

Control-Oriented Learning on the Fly

Melkior Ornik, Arie Israel, and Ufuk Topcu

Abstract

This paper focuses on developing a strategy for control of systems whose dynamics are almost entirely unknown. This situation arises naturally in a scenario where a system undergoes a critical failure. In that case, it is imperative to retain the ability to satisfy basic control objectives in order to avert an imminent catastrophe. A prime example of such an objective is the reach-avoid problem, where a system needs to move to a certain state in a constrained state space. To deal with limitations on our knowledge of system dynamics, we develop a theory of myopic control. The primary goal of myopic control is to, at any given time, optimize the current direction of the system trajectory, given solely the information obtained about the system until that time. We propose an algorithm that uses small perturbations in the control effort to learn local dynamics while simultaneously ensuring that the system moves in a direction that appears to be nearly optimal, and provide hard bounds for its suboptimality. We additionally verify the usefulness of the algorithm on a simulation of a damaged aircraft seeking to avoid a crash, as well as on an example of a Van der Pol oscillator.

NOTICE OF PREVIOUS PUBLICATION

This manuscript is a significantly extended version of M. Ornik, A. Israel, U. Topcu, “Myopic Control of Systems with Unknown Dynamics” [1]. In addition to providing the detailed proofs omitted from that paper for lack of space, this manuscript contains novel sections discussing the existence of a myopically-optimal control law, design of goodness functions, robustness to disturbances, promising directions of future work, and a simulation presenting a myopic approach to control of a Van der Pol oscillator. In particular, the content of Section IV, Section VII, Section VIII-B, Section IX, Appendix A, and Appendix B is entirely novel, and multiple other sections have been substantially expanded.

I. INTRODUCTION

In an event of a rapid and unexpected change in the dynamics of a system, the ability to retain certain degree of useful control over a system is crucial. A notable real-world example of such a catastrophic event is that of an aircraft losing its wing after collision with another airplane [2]. In the particular event, the pilot managed to retain enough control of the aircraft to be able to safely land at a nearby airbase. The pilot’s strategy in this specific event depended on his intuition and prior experience at flying the aircraft. This paper seeks to develop a methodical

M. Ornik is with the Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX, 78712 USA e-mail: mornik@ices.utexas.edu.

A. Israel is with the Department of Mathematics, University of Texas at Austin, Austin, TX, 78712 USA e-mail: arie@math.utexas.edu.

U. Topcu is with the Department of Aerospace Engineering and Engineering Mechanics, University of Texas at Austin, Austin, TX, 78712 USA e-mail: utopcu@utexas.edu.

approach to control of an unknown system in real time. The strategy we propose is based on an intuitive approach one might use when trying to drive an unknown vehicle: performing small “wiggles” in controls to see how the system behaves, before deciding on a longer-term control action.

In the described setting of control of an unknown system, the only data available to extract information on the system dynamics can be obtained during the system run. Data-driven learning of the dynamics of unknown systems has been a subject of significant recent research [3], [4], [5], [6], [7]. However, the developed methods are data-intensive or, alternatively, require some a priori knowledge of system dynamics; hence, they are not suitable for the above motivating example. A recent work [8] based on a similar scenario uses differential inclusions to assess the safety of a trajectory governed by unknown dynamics. However, this work merely discusses uncontrolled dynamics, and does not prescribe a control law which would ensure safety of a trajectory. A number of works do deal with control design for unknown dynamics, with or without trying to explicitly learn a model [9], [10], [11], [12]. However, the results of those works again significantly differ from ours, as they generally do not include formal guarantees on short-time, non-asymptotic system performance. Additionally, the control objectives discussed in those papers are markedly different than the ones that naturally arise in the scenario we are exploring, as they are all based on reference tracking. In particular, [9], [10], and [12] employ methods based on transfer functions to solve a tracking problem in an unknown system, while [11] uses neural networks in optimal tracking control.

In contrast to the previously mentioned papers, we are primarily concerned with solving reach-avoid-type problems. This consideration is motivated by practical reasons: in the event of a system failure, satisfying advanced control specifications likely becomes impossible, and the system should concentrate on its survival. In physical systems, it is reasonable to pose the problem of system survival as a question of reaching a certain target set as soon as possible while avoiding obstacles and staying within a predefined safety envelope. In the motivating example of a damaged aircraft, which we will use as a running example throughout this paper, the target set is the airbase runway, and the aircraft should reach it as soon as possible without touching the ground beforehand.

A standard approach to reach-avoid problems in systems with known dynamics [13] is based on constrained optimal control, where the constraints are given by the geometry of the safety envelope and obstacles, and the time to reach a target set is minimized. The solutions to this problem clearly depend on the full dynamics of the system. In our setting we have very little information on the system dynamics. In fact, we only have the information that can be obtained by observing a system trajectory since the beginning of the system run, i.e., since the abrupt change. Hence, it is unlikely that we can find the exact optimal control law to solve a reach-avoid problem.

We instead propose the notion of *myopic control*, and use it to suboptimally solve reach-avoid problems. In this notion, the system uses a control law that seems to work best *at the given time*, without any knowledge on whether that control law will lead to good results in the future. This framework is motivated by the inherent lack of firm knowledge of future system dynamics.

To counterbalance the lack of confidence about the future, a new myopically-optimal control law is calculated and applied every so often, after a short time interval. In order to determine the next myopically-optimal control law, the system continually modifies the previous law by adding a series of *wiggles*: small, short-time controls that do not significantly change the system trajectory, but act as a learning mechanism to determine the local control dynamics

of the system. In order for this mechanism to function, we make a technical assumption that the underlying system is affine in control. We briefly discuss a potential modified strategy for the general case.

The proposed algorithm results in an approximate solution to the myopic optimal control problem, with a degree of suboptimality dependent on the length of the control law update interval and the size of learning controls, i.e., wiggles. Additionally, if there are known bounds on regularity of system dynamics, the parameters of the algorithm can be set to ensure any desired bound on the degree of suboptimality. Finally, since the algorithm does not use any a priori information on the system (except, perhaps, the aforementioned bounds on regularity), but learns its dynamics on the fly, it is entirely robust to disturbances, as it simply treats them as part of the original system dynamics.

The organization of this paper is as follows: Section II formally introduces the class of problems we are interested in solving. Section III then defines the myopic optimal control problem. This problem is founded on a notion of apparent *goodness* of a control law that attempts to satisfy a certain objective. Thus, Section IV discusses design of an appropriate goodness function. Section V provides an overview of the proposed algorithm to solve the myopic optimal control problem. Section VI then provides performance bounds for the proposed algorithm. In Section VII we discuss the robustness of the algorithm in the presence of disturbances. Section VIII provides simulation results for the running example of a damaged aircraft, showing that the algorithm indeed performs as expected, as well as a highly nonlinear example of a Van der Pol oscillator. Substantial open questions identified in previous sections are discussed in Section IX, while Section X provides concluding remarks. The paper additionally contains two appendices. Appendix A contains a brief discussion of some sufficient conditions for the existence of a solution to the myopic optimal control problem. Finally, as the mathematical results used in Section VI are rather technical, their full proofs are contained in Appendix B.

Notation. Set of all continuous functions from set $\mathcal{A} \subseteq \mathbb{R}^n$ to set $\mathcal{B} \subseteq \mathbb{R}^m$ is denoted by $C(\mathcal{A}, \mathcal{B})$. For a set \mathcal{A} , $2^{\mathcal{A}}$ denotes the set of all subsets of \mathcal{A} . The right one-sided derivative of function $f : I \rightarrow \mathbb{R}^n$ at time t is denoted by $df(t+)/dt$. For a vector $v \in \mathbb{R}^n$, $\|v\|$ denotes its 2-norm, and $\|v\|_{\infty}$ its max-norm. If $x, y \in \mathbb{R}^n$, $x \cdot y$ denotes the standard dot product of x and y . Unless stated otherwise, for $x \in \mathbb{R}^n$, x_i , $i \in \{1, 2, \dots, n\}$, denotes the i -th component of x .

II. PROBLEM DEFINITION

We investigate the control system

$$\dot{x} = f(x) + \sum_{i=1}^m g_i(x)u_i \quad (1)$$

which evolves on $\mathcal{X} \subseteq \mathbb{R}^n$, where $f, g_i \in C(\mathcal{X}, \mathcal{X})$ are unknown functions, and $u = (u_1, \dots, u_m)$ is the m -dimensional control input with values in the bounded set $\mathcal{U} \subseteq \mathbb{R}^m$. While it is trivial to modify our proposed algorithm for the case where \mathcal{U} is any manifold with corners in the sense of [14], for the simplicity of our argument we take $\mathcal{U} = \{u \in \mathbb{R}^m \mid \|u\|_{\infty} \leq 1\}$. In the remainder of the text, a solution to (1) with control input u and initial value x_0 is denoted by $\phi_u(\cdot, x_0)$.

We are assuming that, at any time instance, we are able to observe the full state x and the entire control u . The case of partial observations is also of significant interest for physical systems. Even though it is not the primary

focus of this paper, we briefly discuss it in Section IX. Additionally, in a practical sense of control of a real-world system, we are usually only interested in the behavior of a system on a compact state space \mathcal{X} . It thus makes sense to assume that there exist $M_0 \geq 0$ and $M_1 \geq 0$ such that

- $\|f(x)\| \leq M_0, \|g_i(x)\| \leq M_0$ for all $x \in \mathcal{X}, i \in \{1, \dots, m\}$,
- $\|f(x) - f(y)\| \leq M_1\|x - y\|, \|g_i(x) - g_i(y)\| \leq M_1\|x - y\|$ for all $x, y \in \mathcal{X}, i \in \{1, \dots, m\}$.

