finally, we show how additional side information constrains the Taylor expansion to provide tighter over-approximations.

Lemma 1 (OVER-APPROXIMATION OF f AND g_p) *Let the set $\mathcal{E}_j = \{(\tilde{x}^i, C_{\mathcal{F}^i}, C_{\mathcal{G}^i})\}_{i=0}^{j-1}$ be such that $C_{\mathcal{F}^i} = [C_{\mathcal{F}^i}] \in \mathbb{R}^n$ and $C_{\mathcal{G}^i} = [C_{\mathcal{G}^i}] \in \mathbb{R}^{d \times n}$ satisfy $f_k(\tilde{x}^i) \in C_{\mathcal{F}^i}$ and $g_{p,k}(\tilde{x}^i) \in C_{\mathcal{G}^i}$ for all $p \in \mathbb{N}_{[1,d]}$ and $k \in \mathbb{N}_{[1,n]}$. Then, the interval-valued functions $\mathbf{f} = [\mathbf{f}_k] : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\mathbf{g}_p = [\mathbf{g}_{p,k}] : \mathbb{R}^n \rightarrow \mathbb{R}^n$, defined by $\mathbf{f}_k(\mathcal{A}) = \bigcap_{(\tilde{x}^i, C_{\mathcal{F}^i}, \cdot) \in \mathcal{E}_j} C_{\mathcal{F}^i} + [-1, 1]\bar{f}_k \boldsymbol{\eta}^w(\mathcal{A} - \tilde{x}^i)$ and $\mathbf{g}_{p,k}(\mathcal{A}) = \bigcap_{(\tilde{x}^i, \cdot, C_{\mathcal{G}^i}) \in \mathcal{E}_j} C_{\mathcal{G}^i} + [-1, 1]\bar{g}_{p,k} \boldsymbol{\eta}^w(\mathcal{A} - \tilde{x}^i)$, are such that $\mathcal{R}(f_k, \mathcal{A}) \subseteq \mathbf{f}_k(\mathcal{A})$ and $\mathcal{R}(g_{p,k}, \mathcal{A}) \subseteq \mathbf{g}_{p,k}(\mathcal{A})$ for all $\mathcal{A} \subseteq \mathcal{X}$. Furthermore, the function $\boldsymbol{\eta}^w : \mathbb{R}^n \mapsto \mathbb{R}$ can be any tight interval extension of the weighted norm $\|\cdot\|_w$.*

We provide a proof of the lemma and an expression for $\boldsymbol{\eta}^w$ in the extended version of the paper (Djeumou and Topcu, 2022). Intuitively, Lemma 1 states that if f and g_p can be over-approximated at a finite set of data points specified by \mathcal{E}_j , then it is possible to obtain a formula to over-approximate f and g_p on the entire state space \mathcal{X} via the Lipschitz bounds. Lemma 2 enables to construct and refine the bounds of the set \mathcal{E}_j based on the data \mathcal{T}_j and the noise bounds.

Lemma 2 (REFINEMENT VIA CONTRACTOR) *Given a data point $(\tilde{x}^i, \tilde{x}^i, u^i) \in \mathcal{T}_j$, an interval $\mathcal{F}^i = [\mathcal{F}_k^i] \in \mathbb{I}\mathbb{R}^n$ such that $f(\tilde{x}^i) \in \mathcal{F}^i$, and an interval $\mathcal{G}^i = [\mathcal{G}_{p,k}^i] \in \mathbb{I}\mathbb{R}^{d \times n}$ such that $g_{p,k}(\tilde{x}^i) \in \mathcal{G}_{p,k}^i$ for all $p \in \mathbb{N}_{[1,d]}$, $k \in \mathbb{N}_{[1,n]}$. Let the intervals $C_{\mathcal{F}^i} \in \mathbb{I}\mathbb{R}^n$ and $C_{\mathcal{G}^i} \in \mathbb{I}\mathbb{R}^{d \times n}$ defined by*

$$C_{\mathcal{F}_k^i} = \mathcal{F}_k^i \cap \left\{ \tilde{\mathcal{N}}_k^i - \sum_{p=1}^d \mathcal{G}_{p,k}^i u^i[\alpha^p] \right\}, C_{\mathcal{G}_{p,k}^i} = \begin{cases} \left(\left\{ \mathcal{S}_{p-1,k} - \sum_{l=p+1}^d \mathcal{G}_{l,k}^i u^i[\alpha^l] \right\} \cap \mathcal{G}_{p,k}^i u^i[\alpha^p] \right)^{\frac{1}{u^i[\alpha^p]}}, \\ \text{if } u^i[\alpha^p] \neq 0, \\ \mathcal{G}_{p,k}^i, \quad \text{otherwise,} \end{cases}$$