The class of control-affine systems covers a wide array of physical control systems, including standard linear aircraft models [15], [16], [17], and a variety of nonlinear models of mechanical systems [18], e.g., a unicycle with rider [19], quadrotor helicopter [20], and the F-8 Crusader aircraft [21]. Control and optimal control of general control-affine systems have been substantially covered by previous literature (see, e.g., [18], [22], [23], [24]), including investigations of reach-avoid problems [25]. However, motivated by the scenario of an unexpected failure in a physical system, our setting introduces two additional requirements on control design:

R1 Functions f, g_1, \dots, g_m in (1) are unknown at the beginning of the system run. We are certainly allowed to use trajectory data during the system run to determine information on them, but when the run starts, we are aware merely that the system is of the form (1), and potentially know a bound on the functions' norms M_0 and Lipschitz constants M_1 .

R2 There is only one system run, starting from a single predetermined initial state x_0 . All control design needs to be performed during that run, without any repetitions.

We note that the bulk of previous work on systems with disturbances or noise, as discussed in, e.g., [26], [27], does not help to deal with requirement R1. The standard setup in the theory of systems with disturbances requires the controller to know the “general” system dynamics, which are then modified by some unknown, but bounded, disturbance or noise. Requirement R1 is significantly more limiting for control design. It stipulates that the system dynamics are entirely unknown at the beginning of the system run. Potential additional disturbances to system dynamics in our setup are discussed in Section VII.

As previously mentioned, the goal of this paper is to develop a method for solving reach-avoid-type control problems [28], [29] in the control-affine setting (1) under the above requirements R1 and R2. We use the following formulation for the reach-avoid problem:

Problem 1 (Reach-avoid Problem). *Let $x_0 \in \mathcal{X}, T \geq 0, \mathcal{T} \subseteq \mathcal{X}$, and $\mathcal{B} \subseteq \mathcal{X}$. Let $\phi_u(\cdot, x_0) : [0, T] \rightarrow \mathcal{X}$ be the trajectory generated by system (1) under control u and with $\phi_u(0, x_0) = x_0$. Find, if it exists, the control law $u^* : [0, T] \rightarrow \mathcal{U}$ such that the following holds:*

- (i) $\phi_{u^*}(t, x_0) \notin \mathcal{B}$ for all $t \in [0, T]$,
- (ii) there exists $0 \leq T_{u^*}^t \leq T$ such that $\phi_{u^*}(t, x_0) \in \mathcal{T}$ for all $t \in [T_{u^*}^t, T]$, and
- (iii) $T_{u^*}^t$ from (ii) is the minimal such value for all the control laws $u : [0, T] \rightarrow \mathcal{U}$ which satisfy (i) and (ii).

Let us build intuition through a simplification of Problem 1. We note that if there is no particular target set \mathcal{T} , the reach-avoid problem naturally transforms into the problem of remaining outside the undesirable set for as long as possible. This problem is again one natural abstraction of our running example: instead of having to deal with

the technicalities of aircraft landing, we may simply want to pose the question of requiring the damaged aircraft to stay above ground for as long as possible. Let us provide a formal description of this version of the problem.

Definition 2. Let $x_0 \in \mathcal{X}$, and let $\mathcal{B} \subseteq \mathcal{X}$ be closed. Using the notation of Problem 1, the first bad time of the trajectory ϕ_u is given by

$$b(u) = \min_{t \geq 0} \{ \phi_u(t, x_0) \in \mathcal{B} \}.$$

If $\phi_u(t, x_0) \notin \mathcal{B}$ for all $t \geq 0$, we take $b(u) = \infty$.

Problem 3 (Avoid Problem). Let $x_0 \in \mathcal{X}$, and let $\mathcal{B} \subseteq \mathcal{X}$ be closed. Find a control law $u^* : [0, +\infty) \rightarrow \mathcal{U}$ such that

$$b(u^*) = \max_u b(u).$$

This paper would ideally end by providing a control design to solve Problem 1 and Problem 3 which respects requirements R1 (unknown system) and R2 (no repeated runs). However, this outcome is likely impossible. The requirements we have imposed imply that, at the beginning of the system run, we have no particular knowledge of the values of functions $f(x)$ and $g_i(x)$ at any point x , and no way of acquiring significant knowledge about the dynamics in a certain area of \mathcal{X} until the system trajectory reaches that area. In fact, even when the system trajectory reaches a point x , we are still unable to exactly determine the values of $f(x)$ and $g_i(x)$, because a single controlled trajectory provides only limited information on full control dynamics around a point. Thus, it is not reasonable to expect that Problem 1 and Problem 3 can be exactly solved under the requirements of our setting. We now present an intuitive variation of these problems that can be approached using strategies which respect requirements R1 and R2.

III. MYOPIC CONTROL

Suppose that we want to solve the reach-avoid or avoid problem under requirements R1 and R2. Requirement R2 implies that, in order to determine the optimal control law u^* to be used starting at some time t , one may only use the information obtained from the system trajectory until time t . Thus, based on our knowledge of the system dynamics in the interval $[0, t]$, one would have to assess the “quality” of a candidate future control law on the entire interval $(t, +\infty)$. However, we can only do so with very limited confidence, since the dynamics far away from points already visited at some time $0 \leq t' \leq t$ are entirely unknown. Thus, at time t we have no knowledge with certainty on whether a particular control will result in good behavior at time $t + \alpha$, for some large α .

The above discussion behooves us to replace Problem 1 and Problem 3 by a *myopic optimal control problem*, where we want to design a control law such that, at every time instance, the trajectory *appears* to behave as well as possible. For instance, if our ultimate desire is to solve Problem 3, it makes sense to require that the trajectory at any given instance of time is moving away from the undesirable set \mathcal{B} as fast as possible. Clearly, this reasoning may not be ultimately optimal, as moving away from \mathcal{B} as fast as possible at some point could bring the trajectory into a position where the dynamics are such that it is simply forced to enter \mathcal{B} (see Figure 1 for an example). Thus, this notion of “appearing to behave as well as possible” is simply our best guess. Since we do not have (almost)

any knowledge of the dynamics far away from our current position, the best we can do is try to find a seemingly optimal direction of movement at any given time, given the totality of prior information obtained during the system run until that point.

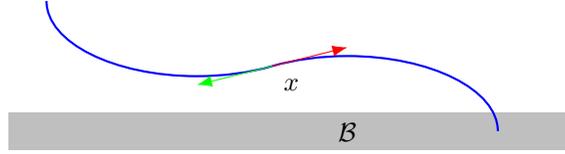


Fig. 1. An example illustrating the built-in imperfection of myopic control. Assume that the system dynamics are such that, unbeknownst to the control designer, the system is restricted to moving along the blue curve. The control objective is to avoid the undesirable set \mathcal{B} , denoted in gray. When the system is at point x , the available direction vectors are denoted by the green and red arrows. From the designer's myopic perspective, it may appear better to control the system in the red direction, because this direction forces the system to move further away from \mathcal{B} . That action will ultimately result in the system moving into the undesirable set. Given the designer's lack of knowledge about the system, such a consequence is unforeseeable.

We propose to formalize the above notion of “appearing to behave as well as possible” using a *goodness function*

$$(\phi, v) \mapsto G(\phi, v), \quad \phi \in \mathcal{F}, v \in \mathbb{R}^n, \quad (2)$$

where

$$\mathcal{F} = \bigcup_{T \geq 0} C([0, T], \mathcal{X}).$$

Function G is intended to quantify how well the trajectory appears to be doing at satisfying its control specifications at a time when its trajectory until time T and velocity at time T are given by ϕ and v , respectively.

Myopic policies and goodness functions have long been used in decision theory, in a wide variety of settings, including resource allocation [30], computer vision [31], and Bayesian learning [32]. However, all these settings differ significantly from the setting of this paper; to the authors' knowledge, myopic decision-making has never been utilized in dealing with control of unknown systems.

Example 4. *Going back to our running example of a damaged aircraft, assume that we want to solve Problem 3, i.e., have the aircraft stay away from the ground for as long as possible. In that case, if x_1 and x_2 denote the aircraft's horizontal and vertical position, respectively, $\mathcal{B} \subseteq \mathbb{R}^2$ is given by $\mathcal{B} = \{(x_1, x_2) \mid x_2 \leq 0\}$. Then, one natural option is for G to measure the slope on which the trajectory is moving towards the boundary of \mathcal{B} , i.e., $x_2 = 0$. Without any additional information on the system, the more negative this slope is, the worse the aircraft naturally appears to be doing. Hence,*

$$G(\phi, v) = \frac{v_2}{v_1},$$

where we disregard the case of $v_1 \leq 0$, is a reasonable goodness function. Figure 2 provides an illustration of this goodness function.

In general, determining an appropriate goodness function depends on our information on the system and the control specifications. We briefly discuss this problem, and provide further examples, in Section IV.

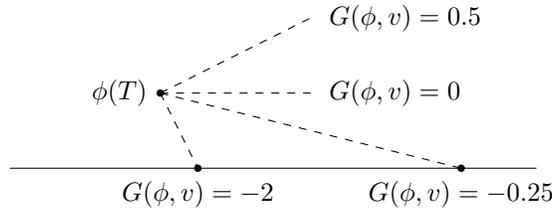


Fig. 2. An illustration of the goodness function G from Example 4. Dashed lines represent the possible tangents to the trajectory of system (1) at $\phi(T)$, and thus, our local predictions of future movement of the system state. These tangents are evaluated by G depending on their slope towards the undesirable set \mathcal{B} , whose edge is represented by a solid line.

Having discussed how a goodness function provides a measure of apparent quality of a trajectory at a given point (with respect to, for instance, specifications of Problem 1 and Problem 3), we now introduce the myopic optimal control problem. As a slight deviation from standard notation, in order to emphasize that $\phi \in \mathcal{F}$ has $[0, T]$ as its domain, we will occasionally write such functions as $\phi|_{[0, T]}$.

Problem 5 (Myopic Optimal Control Problem). *Let $x_0 \in \mathcal{X}$. Find a control law $u^* : [0, +\infty) \rightarrow \mathcal{U}$ such that, for all $t \geq 0$, if $x = \phi_{u^*}(t, x_0)$, then*

$$G\left(\phi_{u^*}(\cdot, x_0)|_{[0, t]}, \frac{d\phi_{u^*}(0+, x)}{dt}\right) = \max_{u \in \mathcal{U}} G\left(\phi_u(\cdot, x_0)|_{[0, t]}, \frac{d\phi_u(0+, x)}{dt}\right). \quad (3)$$

We note that in (3) we are slightly abusing notation: if $u \in \mathcal{U}$, then u is not formally a control input as a function of time. Thus, by $d\phi_u(0+, x)/dt$ we denote the value of the right-hand side of (1) for the particular $x \in X$ and $u \in \mathcal{U}$. For the sake of narrative, we will usually omit a formal distinction between a control law u and a control input $u \in \mathcal{U}$ when such a distinction is obvious from the context. Additionally, while the notion of a solution to (1), and hence to (3) as well, is questionable in the case of general measurable functions u^* , in the remainder of the paper we will be working with piecewise-continuous control laws, and in that sense we will not have any issues.

In subsequent sections of this paper, we will propose an approximate piecewise-constant solution to Problem 5. This solution can be made arbitrarily close to optimal, in the sense that for any $T > 0$ and $\mu > 0$, Algorithm 10 can generate a piecewise-constant control law that satisfies

$$\left| G\left(\phi_{u^*}(\cdot, x_0)|_{[0, t]}, \frac{d\phi_{u^*}(0+, x)}{dt}\right) - \max_{u \in \mathcal{U}} G\left(\phi_u(\cdot, x_0)|_{[0, t]}, \frac{d\phi_u(0+, x)}{dt}\right) \right| < \mu \quad (4)$$

for all $t \geq T$. We note that this result does not directly guarantee existence of an *exact* solution to Problem 5. Some sufficient conditions for the existence of such a solution can be obtained. Nonetheless, because of our inability to exactly learn the system dynamics at even a single point from a single trajectory, it is not realistic to expect that any method of finding an exact solution to Problem 5 would respect requirements R1 and R2. Discussing such solutions is thus not in the focus of this paper. However, we refer the reader to a brief discussion of sufficient conditions for the existence of a solution to Problem 5 contained in Appendix A.

Before proposing a strategy to find an approximate piecewise-constant solution to Problem 5 (while respecting requirements R1 and R2), let us briefly return to the issue of determining an appropriate goodness function, given

the control objective and system information.

IV. MEASURING GOODNESS

Determining an appropriate goodness function G is generally not easy. We provided a very simple possible goodness function for one scenario in Example 4. Let us provide three more examples of natural goodness functions, two for Problem 1 and one for Problem 3.

Example 6. *As the sole objective of Problem 3 is not to enter the bad set \mathcal{B} , it makes sense to desire the trajectory to be moving away from \mathcal{B} as fast as possible. Without any additional knowledge, our best approximation of the system state at time Δt after it is at point x is $x + v\Delta t$, where v is the velocity vector at the given time. Hence, if $d_{\mathcal{B}} : \mathcal{X} \rightarrow [0, +\infty)$ denotes the usual distance from a point to set \mathcal{B} , one reasonable option for G is to take*

$$G(\phi|_{[0,t]}, v) = \liminf_{\Delta t \rightarrow 0^+} \frac{d_{\mathcal{B}}(\phi(t) + v\Delta t) - d_{\mathcal{B}}(\phi(t))}{\Delta t}.$$

This approach to a goodness function is similar to the one developed in [33] in the context of viability theory, where the distance from the undesirable set is considered as an indicator of a trajectory robustness. Since in our case we are interested in measuring the “robustness” (i.e., goodness) of our prediction of the future trajectory, we replace the distance from the undesirable set by the rate of change of the distance.

Remark 7. *We note that function G in Example 6 is not a generalization of the goodness function from Example 4: adapted to set \mathcal{B} from Example 4, the goodness function from Example 6 would be given by v_2 , instead of v_2/v_1 as in Example 4. Both of these goodness functions certainly possess a reasonable rationale: while the function from Example 4 measures how far we expect the trajectory to get before it hits \mathcal{B} , function from Example 6 measures how fast we expect the trajectory to hit \mathcal{B} . This difference between two different intuitively reasonable goodness functions for the same problem serves to confirm our previous point that it is difficult to develop the notion of an objectively correct goodness function. As mentioned, while we propose several goodness functions for the reach-avoid and avoid problems in this section, a formal investigation of a notion of a correct goodness function is not a focus of this paper.*

The goodness function proposed in Example 6 will form the basis of the goodness function for the example of a damaged aircraft in Section VIII-A. However, because of the particularities of aircraft design, the control specifications in that example are more complicated than just stating that the aircraft must stay at a positive altitude. Additionally, due to the limitations of aircraft design, certain states cannot be directly influenced by control actions. Hence, the goodness function in Section VIII-A is more difficult than the one presented in Example 6. While we refer the reader to Section VIII-A for details, we emphasize that the goodness function in Example 6 provides a basis for the developments of that section. We now proceed to propose a goodness function for the full reach-avoid problem.

Example 8. *Framework of Problem 1 is significantly more complicated than Problem 3, as there are two objectives that need to be satisfied at the same time. One naive reasonable goodness function uses the following motivation:*

let us partition the safe set $\mathcal{X} \setminus \mathcal{B}$ into a “boundary zone” \mathcal{Z}^B and a “interior zone” \mathcal{Z}^I . The motivation for naming of the two sets is that, if a system state is inside \mathcal{Z}^B , it is considered at risk of leaving the safe set, and the primary objective becomes to enter the safer interior zone as soon as possible. If a system state is inside the interior zone, it is not considered at risk, and the primary objective is to reach the target set \mathcal{T} . Then, a possible goodness function is given by

$$G(\phi|_{[0,t]}, v) = -\liminf_{\Delta t \rightarrow 0^+} \frac{d_{\mathcal{Z}^I}(\phi(t) + v\Delta t) - d_{\mathcal{Z}^I}(\phi(t))}{\Delta t}$$

if $\phi(t) \in \mathcal{Z}^B$, and

$$G(\phi|_{[0,t]}, v) = -\liminf_{\Delta t \rightarrow 0^+} \frac{d_{\mathcal{T}}(\phi(t) + v\Delta t) - d_{\mathcal{T}}(\phi(t))}{\Delta t}$$

if $\phi(t) \in \mathcal{Z}^I$. Again, the above goodness function is not the only reasonable one, and has its drawbacks. For instance, it is not continuous in $\phi(t)$, which will become important when determining bounds for suboptimality of a solution to Problem 5 in Section VI. A more subtle goodness function may use a measure of how much a system state is at risk (for instance, distance from \mathcal{B}), and then use a mixed approach between trying to enter set \mathcal{T} and not leave $\mathcal{X} \setminus \mathcal{B}$:

$$G(\phi|_{[0,t]}, v) = (1 - \lambda) \liminf_{\Delta t \rightarrow 0^+} \frac{d_{\mathcal{B}}(\phi(t) + v\Delta t) - d_{\mathcal{B}}(\phi(t))}{\Delta t} - \lambda \liminf_{t \rightarrow 0^+} \frac{d_{\mathcal{T}}(\phi(t) + v\Delta t) - d_{\mathcal{T}}(\phi(t))}{\Delta t},$$

where

$$\lambda = \frac{d_{\mathcal{B}}(\phi(t))}{\max_y d_{\mathcal{B}}(y)}.$$

As the above examples show, there are multiple reasonable goodness functions for a particular control objective, and the choice of which function to use rests with the designer. Let us now briefly mention a particularly attractive facet of this freedom to choose a goodness function. The choice of goodness functions is a natural vessel for the incorporation of any kind of side information on the system. For instance, if we were to know that the control dynamics evolve on a particular subset of \mathcal{X} instead of entire \mathcal{X} , it might make sense to choose the goodness function in a way that is conscious of such information. For example, Figure 3 illustrates a situation where the system is known to evolve on a particular subset of \mathbb{R}^2 , and while the naive goodness function would choose a direction that points towards the desired target set, but ultimately leads nowhere, an informed goodness function would choose a direction which seems counterintuitive, but ultimately leads towards the target set.