$$\mathcal{S}_{0,k} = \left\{ \tilde{\mathcal{N}}_k^i - C_{\mathcal{F}_k^i} \right\} \cap \left\{ \sum_{p=1}^d \mathcal{G}_{p,k}^i u^i[\alpha^p] \right\}, \mathcal{S}_{p,k} = \left\{ \mathcal{S}_{p-1,k} - C_{\mathcal{G}_{p,k}^i} u^i[\alpha^p] \right\} \cap \left\{ \sum_{l=p+1}^d \mathcal{G}_{l,k}^i u^i[\alpha^l] \right\},$$

for successive values of $k \in \mathbb{N}_{[1,n]}$ and for all $p \in \mathbb{N}_{[1,d]}$ with $\tilde{\mathcal{N}}_k^i = [\tilde{x}_i - \bar{\eta}, \tilde{x}_i + \bar{\eta}]$. Then, $C_{\mathcal{F}^i}$ and $C_{\mathcal{G}^i}$ are the smallest intervals enclosing $f(\tilde{x}^i)$ and $g_p(\tilde{x}^i)$, given only the data $(\tilde{x}^i, \tilde{x}^i, u^i)$, \mathcal{F}^i , \mathcal{G}^i .

Algorithm 1 Construct: Compute \mathcal{E}_j required to over-approximate f and g_p at each data point of a given trajectory.

Input: Dataset \mathcal{T}_j and a parameter $M > 0$.

Output: $\mathcal{E}_j = \{(\tilde{x}^i, C_{\mathcal{F}^i}, C_{\mathcal{G}^i})\}_{i=0}^{j-1}$.

- 1: $\mathcal{A} \leftarrow \mathcal{X}$, $\mathcal{R}^{f\mathcal{A}}, \mathcal{R}^{G\mathcal{A}} \leftarrow [-M, M]^n$
- 2: Define $\tilde{x}^0 \in \mathcal{A}$, $C_{\mathcal{F}^0} \leftarrow \mathcal{R}^{f\mathcal{A}}$, $C_{\mathcal{G}^0} \leftarrow \mathcal{R}^{G\mathcal{A}}$
- 3: **for** $i \in \mathbb{N}_{[1,j]} \wedge (\tilde{x}^i, \tilde{x}^i, u^i) \in \mathcal{T}_j$ **do**
- 4: $\mathcal{E}_i \leftarrow \text{Refine}((\tilde{x}^i, \tilde{x}^i, u^i), \mathcal{E}_{i-1}, \mathcal{T}_i)$
- 5: **end for**
- 6: **return** \mathcal{E}_j

Algorithm 2 Refine: Update \mathcal{E}_j with data.

Input: A point $(\tilde{x}^j, \tilde{x}^j, u^j)$, \mathcal{E}_j containing past over-approximations, \mathcal{T}_j , and the noise bound $\bar{\eta}$.

Output: $\mathcal{E}_{j+1} = \{(\tilde{x}^i, C_{\mathcal{F}^i}, C_{\mathcal{G}^i})\}_{i=0}^j$.

- 1: Compute $\mathcal{F}^j = \mathbf{f}(\tilde{x}^j)$, $\mathcal{G}^j = [g_{p,k}(\tilde{x}^j)]$ via Lemma 1 and \mathcal{E}_j
- 2: Compute $C_{\mathcal{F}^j}, C_{\mathcal{G}^j}$ via Lemma 2
- 3: **for** $(\tilde{x}^i, \tilde{x}^i, u^i) \in \mathcal{T}_{j+1}$ **do**
- 4: Execute 1–2 with $j = i$, $\mathcal{E}_j = \mathcal{E}_{j+1}$
- 5: **end for**
- 6: **return** \mathcal{E}_{j+1}

The proof of the lemma is provided in the extended version of the paper (Djeumou and Topcu, 2022). Lemma 2 provides tighter sets $C_{\mathcal{F}^i} \subseteq \mathcal{F}^i$ and $C_{\mathcal{G}^i} \subseteq \mathcal{G}^i$ that prune out from \mathcal{F}^i and \mathcal{G}^i some values $f(\tilde{x}^i)$ and $g_p(\tilde{x}^i)$ that do not satisfy the dynamics constraint $\dot{x}^i = f(x^i) + \sum_{p=1}^d g_p(x^i)u^i[\alpha^p]$. Specifically, one can observe that Lemma 2 formalizes the output of HC4-Revise, when contracting \mathcal{F}^i and \mathcal{G}^i under the constraints of the dynamics. Next, we develop Algorithm 1 that utilizes \mathcal{T}_j to construct and refine the set \mathcal{E}_j required in Lemma 1 to over-approximate f and g_p .

Theorem 1 (DATA-DRIVEN DIFFERENTIAL INCLUSION) *Given a dataset \mathcal{T}_j , the bounds \bar{f}_k and $\bar{g}_{p,k}$, it holds that the unknown vector field of the dynamics (1) satisfies*

$$\dot{x} \in \mathbf{h}(x, u) \triangleq \mathbf{f}(x) + \sum_{p=1}^d \mathbf{g}_p(x)u[\alpha^p], \quad (3)$$

where \mathbf{f} and \mathbf{g}_p are obtained from Lemma 1 with \mathcal{E}_j taken as the output of Algorithm 1.

Remark 1 (PERSISTENT EXCITATION) *The tightness of the differential inclusion (3) depends on how much information on f and g_p can be obtained from \mathcal{T}_j . This is the classical observability problem, sometimes referred to as persistent excitation (Willems et al., 2004). Thus, the learning algorithm should sometimes take suboptimal actions through persistent excitations of the system.*

Finally, we compute an over-approximation of the reachable sets of all dynamics described by the differential inclusion (3). Theorem 2 provides a closed-form expression for such a set.

Theorem 2 (DATA-DRIVEN REACHABLE SET OVER-APPROXIMATION) *Given the dataset \mathcal{T}_j , a constant control signal $u : t \mapsto u^q$ on the interval $[t_q, t_{q+1}]$ with $u^q \in \mathcal{U}$, and the uncertain set $\mathcal{R}^q \in \mathbb{I}\mathbb{R}^n$ of states x^q at time t_q , a closed-form expression for $\mathcal{R}^{q+1} \supseteq \{x(t_{q+1}; u^q, x^q) \in \mathcal{X} | x^q \in \mathcal{R}^q\}$, which over-approximates the reachable set at t_{q+1} for all $x^q \in \mathcal{R}^q$, is given by*

$$\mathcal{R}^{q+1} = \mathcal{R}^q + \mathbf{h}(\mathcal{R}^q, u^q)\Delta t + \left(\mathcal{J}^f + \sum_{p=1}^d \mathcal{J}^{g_p} u^q[\alpha^p]\right) \mathbf{h}(\mathcal{P}^q, u^q) \Delta t^2 / 2, \quad (4)$$

where the matrices $\mathcal{J}^f = [\mathcal{J}_{k,l}^f] \in \mathbb{I}\mathbb{R}^{n \times n}$ and $\mathcal{J}^{g_p} = [\mathcal{J}_{k,l}^{g_p}] \in \mathbb{I}\mathbb{R}^{n \times n}$, over-approximations of the Jacobian of f and g_p , are such that $\mathcal{J}_{k,l}^f = [-1, 1] w_k \bar{f}_k$ and $\mathcal{J}_{k,l}^{g_p} = [-1, 1] w_k \bar{g}_{p,k}$ for all $p \in \mathbb{N}_{[1,d]}$ and $k, l \in \mathbb{N}_{[1,n]}$. Further, the set \mathcal{P}^q , a rough enclosure of $\{x(t_{q+1}; u^q, x^q) \in \mathcal{X} | x^q \in \mathcal{R}^q\}$, is a solution of the fixpoint equation $\mathcal{R}^q + [0, \Delta t] \mathbf{h}(\mathcal{P}^q, u^q) \subseteq \mathcal{P}^q$.

We provide the proof of Theorem 2 in the extended paper (Djeumou and Topcu, 2022). It merges the differential inclusion with an interval Taylor-based expansion of order 2 to obtain the result. We note that heuristic algorithms (Nedialkov et al., 1999) can be used to solve the fixpoint equation.

Theorem 2, as it is, does not incorporate side information other the regularity assumption. We describe in the following how to incorporate a-priori knowledge to tighten \mathcal{R}^{q+1} given by (4).

Side information 1 (PARTIAL DYNAMICS KNOWLEDGE) *The vector field of (1) contains both known and unknown terms. That is, $\dot{x} = \sum_{s=1}^S f^s(x) \cdot f^s(x) + \sum_{p=1}^d \sum_{s=1}^S g_p^s(x) \cdot g_p^s(x) u[\alpha^p]$, where \cdot denotes the elementwise product between vectors, f^s and $g_p^s(x)$ are known differentiable functions, and f^s, g_p^s are unknown Lipschitz functions satisfying Assumption 1.*

Given \mathcal{E}_j^s containing past over-approximations of f^s, g_p^s and a new data point $(\tilde{x}^j, \tilde{x}^j, u^j)$, the refinement (Algorithm 2) is adapted to compute in line 1 over-approximations $\mathbf{f}^s(\tilde{x}^j)$ and $\mathbf{g}_p^s(\tilde{x}^j)$ via Lemma 1 and \mathcal{E}_j^s . Then, line 2 is modified such that each $\mathbf{f}^s(\tilde{x}^j)$ and $\mathbf{g}_p^s(\tilde{x}^j)$ are contracted according to the new dynamics' constraint $\dot{x}^j = \sum_{s=1}^S f^s(x^j) \cdot f^s(x^j) + \sum_{p=1}^d \sum_{s=1}^S g_p^s(x^j) \cdot g_p^s(x^j) u^j[\alpha^p]$. The contracted sets can be obtained straightforwardly by slight changes in the scheme described by Lemma 2 or by calling an algorithm such as `HC4-Revise`. Thus, the new differential inclusion (3) is given by $\mathbf{h}(x, u) = \sum_{s=1}^S \mathbf{f}^s(x) \cdot \mathbf{f}^s(x) + \sum_{p=1}^d \sum_{s=1}^S \mathbf{g}_p^s(x) \cdot \mathbf{g}_p^s(x) u[\alpha^p]$, where \mathbf{f}^s and \mathbf{g}_p^s are interval extensions of known f^s and g_p^s . Furthermore, we compute the new Jacobian terms \mathcal{J}^f and \mathcal{J}^{g_p} used in \mathcal{R}^{q+1} by applying chain rules and exploiting the Lipschitz bounds.