Remark 9. In all of the above examples, the proposed goodness functions are memoryless: $G(\phi|_{[0,T]}, v)$ is actually merely a function of $(\phi(T), v)$. However, the proposed theory allows for more complicated measures of goodness, incorporating knowledge obtained earlier in the system run. For instance, the fact that a goodness function is permitted to be memory-conscious allows to directly deal with temporal logic specifications. We consider linear temporal logic (LTL) specifications given over a collection $\mathcal{P} = \{\mathcal{A}_k \mid k \in K\}$ of subsets of \mathcal{X} . An example of such a specification is “visit \mathcal{A}_3 after \mathcal{A}_2 , and do not enter \mathcal{A}_4 before visiting \mathcal{A}_1 ”. We omit a formal description of LTL, and direct the reader to [34]. Informally, an LTL formula determines all allowed discrete behaviors of a trajectory, where the discrete behavior of a trajectory describes the order in which a system trajectory visits the subsets in \mathcal{P} . For any desired discrete behavior, the problem of finding a control law such that the system trajectory

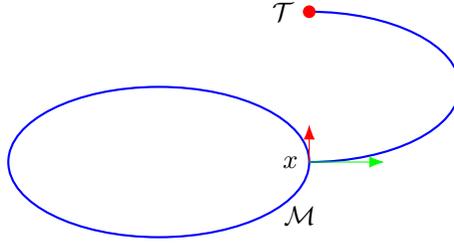


Fig. 3. An example of the role of side information when designing a goodness function. The control objective is to reach the set \mathcal{T} , and system trajectories are, unlike in Figure 1, a priori known to evolve on the set $\mathcal{M} \subseteq \mathbb{R}^2$ (drawn in blue). Assume that the system trajectory is currently (at time t) at position x , and available velocity vectors at that point are given by the red and green arrows. While an uninformed goodness function might prefer the red arrow, which points directly towards \mathcal{T} , the available side information urges the designer to apply such a goodness function that the trajectory continue in the direction indicated by the green arrow.

has that behavior can be posed as a series of reach-avoid problems [35]. In other words, the system is first required to reach some target set \mathcal{T}_1 while avoiding an undesirable set \mathcal{B}_1 , then required to reach set \mathcal{T}_2 while avoiding \mathcal{B}_2 , etc. Guided by our previous examples, a reasonable goodness function to solve such a problem is given by $G(\phi|_{[0,t]}, v) = G_i(\phi|_{[0,t]}, v)$, where $i \in \mathbb{N}$ is the smallest index such that $\{\phi|_{[0,t]}(\tau) \mid \tau \in [0, t]\} \cap \mathcal{T}_i = \emptyset$, and G_i is a goodness function for Problem 1 with target set \mathcal{T}_i and undesirable set \mathcal{B}_i (for instance, G_i can be set to one of the two goodness functions proposed in Example 8).

The development of general goodness functions is not in the focus of this paper. However, the examples above show that goodness functions can be successfully designed for a variety of reach-avoid-type problems. We now turn towards our proposed algorithm for the approximate solution of Problem 5.

V. LEARN-CONTROL ALGORITHM

We intend to approximately solve Problem 5 by the use of short-time piecewise-constant controls. The proposed strategy relies on repeatedly learning the local dynamics at points on the observed system trajectory, and using those dynamics to simultaneously apply a myopically-optimal control law.

In order to learn the local dynamics around a single point on the trajectory, the controller can apply any $m + 1$ affinely independent constant controls for a short period of time. Since system (1) is control-affine, we can use the observed changes in states to approximate the function $u \mapsto v_x(u)$ defined by

$$v_x(u) = f(x) + \sum_{i=1}^m g_i(x)u_i,$$

where x is a point on the trajectory at the end of this learning process.

After obtaining an approximation for v_x , if ϕ denotes the trajectory until the end of the learning process, by determining $\operatorname{argmax}_u G(\phi, v_x(u))$, a constant control u^* that appears to best satisfy the control specifications at point x can be found. This control is then used during the next iteration of the learning process. In other words, the next learning process uses an affinely independent set of controls $u^* + \Delta u^0, \dots, u^* + \Delta u^m$, where Δu^i are small enough to ensure that their application still results in near-optimal goodness.

The described method bears some resemblance to the approach in [36], which also uses short-time controls to achieve real-time (nearly) optimal behavior. However, the dynamics in [36] are presumed to be known. Thus, there is no learning element to the control strategy. Additionally, the control objective is to solve a standard optimal control problem, instead of the myopic version naturally imposed by the requirements of our scenario.

The above description of the control strategy is formally given as Algorithm 10.

Algorithm 10.

```

1   $u^* := 0$ 
2   $t_0 := 0$ 
3  repeat
4      Let  $x^0$  be the current state of the system.
5      Take  $\Delta u^0 = 0$  and choose  $\Delta u^1 = \pm \delta e_1, \dots, \Delta u^m = \pm \delta e_m$  such that  $u^* + \Delta u^1, \dots, u^* + \Delta u^m \in \mathcal{U}$ .
6      for  $j = 0, \dots, m$ 
7          Apply control  $u^* + \Delta u^j$  in the interval of time  $[t_0 + j\varepsilon, t_0 + (j + 1)\varepsilon]$ .
8          Let  $x^{j+1}$  be the system state at the end of that period.
9      end for
10      $x := x^{m+1}$ 
11     Define function  $\tilde{v} : \mathcal{U} \rightarrow \mathbb{R}^n$  as follows:
           If  $\lambda_0, \dots, \lambda_m$  are unique coefficients with  $\sum \lambda_j = 1$  such that  $u = \sum \lambda_j (u^* + \Delta u^j)$ , let  $\tilde{v}(u) = \sum \lambda_j (x^{j+1} - x^j) / \varepsilon$ .
12     Take  $u^* \in \operatorname{argmax} G(\phi|_{[0, t_0 + (m+1)\varepsilon]}, \tilde{v}(u))$ 
13      $t_0 := t_0 + (m + 1)\varepsilon$ 
14 until the end of system run

```

We note that while Algorithm 10 prescribes a particular form for Δ^i in line 5, the local dynamics of the system can be learned by using any Δ^i such that $u^* + \Delta u^1, \dots, u^* + \Delta u^m \in \mathcal{U}$ is an affinely independent set. Algorithm 10 includes a particular form for Δ^i in order to more easily establish a bound on the degree of suboptimality of the resulting control law. Such a bound, given in the following section, is dependent on parameters $0 < \delta \leq 1$ and $\varepsilon > 0$. With an appropriate choice of δ and ε , we can ensure that Algorithm 10 produces a control law which comes arbitrarily close to solving Problem 5. This choice of δ and ε will naturally depend on the bounds M_0 and M_1 on the norms and Lipschitz constants of f and g_i , and the regularity of goodness function G .

Remark 11. *Algorithm 10 does not distinguish between the learning and control phases of the algorithm: the algorithm learns the local dynamics as a result of performing slight perturbations of the previously established optimal control law. However, we note that these two phases can be decoupled by a minor modification of the above algorithm. After learning the local dynamics through consecutively applying any $m+1$ affinely independent controls in a short time period $\varepsilon' < \varepsilon$, the system can then apply the optimal control law derived from these dynamics for*

the remaining $\varepsilon - \varepsilon'$ time in the cycle, after which it begins a new cycle by learning the new local dynamics. While this modification results in a short ε' period of time inside every iteration of the learn-control cycle without any guarantees on the degree of myopic suboptimality, it offers computational benefits in the case where the optimal control law does not rapidly change over time.

In light of this modification, we note that the performance bounds presented in the subsequent section apply to Algorithm 10 as originally presented. Analogous proofs could be constructed for the modified algorithm as above. For computational reasons, the algorithm used in simulations in Section VIII-A contains the modification described above, while the work in Section VIII-B uses Algorithm 10 directly.

VI. PERFORMANCE BOUNDS

A. Learning

We claim that the procedure in lines 5–11 of Algorithm 10 produces a good approximation $\tilde{v} : \mathcal{U} \rightarrow \mathbb{R}^n$ of the map $v_x : \mathcal{U} \rightarrow \mathbb{R}^n$, with $x = x^{m+1}$ defined as in the algorithm.

We first note that, for every point $u^* \in \mathcal{U} = [-1, 1]^m$, if $u_i^* \geq 0$, then $-1 \leq u_i^* - \delta < 1$, and if $u_i^* < 0$, then $-1 < u_i^* + \delta < 1$. Hence, for all $u^* \in \mathcal{U}$ we can indeed choose $\Delta u^i = \delta e_i$ or $\Delta u^i = -\delta e_i$ such that $u^* + \Delta u^i \in \mathcal{U}$, as stipulated by line 5 of Algorithm 10. We also note that $u^* + \Delta u^i$, $i \in \{0, \dots, m\}$, trivially form an affinely independent set.

We first describe the remainder of the procedure informally. We denote $u^* + \Delta u^i$ by u^i .

For each $j \in \{0, 1, \dots, m\}$, vector

$$\frac{x^{j+1} - x^j}{\varepsilon} = \frac{\phi_{u^j}(\varepsilon, x^j) - \phi_{u^j}(0, x^j)}{\varepsilon} \quad (5)$$

approximates the value $d\phi_{u^j}(\varepsilon, x^j)/dt = v_{x^{j+1}}(u^j)$, where we recall that $v_y(u)$ is defined by $v_y(u) = f(y) + \sum g_i(y)u_i$. Additionally, since x^{j+1} and $x = x^{m+1}$ are not far apart (because x^{m+1} is the state of the trajectory just $(m-j)\varepsilon$ later than x^{j+1}), $v_{x^{j+1}}(u^j) \approx v_{x^{m+1}}(u^j) = v_x(u^j)$.