Side information 2 (ALGEBRAIC CONSTRAINTS) *We are given a constraint $r(\dot{x}(\cdot), x(\cdot)) \geq 0$ where r is a differentiable function. Such a constraint may derive from conservation laws of physics.*

This side information provides tighter over-approximations of $\mathbf{f}, \mathbf{g}_p, \mathcal{J}^f$, and \mathcal{J}^{g_p} locally. Thus, it enables to obtain a tighter \mathcal{R}^{q+1} . Specifically, this constraint can be formulated as the new constraint $w(f(x), [g_{p,k}(x)], u, x) \geq 0$ by substitution of \dot{x} with the right hand side of (1). Under equality constraints, another constraint $z(f(x), [g_{p,k}(x)], \frac{\partial f}{\partial x}(x), [\frac{\partial g_{p,k}}{\partial x}(x)], u, x) = 0$ can be derived by differentiating w with respect to x . The new constraints w and z can be incorporated in the computation of \mathcal{R}^{q+1} through contractor programming. More specifically, Algorithm 2 is adapted in line 2 to further contract $C_{\mathcal{F}^j}, C_{\mathcal{G}^j}$ with respect to the constraints w and z . Similarly, the constraint z can be used to contract the interval extensions of the Jacobian \mathcal{J}^f and \mathcal{J}^{g_p} .

5. Approximate Optimal Control

In this section, we develop an algorithm that computes approximate solutions to the optimal control problem (2) using over-approximations of the reachable sets. Further, we characterize the suboptimality of the approximate solutions with respect to the case of known dynamics.

The nonconvexity in (2) is due to the nonconvexity of $x^{q+1} = x(t_{q+1}; x^q, u^q)$ and the possibly nonconvex cost c . We replace x^{q+1} with $x^{q+1} = \hat{h}^\theta(x^q, u^q) \in \mathcal{R}^{q+1}$, where the function \hat{h}^θ , parameterized with $\theta \in \mathbb{R}^n$, is a trajectory picked inside \mathcal{R}^{q+1} . For example, a straightforward choice can be $\hat{h}^\theta(x^q, u^q) = \theta \overline{\mathcal{R}^{q+1}} + (1 - \theta) \underline{\mathcal{R}^{q+1}}$, for $\theta \in [0, 1]^n$ or θ can also be optimistically optimized to minimize the cost. Then, we solve the nonconvex problem by sequentially linearizing x^{q+1} and the cost c around the solution of the s^{th} iteration. This results into a *convex subproblem* that is solved to full optimality. The obtained solutions are then used at the $(s + 1)^{\text{th}}$ iteration.

Linearization. Let $\mathbf{x} = [x^{j+1}; \dots; x^{j+N+1}] \in \mathbb{R}^{nN}$ and $\mathbf{u} = [u^j; \dots; u^{j+N}] \in \mathbb{R}^{mN}$. We denote the solutions of the s^{th} iteration by $\mathbf{x}^s = [x^{j+1,s}; \dots; x^{j+N+1,s}]$ and $\mathbf{u}^s = [u^{j,s}; \dots; u^{j+N,s}]$. Then, we can approximate the gradient of h^θ (or $x(t_{q+1}; x^q, u^q)$) around $\mathbf{x}^s, \mathbf{u}^s$ as follows:

$$\begin{aligned} A^{q,s} &= \left. \frac{\partial h^\theta(x^q, u^q)}{\partial x^q} \right|_{x^{q,s}, u^{q,s}} \in \mathbb{I} + (\mathcal{J}^f(x^{q,s}) + \sum_{p=1}^d \mathcal{J}^{g_p}(x^{q,s}) u^{q,s}[\alpha^p]) \Delta t = \mathcal{A}^{q,s}, \\ B^{q,s} &= \left. \frac{\partial h^\theta(x^q, u^q)}{\partial u^q} \right|_{x^{q,s}, u^{q,s}} \in \sum_{p=1}^d \left[\Delta t g_{p,k}(x^{q,s}) \frac{\partial u[\alpha^p]}{\partial u_l} \Big|_{u^{q,s}} \right]_{k,l} = \mathcal{B}^{q,s}, \end{aligned}$$

where \mathbb{I} is the identity matrix of appropriate dimensions. The jacobian $\mathcal{J}^f(x^{q,s}), \mathcal{J}^{g_p}(x^{q,s})$ are exactly \mathcal{J}^f and \mathcal{J}^{g_p} when no extra side information is given. With side information, the matrices are computed through chain rules as described in side information 1. Note that since we neglect the term in Δt^2 , $\mathcal{A}^{q,s}$ and $\mathcal{B}^{q,s}$ are approximations of the actual range of the gradients of h^θ for all θ .

Next, we define the variables $\Delta \mathbf{x} = \mathbf{x} - \mathbf{x}^s$, $\Delta x^q = x^q - x^{q,s}$, $\Delta \mathbf{u} = \mathbf{u} - \mathbf{u}^s$, and $\Delta u^q = u^q - u^{q,s}$ in terms of the unknown solutions of the current iteration \mathbf{x} and \mathbf{u} . Thus, at the $(s + 1)^{\text{th}}$ iteration, the first-order approximation of $x^{q+1} = \hat{h}^\theta(x^q, u^q)$ around the previous solution $(x^{q,s}, u^{q,s})$ is

$$x^{q+1,s} + \Delta x^{q+1} \in \hat{h}^\theta(x^{q,s}, u^{q,s}) + \mathcal{A}^{q,s} \Delta x^q + \mathcal{B}^{q,s} \Delta u^q + v^q, \quad (5)$$

where $\mathbf{v} = [v^j; \dots; v^{j+N}]$ are penalty variables that enable the linearization to be always feasible. Further, to ensure that the variable v^q is used only when necessary, we augment the cost function with the sufficiently large penalization weight $\lambda > 0$. Thus, the solution for the $(s + 1)^{\text{th}}$ iteration, optimizes the penalized and linearized cost given by $L^s(\Delta \mathbf{x}, \Delta \mathbf{u}) = \sum_{q=j}^{j+N} (c(x^{q,s}, u^{q,s}, x^{q+1,s}) + \nabla c(x^{q,s}, u^{q,s}, x^{q+1,s})[\Delta \mathbf{x}; \Delta \mathbf{u}]) + \lambda \sum_{q=j}^{j+N} \|v^q\|$, where we also linearize the possibly nonconvex function c given that ∇c is its gradient, and $\|\cdot\|$ can be either the infinity norm or 1-norm. In order to verify the linearization accuracy, we also define the nonlinear *realized* cost $J(\mathbf{x}, \mathbf{u}) = \sum_{q=j}^{j+N} c(x^q, u^q, x^{q+1}) + \lambda \sum_{q=j}^{j+N} \|x^{q+1} - \hat{h}^\theta(x^q, u^q)\|$.

Trust region constraints and linearized problem. We impose the trust region constraint $\|\Delta \mathbf{u}\| \leq r^s$ to ensure that \mathbf{u} does not deviate significantly from the control input \mathbf{u}^s obtained in the previous iteration, where r^s will be updated at each iteration so that the \mathbf{x} remains close to \mathbf{x}^s . This update rule enables to keep the solutions within the region where the linearization is accurate. As a consequence, each iteration of our algorithm solves the following linear optimization problem:

$$\begin{aligned} &\underset{\Delta \mathbf{u}, \Delta \mathbf{x}}{\text{minimize}} && L^s(\Delta \mathbf{x}, \Delta \mathbf{u}) \\ &\text{subject to} && (5), \|\Delta \mathbf{u}\| \leq r^s, \mathbf{u}^s + \Delta \mathbf{u} \in \mathcal{U}^N, \mathbf{x}^s + \Delta \mathbf{x} \in \mathcal{X}^N. \end{aligned} \quad (6)$$

The optimal solution of the linearized problem is either accepted and used in the next iteration or rejected until convergence. When the linearization is considered accurate, i.e., the realized cost J and linearized cost L^s are similar, the solution is accepted and the trust region is expanded. Otherwise, the solution is rejected and the trust region is contracted.

Theorem 3 (SUBOPTIMALITY BOUND) Assume that L_c with the 2-norm is the Lipschitz constant of the cost c on $\mathcal{X} \times \mathcal{U} \times \mathcal{X}$. Let C_j^* and \hat{C}_j be the optimal costs of the control problem (2) following the sequential optimization scheme when the dynamics are known, e.g., $x^{q+1} = x(\cdot, x^q, u^q)$ is known, and the dynamics are unknown, e.g., $x^{q+1} = h^\theta(x^q, u^q) \in \mathcal{R}^{q+1}$. Then, $|C_j^* - \hat{C}_j| \leq L_c \left(\|\text{wd}(\mathcal{R}_{\mathcal{U}}^{j+N+1})\|_2 + \sum_{q=j+1}^{j+N} 2\|\text{wd}(\mathcal{R}_{\mathcal{U}}^q)\|_2 \right)$ holds with $\text{wd}(\mathcal{A}) = \bar{\mathcal{A}} - \underline{\mathcal{A}}$ being the width of the interval \mathcal{A} . The interval $\mathcal{R}_{\mathcal{U}}^{q+1}$ is the over-approximation of the reachable set at time index t_{q+1} from the initial uncertain set $\mathcal{R}_{\mathcal{U}}^q$ (with $\mathcal{R}_{\mathcal{U}}^j = \tilde{x}^j$) and for all $u^q \in \mathcal{U}$.