Finally, since u^0, \dots, u^m are affinely independent, for every u there exist unique $\lambda_0, \dots, \lambda_m \in \mathbb{R}$ such that $u = \lambda_0 u^0 + \dots + \lambda_m u^m$ and $\lambda_0 + \dots + \lambda_m = 1$. Thus,

$$f(x) + \sum_{i=1}^m g_i(x)u_i = \sum_{j=0}^m \lambda_j \left(f(x) + \sum_{i=1}^m g_i(x)u_i^j \right),$$

i.e. $v_x(u) = \sum \lambda_j v_x(u^j)$. Since we already have approximations for $v_x(u^j)$ from (5) and the subsequent discussion, we can thus approximate $v_x(u)$ by taking

$$v_x(u) \approx \sum_{j=0}^m \lambda_j \frac{x^{j+1} - x^j}{\varepsilon}. \quad (6)$$

It can be formally shown that approximation (6) produces an error no worse than linear in ε . This result is given in Theorem 13, while foregoing technical claims are given in Lemma 12.

Lemma 12. *Let u be the control law produced in one **repeat** loop of Algorithm 10 (i.e., u is the concatenation of u^0, u^1, \dots, u^m). Also, let x^j be defined as in Algorithm 10, where all x^j belong to the same iteration of the **repeat** loop. Then, the following holds:*

(i) For all $t_1, t_2 \in [0, (m+1)\varepsilon)$,

$$\|\phi_u(t_1, x^0) - \phi_u(t_2, x^0)\| \leq M_0(m+1)|t_1 - t_2|.$$

In particular, $\|x^j - x^k\| \leq M_0(m+1)|j - k|\varepsilon$ for all $j, k \in \{0, \dots, m+1\}$.

(ii) For all $j \in \{0, \dots, m\}$,

$$\left\| \frac{x^{j+1} - x^j}{\varepsilon} - v_{x^{j+1}}(u^j) \right\| \leq \frac{M_0 M_1 (m+1)^2 \varepsilon}{2}.$$

(iii) For all $j \in \{0, \dots, m\}$,

$$\|v_{x^{j+1}}(u^j) - v_x(u^j)\| \leq M_0 M_1 (m+1)^3 \varepsilon.$$

Theorem 13. Let $x^0, \dots, x^{m+1} = x$ be as in Algorithm 10, and let $u \in \mathcal{U}$. Let $\lambda_0, \dots, \lambda_m \in \mathbb{R}$ be such that $u = \lambda_0 u^0 + \dots + \lambda_m u^m$ and $\lambda_0 + \dots + \lambda_m = 1$. Then,

$$\left\| v_x(u) - \sum_{j=0}^m \lambda_j \frac{x_{j+1} - x_j}{\varepsilon} \right\| \leq 2M_0 M_1 (m+1)^3 \varepsilon \frac{4m^{\frac{3}{2}} + \delta}{\delta}.$$

For the sake of readability, the proofs of Lemma 12 and Theorem 13 are deferred to Appendix B. From these two results, we conclude that the procedure in lines 5–11 indeed obtains reasonable approximations of the dynamics (1) at point $x = x^{m+1}$. In the next part, we propose how to use these learned dynamics to develop a control law u^* which approximately solves Problem 5.

B. Control

From Theorem 13 it follows that Algorithm 10 can calculate $v_{x_{m+1}}(u) = f(x_{m+1}) + \sum_{i=1}^m g_i(x_{m+1})u_i$ for any value $u \in \mathcal{U}$ with arbitrary precision. Thus, we are able to accurately calculate $G(\phi|_{[0, t_0 + m + 1]\varepsilon}, v_{x^{m+1}}(u))$ for any control value u . Additionally, for a fixed x and ϕ , $u \mapsto G(\phi, v_x(u))$ is a real function of a bounded variable $u \in \mathcal{U}$. Depending on the properties of G , max and argmax of this function can be hence found by an analytic or numerical optimization method; we briefly discuss numerical considerations of the solvability of this optimization problem in Section IX. Thus, the above approximations notwithstanding, we can find an optimal control input u^* to be applied when the previous system trajectory is given by ϕ , and the trajectory is at x .

As shown in Algorithm 10, our plan is to use the found optimal control input u^* for a short time $(m+1)\varepsilon$ after the system is at point x^{m+1} . During that time, the system will learn new local dynamics around the new x^{m+1} , and the whole procedure will be repeated again. We note that this process does not necessarily generate an optimal control law, as the best control input u^* for a certain point x^{m+1} may no longer be the best immediately after the system leaves the said point. However, if the function G is “tame enough”, u^* will still be a good approximation of the optimal control input. We will show that this fact remains true even if we take into account that the input u^* was calculated based on a slightly incorrect learned dynamics at x^{m+1} . We first introduce a measure of tameness of G .

Definition 14. Function $G : \mathcal{F} \times \mathbb{R}^n \rightarrow \mathbb{R}$ has Lipschitz constant L if for all $\phi_1|_{[0, T_1]}, \phi_2|_{[0, T_2]} \in \mathcal{F}$ and $v_1, v_2 \in \mathbb{R}^n$, the following holds:

$$\left| G(\phi_1|_{[0, T_1]}, v_1) - G(\phi_2|_{[0, T_2]}, v_2) \right| \leq L \left(d(\phi_1|_{[0, T_1]}, \phi_2|_{[0, T_2]}) + \|v_1 - v_2\| \right), \quad (7)$$

where

$$d(\phi_1|_{[0,T_1]}, \phi_2|_{[0,T_2]}) = |T_1 - T_2| + \max_{t \in [0, \min(T_1, T_2)]} \|\phi_1(t) - \phi_2(t)\|. \quad (8)$$

Definition 14 explicitly makes note of the usually standard notion of a Lipschitz constant, because d as defined in (8) is not a proper metric on \mathcal{F} . Namely, it does not satisfy the triangle inequality. Thus, the existence of a Lipschitz constant as in Definition 14 does not imply continuity of G in the standard sense.

The claims of the beginning of this section are now formalized by the following result, the proof of which is in Appendix B.

Theorem 15. *Let $\phi_1|_{[0,T_1]}, \phi_2|_{[0,T_2]} \in \mathcal{F}$, and let $\nu > 0$. Define $x = \phi_1(T_1)$ and $y = \phi_2(T_2)$. Assume that G has Lipschitz constant L in the sense of Definition 14, and let u^* be the optimal control at x under approximately learned local dynamics with the previous trajectory $\phi_1|_{[0,T_1]}$, i.e.,*

$$G\left(\phi_1|_{[0,T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i u_i^*\right) = \max_{u \in \mathcal{U}} G\left(\phi_1|_{[0,T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i u_i\right),$$

where $\|\tilde{f} + \sum_{i=1}^m \tilde{g}_i u_i - (f(x) + \sum_{i=1}^m g_i(x)u_i)\| \leq \nu$ for all $u \in \mathcal{U}$. Then,

$$\begin{aligned} & \left| \max_u G(\phi_2|_{[0,T_2]}, v_y(u)) - G(\phi_2|_{[0,T_2]}, v_y(u^*)) \right| \\ & \leq 2Ld(\phi_1|_{[0,T_1]}, \phi_2|_{[0,T_2]}) + 2LM_1(m+1)\|x-y\| + 2L\nu. \end{aligned}$$

Remark 16. *We note that the claim of Theorem 15, as well as all the results that will be presented in the remainder of this paper, can be easily adapted to hold under a relaxed assumption that G has a local Lipschitz constant L , i.e., that (7) is only guaranteed to hold when $d(\phi_1|_{[0,T_1]}, \phi_2|_{[0,T_2]}) < \psi$ for some $\psi > 0$. However, this relaxation brings more technicality into the obtained bounds, and we hence omit further discussion on this topic.*

In the context of Algorithm 10, the result of Theorem 15 can be interpreted in the following way: let us take T_1 to be the time at the beginning of one **repeat** loop in the algorithm, $\phi_1|_{[0,T_1]}$ to be the trajectory of the system until time T_1 , T_2 to be any time in the interval $[T_1, T_1 + (m+1)\varepsilon)$, and $\phi_2|_{[0,T_2]}$ to be the trajectory of the system until time T_2 . Then, Theorem 15 essentially proves that Algorithm 10 results in an almost optimal policy with respect to (3). There is only one major point that remains to be discussed. In Algorithm 10, we do not apply *just* u^* as the control. In order to facilitate learning, we modify this u^* by some small Δu^i . We need to show that such a small control perturbation will not significantly impact the degree of suboptimality of Algorithm 10. Corollary 17 (with the proof in Appendix B) deals with this issue.

Corollary 17. *Assume the same notation as in Theorem 15. Let $\delta > 0$ and let $\tilde{u} \in \mathbb{R}^m$, $\|\tilde{u}\| \leq \delta$. Then,*

$$\begin{aligned} & \left| \max_u G(\phi_2|_{[0,T_2]}, v_y(u)) - G(\phi_2|_{[0,T_2]}, v_y(u^* + \tilde{u})) \right| \\ & \leq 2Ld(\phi_1|_{[0,T_1]}, \phi_2|_{[0,T_2]}) + 2LM_1(m+1)\|x-y\| + 2L\nu + LM_0(m+1)\delta. \end{aligned}$$

Finally, the following result translates the bounds in Theorem 15 and Corollary 17 into a bound in terms of algorithm parameters ε and δ . Thus, Theorem 18, with the proof in Appendix B, finally provides a bound on the

degree of suboptimality of Algorithm 10. For notational purposes, we assume that the run of system (1) is of infinite length. However, the same result clearly holds for any finite length.