We provide the proof of Theorem 3 in the extended paper (Djeumou and Topcu, 2022). Theorem 3 provides that the suboptimality bound is proportional to the width of the over-approximation of the reachable set. Thus, our algorithm achieves near-optimal control with more data along the trajectory and more side information, as the over-approximations become tighter.

6. Numerical Experiments

In this section, we empirically demonstrate that the algorithm, using data from only the current trial and the least amount of side information necessary to learn, can achieve performance comparable to the highly-tuned implementations of D4PG (Hoffman et al., 2020) and SAC (Yarats and Kostrikov, 2020) trained over *ten million* of interactions with the environments. We emphasize that the comparison is unfair to our algorithm since, at each evaluating episode, it learns from only the *thousand* data obtained during the episode. Further, we show in an F-16 aircraft simulator, a 13-states and 4-control inputs nonlinear dynamics with polynomial control, that (a) The algorithm outperforms system identification approaches such as SINDYC (Kaiser et al., 2018); (b) The algorithm can meet real-time requirements. We provide further details on the numerical experiments in the extended paper (Djeumou and Topcu, 2022). A video of the simulations is at <https://tinyurl.com/hdem8x76>, and the code at <https://github.com/wuwushrek/datacontrolreach.git>.

Experiments in MuJoCo. The equations of motion for multi-joint dynamical systems in the MuJoCo environment are as follows: $M(q)\ddot{q} + b(\dot{q}, q) = h(u) + J_c^T(q)F_c(\dot{q}, q, u)$, where q is the system’s state, $M(q)$ is the inertial matrix, $b(\dot{q}, q)$ contains Coriolis, centrifugal, gravitational and passive forces, $J_c^T(q)$ is the contact Jacobian matrix, and $F_c(\dot{q}, q, u)$ is the contact force.

For each environment, the cost function is provided by MuJoCo, and we perform numeric differentiation in order to find its gradient. The Lipschitz bounds are under-estimated using only 1000 data points obtained prior to the on-the-fly control. The Reacher environment does not consider any side information other than the Lipschitz bounds, while Swimmer and Cheetah consider that $M(q)$ is known (Side information 1) in order to start learning. Indeed, without such side information, our algorithm fails to learn to control due to the large over-approximations of reachable sets. $M(q)$ is typically obtained for a robot through Euler-Lagrange formulation that uses the kinetic and potential energy. Further, we reduce the over-approximation of the contact force F_c by considering the Coulomb law of friction. That is, via Side information 2, we impose the constraints $F_c^1 \geq 0$ and $F_c^1 \geq \sqrt{(F_c^2)^2 \mu_1 + (F_c^3)^2 \mu_2}$ at each contact point, where F_c^1 is the normal force value, F_c^2 and F_c^3 are the tangential forces, and μ_1, μ_2 are the friction coefficients.

Figure 2 demonstrates that it is possible to learn to control with only data from a single episode by leveraging side information. In Cheetah, more side information can improve our algorithm’s performance. We reduced the time-step value of Reacher to accommodate our algorithm and observed that D4PG was unable to learn the task solely due to such a change, while SAC was not affected.




Figure 2: From left to right, we plot the (average) *immediate reward* over 100 episodes. The experiments show that our algorithm can yield performance comparable to D4PG and SAC.

Data-driven control of an F-16 aircraft. We consider a scenario involving an F-16 aircraft (Heidlauf et al., 2018) diving towards the ground at a low altitude and a high downward pitch angle. We show how our algorithm can prevent a ground collision using only the measurements obtained during the dive and elementary laws of physics as side information. We compare our algorithm with the linear-quadratic regulator (LQR) of the simulator, a pre-trained neural network for the task, and SINDYc achieving sparse system identification from a library of functions.

Our algorithm considers the structural knowledge of rigid-body dynamics while assuming that the aerodynamics forces and moments are completely unknown. In other words, the effect of the control inputs on the aircraft is unknown. For example, from the first principles, the lateral velocity’s derivative is given by $rv - qw - g \sin \theta + F_u/m$, where the structure is generic but the aerodynamic force F_u (specific to the aircraft) is unknown. We use the library `PySINDY` (de Silva et al., 2020) for the comparison with system identification. We considered monomials (up to degree 6), sines and cosines of the state, and the products of these functions with the control inputs as the library functions. We provide the noisy measurements of the state and its derivatives to both SINDYc and our algorithm. Our algorithm uses Lipschitz bounds estimated using 1000 data points. Finally, the neural network baseline was trained via policy optimization. Figure 3 empirically demonstrates the effectiveness of the proposed approach.




Figure 3: Our algorithm enables the F-16 to avoid the ground collision while the embedded LQR controller and SINDYc fail to avoid the crash. Further, it can be applied in real time since the compute time is less than the *control time step* enforced by the simulator.

7. Conclusion

This paper develops a learning-based, data-efficient control algorithm for unknown systems using streaming data from an ongoing trial and available side information. The experiments demonstrate that it is possible, with data from a single episode and side information, to achieve performance comparable to learning algorithms trained over millions of environment interactions. Further, we empirically show that the algorithm is fast and can be used in a scenario with real-time constraints.

Acknowledgments

This work was supported in part by AFOSR FA9550-19-1-0005 and NSF 1646522.

References

- Amir Ali Ahmadi and Bachir El Khadir. Learning dynamical systems with side information. In *LADC*, 2020.
- Gabriel Barth-Maron, Matthew W. Hoffman, David Budden, Will Dabney, Dan Horgan, TB Dhruva, Alistair Muldal, Nicolas Manfred Otto Heess, and Timothy P. Lillicrap. Distributed distributional deterministic policy gradients. *ArXiv*, abs/1804.08617, 2018.
- Frédéric Benhamou, Frédéric Goualard, Laurent Granvilliers, and Jean-François Puget. Revising hull and box consistency. In *ICLP*, 1999.
- Julian Berberich, Johannes Köhler, Matthias A Müller, and Frank Allgöwer. Data-driven model predictive control with stability and robustness guarantees. *IEEE Transactions on Automatic Control*, 66(4):1702–1717, 2020a.
- Julian Berberich, Carsten W. Scherer, and Frank Allgöwer. Combining prior knowledge and data for robust controller design. *ArXiv*, abs/2009.05253, 2020b.
- Martin Berz and Kyoko Makino. Verified integration of odes and flows using differential algebraic methods on high-order taylor models. *Reliable Computing*, 4:361–369, 1998.
- Joschka Boedecker, Jost Tobias Springenberg, Jan Wülfing, and Martin Riedmiller. Approximate real-time optimal control based on sparse gaussian process models. In *2014 IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, pages 1–8. IEEE, 2014.
- Steven L Brunton, Joshua L Proctor, and J Nathan Kutz. Sparse identification of nonlinear dynamics with control (sindy). *IFAC-PapersOnLine*, 49(18):710–715, 2016.
- Gilles Chabert and Luc Jaulin. Contractor programming. *Artificial Intelligence*, 173:1079–1100, 2009.
- Ankush Chakrabarty, Arvind Raghunathan, Stefano Di Cairano, and Claus Danielson. Data-driven estimation of backward reachable and invariant sets for unmodeled systems via active learning. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 372–377. IEEE, 2018.
- Jeremy Coulson, John Lygeros, and Florian Dörfler. Data-enabled predictive control: In the shallows of the deepc. *2019 18th European Control Conference (ECC)*, pages 307–312, 2019.
- Brian de Silva, Kathleen P. Champion, Markus Quade, Jean-Christophe Loiseau, J. Nathan Kutz, and Steven L. Brunton. Pysindy: A python package for the sparse identification of nonlinear dynamical systems from data. *Journal of Open Source Software*, 5:2104, 2020.
- Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472. Citeseer, 2011.