Theorem 18. *Let $x_0 \in \mathcal{X}$. Let $u^+ : [0, +\infty) \rightarrow \mathcal{U}$ be the control law used in Algorithm 10. Assume that G has Lipschitz constant L . Then, for all $t \geq (m+1)\varepsilon$,*

$$\left| G \left(\phi_{u^+}(\cdot, x_0)|_{[0,t]}, \frac{d\phi_{u^+}(0+, x)}{dt} \right) - \max_{u \in \mathcal{U}} G \left(\phi_{u^+}(\cdot, x_0)|_{[0,t]}, \frac{d\phi_u(0+, x)}{dt} \right) \right| \leq 6L(M_0 + 1)(M_1 + 1)(m + 1)^3 \left(1 + \frac{4m\sqrt{m}}{\delta} \right) \varepsilon + LM_0(m + 1)\delta, \quad (9)$$

where $x = \phi_{u^+}(t, x_0)$.

Theorem 18 is the central result of the theoretical discussions of this paper. It shows that, for any $\eta > 0$, if ε and δ are chosen so that $6L(M_0 + 1)(M_1 + 1)(m + 1)^3(1 + 4m\sqrt{m}/\delta)\varepsilon + LM_0(m + 1)\delta \leq \eta$, Algorithm 10 will result in a control law that approximately satisfies (3) in Problem 5, with an error no larger than η . As can be seen from the proofs, the bound in Theorem 18 is very rough and can certainly be improved. However, an improvement on this bound is not in the focus of this paper.

Making use of the bound in Theorem 18 requires two elements: the controller's a priori knowledge of constants M_0 , M_1 , and L , and the ability to make $\varepsilon > 0$ and $\delta > 0$ as small as possible. It is difficult to escape the first requirement, although one could potentially think of a heuristic approach to learning M_0 , M_1 , and L on the go. As for the second issue, in order to achieve arbitrarily good performance bounds on Algorithm 10, the control actuators must allow for an arbitrarily good time resolution (i.e., shortest amount of time ε that a control can be applied for) and actuation resolution (i.e., smallest amount of control δ that can be applied). While minimal ε and δ that can be applied will depend on the actuator type, improvement on both of these resolutions is a topic of significant research in the engineering community [37], [38].

Theorem 18 completes this section. Before proceeding to a presentation of simulation results and a discussion of some open questions, let us briefly discuss how Algorithm 10 responds to disturbances.

VII. ROBUSTNESS TO DISTURBANCES

As mentioned in the introduction, a beneficial consequence of the fact that Algorithm 10 learns the system dynamics during the system run is that it is able to easily account for the presence of disturbances. Simply put, with a trivial modification to the algorithm, the disturbance can be learned the same way as the "intended" underlying dynamics. In fact, the algorithm will not even know which part of the dynamics was created by the disturbance. Let us formalize the above discussion.

Consider a disturbance function N which includes both time-varying and state-varying elements, i.e., $N : [0, +\infty) \times \mathcal{X} \rightarrow \mathbb{R}^n$. We assume that N is a Lipschitz continuous function with $\|N(x)\| \leq M_0^*$ and $\|N(x) - N(y)\| \leq M_1^* \|x - y\|$ for all $x, y \in \mathbb{R}^n$. After including this disturbance, dynamics (1) are now replaced by

$$\dot{x} = f(x) + \sum_{i=1}^m g_i(x)u_i + N(t, x). \quad (10)$$

We assume that dynamics given by (10) still evolve on \mathcal{X} . Because of the time-varying element of N , (10) does not automatically fall within the class of systems in (1). However, by appending an additional variable x_{n+1} given by the dynamics $dx_{n+1}/dt = 1$ and $x_{n+1}(0) = 0$, and with an obvious slight abuse of notation, we obtain dynamics

$$\dot{x} = \begin{bmatrix} f(x) + N(x) \\ 1 \end{bmatrix} + \sum_{i=1}^m \begin{bmatrix} g_i(x) \\ 0 \end{bmatrix} u_i. \quad (11)$$

System (11) falls into the control-affine category (1), and we can thus directly apply Algorithm 10. The only changes are in the bounds M_0, M_1 : instead of $\|f\|, \|g_i\| \leq M_0$, it now holds that

$$\left\| \begin{bmatrix} f + N \\ 1 \end{bmatrix} \right\|, \left\| \begin{bmatrix} g_i \\ 0 \end{bmatrix} \right\| \leq \max(M_0 + M_0^*, 1),$$

and, analogously, the Lipschitz constants of the new functions in (11) are bounded by $M_1 + M_1^*$. Thus, if we possess a priori bounds on M_0, M_1, M_0^* , and M_1^* , ε and δ in Algorithm 10 will need to be made smaller in order to guarantee the same degree of suboptimality as in the case without disturbances. With that small caveat, Algorithm 10 performs the same in the presence of disturbances as it does without it. In addition, we note that the same above explanation proves the correctness of Algorithm 10 in the generalized case of non-autonomous control-affine systems

$$\dot{x}(t) = f(t, x) + \sum_{i=1}^m g_i(t, x)u_i,$$

if maps f and g_i are Lipschitz continuous both with respect to time and position.

VIII. SIMULATION RESULTS

A. Damaged Aircraft Example

The simulation work presented in this section is a version of the running example of this paper: a damaged aircraft that is attempting to avoid a crash. We base our model on the linearized dynamics of an undamaged Boeing 747 in landing configuration at sea level [39], [40]. In order to model aircraft damage and introduce nonlinearity into the model, we modify the standard dynamics into the following system:

$$\begin{aligned} \dot{x}_1 &= -0.021x_1 + 0.122x_2 - 0.322x_3 + 0.01u_1 + u_2 \\ \dot{x}_2 &= -0.209x_1 - 0.53x_2 + 2.21x_3 - 0.064u_1 - 0.044u_2 \\ \dot{x}_3 &= 0.017x_1 + 0.01 \cos(x_1)x_1 - 0.164x_2 + 0.15 \sin(x_1)x_2 \\ &\quad - 0.421x_3 - 0.378u_1 + 0.544u_2 + 0.5 \sin(x_2)u_2 \\ \dot{x}_4 &= x_3 \\ \dot{x}_5 &= -x_2 + 2.21x_4, \end{aligned} \quad (12)$$

where $x = (w_l, w_v, q, \theta, h)$ and $u = (\delta_e, \delta_t)$. Variable w_l represents the deviation from the steady-state longitudinal speed of the aircraft, w_v the downwards-pointing vertical speed, q the pitch rate, θ the pitch angle and h the altitude. Additionally, inputs δ_e and δ_t are the deviations from the elevator and thrust values used for level flight.

All values are in feet, seconds and centiradians. The limits on control inputs are set to $u_1 = \delta_e \in [-30, 30]$ and $u_2 = \delta_t \in [-1, 1]$, roughly informed by the descriptions in [41], [42].

System (12) gives the equations of motions that the simulated aircraft will follow. The only difference between (12) and the model of an undamaged aircraft from [39], [40] is in the presence of nonlinear factors, which do not appear in [39], [40]. While we recognize that (12) is not necessarily realistic for any particular type of aircraft damage, and the introduction of trigonometric terms was partly motivated by the desire to showcase how our control strategy handles nonlinearities, dynamics (12) were generally chosen in order to emulate damage to the aircraft's horizontal stabilizer, which results in the pitch rate behaving more erratically. We also remind the reader that our controller design does not start with any a priori knowledge of the system dynamics, and our controller will thus not be making use of any similarities between the undamaged dynamics from [39], [40] and actual dynamics (12).

As mentioned throughout the previous sections, the setting investigated in this example is that the aircraft suffered damaged while in flight, and its main objective is to remain in the air. We now make this scenario more precise: the initial state (i.e., the state at which the abrupt in change in system dynamics occurred) is $(0, 0, 0, 0, 100)$, that is, the aircraft was in level flight at an altitude of 100 feet. The primary goal of the aircraft is:

- (i) $x_5(t) > 0$ for all $t \geq 0$.

We also impose further specifications:

- (ii.a) $x_4(t) \in [-40, 40]$ for all $t \geq 0$ (high pitch angles are unsafe for the aircraft; the bound is informed by [43]),
 - (ii.b) $\dot{x}_5(t) \in [-30, 30]$ for all $t \geq 0$ (high climb rates are unsafe for the aircraft; the bound is informed by [43]),
- and
- (iii) $x_5(t) \in [900, 1100]$ for all $t \geq T$, with $T > 0$ as small as possible (because of the damage and the fact that the aircraft is in landing configuration, it may not be either possible nor smart to move to a high altitude, but it is also not desirable to stay very close to the ground).

While we want to ensure that all above restrictions are satisfied, it is possible that at certain times, the above specifications will clash with each other. We impose the following priority, ranked in descending order: (i), (ii.a) and (ii.b), (iii).