- Marc Peter Deisenroth, Gerhard Neumann, Jan Peters, et al. A survey on policy search for robotics. *Foundations and trends in Robotics*, 2(1-2):388–403, 2013.
- Alex Devonport and Murat Arcaç. Data-driven reachable set computation using adaptive gaussian process classification and monte carlo methods. In *2020 American Control Conference (ACC)*, pages 2629–2634. IEEE, 2020.
- Franck Djeumou and Ufuk Topcu. Learning to reach, swim, walk and fly in one trial: Data-driven control with scarce data and side information. *ArXiv*, abs/2106.10533, 2022.
- Franck Djeumou, Abraham P. Vinod, Éric Goubault, Sylvie Putot, and Ufuk Topcu. On-the-fly control of unknown systems: From side information to performance guarantees through reachability. *ArXiv*, abs/2011.05524, 2020.
- Franck Djeumou, Abraham P. Vinod, Eric Goubault, Sylvie Putot, and Ufuk Topcu. On-the-fly control of unknown smooth systems from limited data. In *2021 American Control Conference (ACC)*, pages 3656–3663. IEEE, 2021.
- Aditya Gahlawat, Pan Zhao, Andrew Patterson, Naira Hovakimyan, and Evangelos A. Theodorou. L1-gp: L1 adaptive control with bayesian learning. In *LADC*, 2020.
- Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous deep q-learning with model-based acceleration. In *International Conference on Machine Learning*, pages 2829–2838. PMLR, 2016.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *International Conference on Machine Learning (ICML)*, 2018.
- Sofie Haesaert, Paul MJ Van den Hof, and Alessandro Abate. Data-driven and model-based verification via bayesian identification and reachability analysis. *Automatica*, 79:115–126, 2017.
- Peter Heidlauf, Alexander Collins, Michael Bolender, and Stanley Bak. Verification challenges in f-16 ground collision avoidance and other automated maneuvers. In *ARCH@ADHS*, pages 208–217, 2018.
- Matt Hoffman, Bobak Shahriari, John Aslanides, Gabriel Barth-Maron, Feryal Behbahani, Tamara Norman, Abbas Abdolmaleki, Albin Cassirer, Fan Yang, Kate Baumli, Sarah Henderson, Alex Novikov, Sergio Gómez Colmenarejo, Serkan Cabi, Çağlar Gulcehre, Tom Le Paine, Andrew Cowie, Ziyu Wang, Bilal Piot, and Nando de Freitas. Acme: A research framework for distributed reinforcement learning. *ArXiv*, abs/2006.00979, 2020.
- Eurika Kaiser, J. Nathan Kutz, and Steven L. Brunton. Sparse identification of nonlinear dynamics for model predictive control in the low-data limit. *Proceedings. Mathematical, Physical, and Engineering Sciences*, 474, 2018.
- Milan Korda and Igor Mezić. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, 93:149–160, 2018.

- Andreas Krause and Cheng S Ong. Contextual gaussian process bandit optimization. In *Advances in NIPS*, pages 2447–2455, 2011.
- J Nathan Kutz. Deep learning in fluid dynamics. *Journal of Fluid Mechanics*, 814:1–4, 2017.
- Sergey Levine and Pieter Abbeel. Learning neural network policies with guided policy search under unknown dynamics. In *NIPS*, volume 27, pages 1071–1079. Citeseer, 2014.
- Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Manfred Otto Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *CoRR*, abs/1509.02971, 2016.
- Yuanqi Mao, Michael Szmuk, Xiangru Xu, and Behcet Acikmese. Successive convexification: A superlinearly convergent algorithm for non-convex optimal control problems. *ArXiv*, abs/1804.06539, 2019.
- Ivan Markovsky and Florian Dörfler. Data-driven dynamic interpolation and approximation. *Vrije Univ. Brussel, Tech. Rep*, 2021.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Belle-mare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charlie Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518:529–533, 2015.
- Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, pages 1928–1937. PMLR, 2016.
- Ramon E Moore. *Interval analysis*, volume 4. Prentice-Hall Englewood Cliffs, 1966.
- Anusha Nagabandi, Gregory Kahn, Ronald S Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7559–7566. IEEE, 2018.
- Nedialko S Nedialkov, Kenneth R Jackson, and George F Corliss. Validated solutions of initial value problems for ordinary differential equations. *Applied Mathematics and Computation*, 105(1):21–68, 1999.
- Junhyuk Oh, Valliappa Chockalingam, Honglak Lee, et al. Control of memory, active perception, and action in minecraft. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2016.
- Joshua L Proctor, Steven L Brunton, and J Nathan Kutz. Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems*, 15(1):142–161, 2016.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.

- John Schulman, Philipp Moritz, Sergey Levine, Michael I. Jordan, and P. Abbeel. High-dimensional continuous control using generalized advantage estimation. *CoRR*, abs/1506.02438, 2016.
- Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, Timothy P. Lillicrap, and Martin A. Riedmiller. Deepmind control suite. *ArXiv*, abs/1801.00690, 2018.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.
- Gilles Trombettoni, Yves Papegay, Gilles Chabert, and Odile Pourtallier. A box-consistency contractor based on extremal functions. In *International Conference on Principles and Practice of Constraint Programming*, pages 491–498. Springer, 2010.
- Pascal Van Hentenryck, Laurent Michel, and Yves Deville. *Numerica: a modeling language for global optimization*. MIT press, 1997.
- Henk J van Waarde, M Kanat Camlibel, and Mehran Mesbahi. From noisy data to feedback controllers: non-conservative design via a matrix s-lemma. *IEEE Transactions on Automatic Control*, 2020.
- Henk J Van Waarde, Jaap Eising, Harry L Trentelman, and M Kanat Camlibel. Data informativity: a new perspective on data-driven analysis and control. *IEEE Transactions on Automatic Control*, 65(11):4753–4768, 2020.
- Jan C. Willems, Paolo Rapisarda, Ivan Markovsky, and Bart De Moor. A note on persistency of excitation. *2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601)*, 3:2630–2631 Vol.3, 2004.
- Denis Yarats and Ilya Kostrikov. Soft actor-critic (sac) implementation in pytorch. https://github.com/denisyarats/pytorch_sac, 2020.