We now design a goodness function corresponding to the above specifications. As we described in Section IV, there is no formally correct way for developing a goodness function. Its design depends primarily on the geometric intuition of the problem. Hence, we first give a rough idea behind the goodness function that we wish to develop:

if $h = x_5 < 100$:	G is such that x_5 quickly increases,
else if $ x_4 > 40$ or $ \dot{x}_5 > 30$:	G is such that $ x_4 $, $ \dot{x}_5 $ or both (as needed) quickly decrease,
else if $x_5 \notin [900, 1100]$:	G is such that x_5 quickly approaches 1000,
else:	G is such that $ x_4 $ and $ \dot{x}_5 $ remain small.

The above informal design remains to be converted into a goodness function G . In particular, based on the above motivation, we need to define six functions that will result in following motions: x_5 quickly increasing, x_5 quickly decreasing, $|x_4|$ quickly decreasing, $|\dot{x}_5|$ quickly decreasing, $|x_4|$ remaining around 0, $|\dot{x}_5|$ remaining around 0.

Defining the above functions would not be difficult if we had full knowledge of the system dynamics (12). However, since the setting of this example deals with an unknown abrupt change in system dynamics, these dynamics are “hidden” from us when designing the goodness function. Instead, the only a priori knowledge about system dynamics comes from basic physical laws and the physical meaning of x_1, \dots, x_5 . In particular, we know the following:

- 1° Increasing $u_1 = \delta_e$ will decrease $x_2 = w_v$ and $x_3 = q$ (by the design of the elevator);
- 2° increasing $u_2 = \delta_t$ will increase $x_1 = w_l$ and $x_3 = q$, and decrease $x_2 = w_v$ (by the function of the aircraft engine controls);
- 3° $\dot{x}_4 = x_3$ (by definition of x_3 and x_4);
- 4° disregarding large deviations in the longitudinal speed, \dot{x}_5 depends on x_2 and x_4 : increasing x_4 will increase \dot{x}_5 , and increasing x_2 will decrease \dot{x}_5 (by definitions of x_2 , x_4 , and x_5 , and the fact that an aircraft will always have a positive longitudinal speed); and
- 5° u_1 and u_2 do not directly increase or decrease x_4 and x_5 , but instead act on them through x_2 and x_3 .

We once again emphasize that our strategy does not require any specific knowledge on the system dynamics (12): facts 1°-5° all follow from physical laws, definitions of the variables, and basic understanding of how aircraft control inputs work.

Let us now go back to designing functions that result in the required six motions. From 5°, we know that we cannot directly increase x_5 through control inputs. Thus, from 4°, in order to increase x_5 (and \dot{x}_5) we know that we want to increase x_4 and decrease x_2 . Since by 5° and 3° we cannot directly influence x_4 through our inputs, but only x_3 , we want to find a control input that results in an increase of x_3 and decrease of x_2 . We do not know exactly which of these two methods we should prioritize, so it makes sense to find an input that maximizes the value $-\dot{x}_2 + \dot{x}_3$. Analogously, in order to decrease \dot{x}_5 and x_5 , we want to maximize $\dot{x}_2 - \dot{x}_3$.

What remains is to find methods for decreasing $|x_4|$ and keeping $|x_4|$ and $|\dot{x}_5|$ around 0. Let us first deal with the tasks pertaining to x_4 . In order to decrease $|x_4|$, we maximize the increase of $x_3 = \dot{x}_4$ in the direction opposite from x_4 , thus providing negative acceleration which will ultimately result in x_4 reducing to 0. In order for $|x_4|$ to remain around 0, we simply use the above method to decrease $|x_4|$ whenever it grows beyond some small number (e.g., $|x_4| > 2$).

By 4°, keeping $|\dot{x}_5|$ around 0 can be accomplished by reducing $|x_4|$, for which we already designed a strategy, and reducing $|x_2|$. We use a similar method as above to reduce $|x_2|$. We first bring the values of x_2 to around 0 by negative acceleration, and then re-apply this method whenever $|x_2|$ grows beyond 2.

Using the above methods, our rough idea of a goodness function G is now formalized as follows:

$$\begin{aligned}
 &\text{if } x_5(t) < 100: && G(x|_{[0,t]}, v) = v_3 - v_2, \\
 &\text{else if } |x_4(t)| > 40 \text{ or if } |\dot{x}_5(t)| > 30 && G(x|_{[0,t]}, v) = m_1 + m_2, \text{ where } m_1 = -v_2 \cdot \text{sign}(x_2(t)) \\
 &\text{or if } x_5(t) \in [900, 1100]: && \text{if } |x_2(t)| > 2 \text{ and } m_1 = 0 \text{ otherwise, and} \\
 & && m_2 = -v_3 \cdot \text{sign}(x_4(t)) \text{ if } |x_4(t)| > 2, \text{ and } m_2 = 0 \text{ otherwise,} \\
 &\text{else:} && G(x|_{[0,t]}, v) = (v_3 - v_2) \cdot \text{sign}(1000 - x_5(t)).
 \end{aligned}$$

Once again, we emphasize that this goodness function is not the only possible one. It is certainly possible to come to an equally reasonable, but different function through different logic. One possible goodness function is described in Example 8, which proposes a continuous goodness function for the reach-avoid problem. However, that function depends on the “mixing” parameter that prioritizes the two objectives, and such a parameter would be difficult to design. On the other hand, our proposed goodness function is simpler to design, but is not continuous. Hence, we are unable to directly use the results of Section VI to determine good choices of cycle length ε and wiggle size δ . Additionally, for computational reasons, instead of applying Algorithm 10 directly as written, we employ it with a minor modification as in Remark 11, where learning and control periods inside one learn-control cycle are decoupled. We chose the length of the learning period to be $\varepsilon' = 10^{-4}$, and the length of the entire learn-control cycle to be 10^{-1} .

The simulation results are presented in Figure 4. From Figure 4, we note that the algorithm, i.e., the myopic control strategy, performs exactly as desired. Within ~ 75 seconds, the aircraft increases its height to roughly 1000 feet. After that, it takes another ~ 75 seconds to safely position its altitude within $[900, 1100]$. From then on, the system stabilizes to an altitude of ~ 960 feet. Additionally, the aircraft’s pitch angle θ and climb rate \dot{h} almost always remain in $[-40, 40]$ and $[-30, 30]$, respectively. While the peaks that result in the trajectory leaving the desired bounds could be avoided by using a different goodness function in the climb period, we did not make such changes in order to emphasize that the control design performs exactly as intended: as soon as the climb rate or angle leave the desired bounds, the controller takes immediate action to return them back into the desired region.

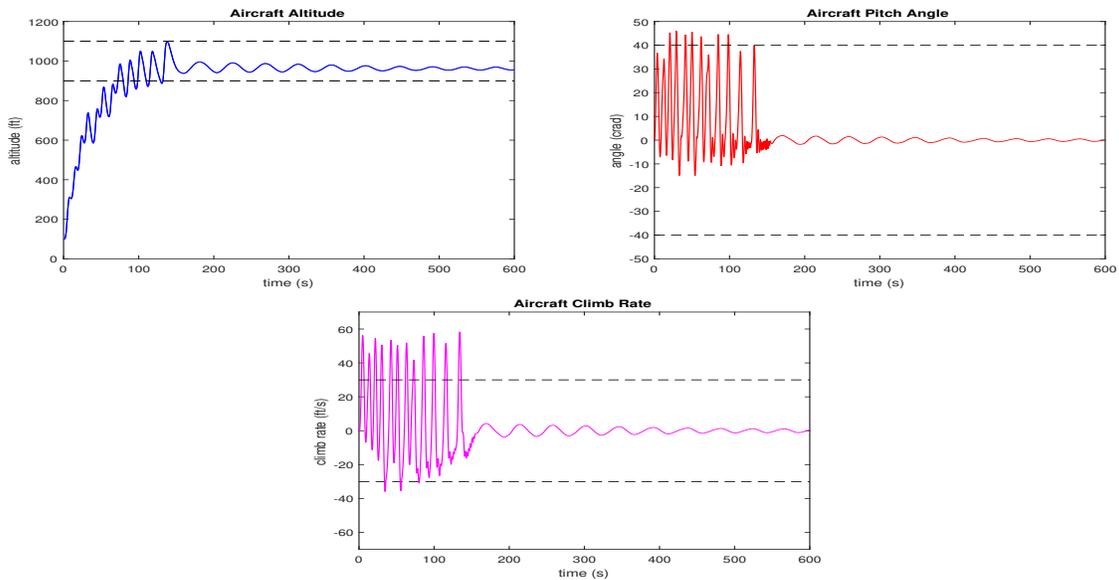


Fig. 4. Graphs of the simulated aircraft trajectory from the setting of Section VIII-A. The aircraft altitude h is given in blue, the pitch angle θ is in red, and the climb rate \dot{h} is in pink. The desired bounds for these values are denoted by dashed lines on the corresponding graph.

We note that the aircraft performs better than intended. In particular, it asymptotically stabilizes to a level flight at a particular altitude, which was neither required by the problem specifications, nor explicitly enforced by the

goodness function. However, it makes natural sense given the system dynamics: when the aircraft is in an almost perfect position (i.e., $|x_2| \leq 2$, $|x_4| \leq 2$, $x_5 \in [900, 1100]$), the goodness function is 0. In that case, the system does not have any preference for the given control. It thus always applies $u = 0$. Then, the system dynamics are given as $\dot{x} = f(x)$, which, in the case of dynamics (12), results in convergence to a local equilibrium. By a smart choice of a goodness function G , the system could be made to asymptotically converge to any altitude. However, for this behavior to be a priori guaranteed, we would need to have significant prior knowledge of dynamics (12).

B. Van der Pol Oscillator

We consider the following control system:

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} = \begin{pmatrix} x_2 \\ -x_1 - 2(1 - x_1^2)x_2 \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} u, \quad (13)$$

with the initial value $x(0) = (1, -2)$ and control input bounds $u \in [-2, 2]$.

Equation (13) models a Van der Pol oscillator with control input. The Van der Pol oscillator is a standard example of a nonlinear control system (see, e.g., [44] for a longer discussion) and is widely used in applications, including human locomotion [45], robotics [46], [47], and laser dynamics [48].

The control objective in this example is to bring the value of x_2 around 0, and control it so that it remains as close as possible to that value. In a sense, this is a ‘‘poor man’s stabilization’’. As myopic control is by its nature not well-suited for asymptotic results, we are not explicitly requiring $x_2(t) \rightarrow 0$ as $t \rightarrow \infty$.

The design of goodness function G proceeds similarly to Section VIII-A. We are not using any particular information on system (13), but we are assuming that it is known that x_2 can be directly influenced by the choice of u . Then, by the same logic as in Example 6 and Section VIII-A, in order to minimize $|x_2|$, it makes sense for G to measure how fast x_2 approaches 0. As in Section VIII-A, this results in a goodness function $G(x|_{[0,t]}, v) = -v_2 \text{sign}(x_2(t))$. However, as all our results in Section VI require G to be continuous, in this example we mollify this function and take

$$G(x|_{[0,t]}, v) = -v_2 \arctan(x_2(t)). \quad (14)$$

Since for every $x_2 \in \mathbb{R} \setminus \{0\}$ the value of $v_2 \arctan(x_2)$ is a positive scalar multiple of $v_2 \text{sign}(x_2)$, this mollification does not change $\text{argmax} G$. The control input which maximizes G will still be given by $u = 2$ if $x_2 < 0$ and $u = -2$ if $x_2 > 0$. However, the penalty for choosing a slightly suboptimal control instead of the optimal one will be lower when x_2 is closer to 0. This makes natural sense: the closer the value of x_2 is to 0, the more satisfied with the current state we are, and we can afford to be more permissive with the direction in which the system is moving.

Before presenting the simulation results, let us briefly explore the performance bounds provided by Section VI. Since the goodness function G from (14) is Lipschitz-continuous, we can apply the results from Section VI. In order to do that, we need bounds on the norm of the system dynamics, as well as Lipschitz constants of system dynamics and goodness function. While the system is technically evolving on all of \mathbb{R}^2 , we will assume that we

are only interested in the dynamics that take place in $\mathcal{X} = [-5, 5] \times [-5, 5]$. Then, it is not difficult to show that $M_0 \leq 250$, $M_1 \leq 99$, $L \leq 29$. By plugging these values into Theorem 18, the myopic suboptimality of the control law generated by Algorithm 10 is bounded by $B_{\varepsilon\delta} \approx 3.5 \cdot 10^7(1 + 4/\delta)\varepsilon + 1.5 \cdot 10^4\delta$. We note that, for a given direction vector $v = \dot{x}(t)$, by (14), x_2 will be moving towards 0 if $G(x|_{[0,t]}, v) > 0$. Thus, as long as $\max_u G(x|_{[0,t]}, v_{x(t)}(u)) > B_{\varepsilon\delta}$ the system will be moving in the right direction.

Since the goodness function G as defined in (14) only depends on $x(t) \in \mathbb{R}^2$ and $v = v(u)$, inequality $\max_u G(x|_{[0,t]}, v_{x(t)}(u)) > B_{\varepsilon\delta}$ defines a region \mathcal{G} of \mathbb{R}^2 in which the system is guaranteed to be progressing in a good direction (i.e., x_2 is approaching 0). However, simple calculations show that in order to guarantee that the trajectory starting from $x_0 = (1, -2)$ will even start moving in the right direction (after the first ε period of time), it is necessary to have $\varepsilon < 10^{-7}$ and $\delta < 10^{-4}$. However, these bounds are derived from Section VI and are very conservative. In fact, simulations show these value can be larger by several orders of magnitude without impacting the performance of the system. Figure 5 shows the simulation results with $\varepsilon = 10^{-4}$ and $\delta = 10^{-3}$.

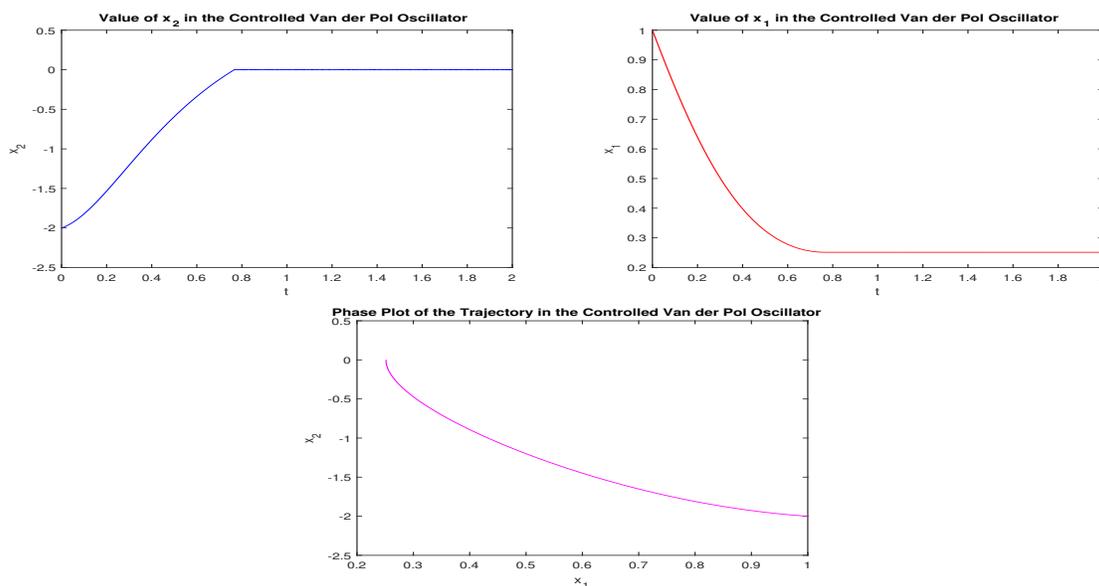


Fig. 5. Graphs of the simulated trajectory from the setting of Section VIII-B. The blue line represents the value of x_2 over time, the red line represents the value of x_1 over time, and the pink line is the phase plot of the trajectory.

As shown in Figure 5, the control strategy performs exactly as desired. The trajectory moves in the right direction from the beginning and reaches $x_2 = 0$. After reaching it, it remains in $|x_2| < 5 \cdot 10^{-3}$, constantly switching the controller in order to remain close to the desired point. We also note that x_1 also remains within 10^{-3} of a particular value. This was not guaranteed nor intended by the design of the goodness function, but is a natural consequence of the fact that x_2 remains around 0 and keeps switching signs. Thus, \dot{x}_1 also remains around 0, and every increase in x_1 is followed by a decrease of roughly the same size.

We once again emphasize that the wiggle size δ and cycle time ε were several orders of magnitude larger than required to guarantee good behaviour by the results of Section VI, and yet, the system works as intended. While

such a result is partly a natural consequence of the liberal bounds on the norms and Lipschitz constants of the system in this example, it additionally shows that the bounds given in Section VI provide much more conservative guarantees than what is obtained in practice.

IX. FUTURE WORK

Before providing a concluding summary of the paper, let us list and briefly discuss some remaining open issues and possible future directions of work based on the contributions of this paper. We generally proceed in the order that the issues appear in the paper.

- The paper makes an assumption on the control-affine structure (1) of the control system. While control-affine systems are a wide and physically interesting class, it is of significant interest to try to generalize the approach of this paper to general nonlinear systems. A naive approach to doing so would be to separate the learning and the control phase of Algorithm 10, along the same lines as in Remark 11. Then, in the learning phase, the algorithm would perform an exhaustive search over the space \mathcal{U} , choose the best control, and apply it for a certain time period in the control phase. After that, the process would be repeated. However, this method would not provide good guarantees for algorithm performance: the number of tested controls in the learning phase, which contains no performance guarantees, would be enormous, and the learning would hence take so long that it could lead to significant performance degradation.

A useful, and seemingly not difficult, extension would be to adapt Algorithm 10 to control systems which, while not control-affine, still have some structure which allows the algorithm to learn control dynamics around a point from only a small number of tested controls. The class of such systems, for example, includes systems polynomial in control (discussed, e.g., in [49]).

- An interesting extension of this work would be to try to account for piecewise-continuity of f, g_i . This setting corresponds to (1) being a hybrid/switched system. Currently, the presented algorithm has no way to detect whether, or make use of the fact that the system made a jump into a different mode of dynamics. The addition of such an ability would potentially also enable identifying which mode of dynamics is the best for a particular control objective.
- While the case of partial observations, i.e., partial knowledge of the system trajectory, is of significant interest for physical systems [50], [51], it is not explored in this paper. A possible approach to that scenario is to use test controls $u^* + \Delta u^0, \dots, u^* + \Delta u^m$ in such a way that, while each control only reveals partial information of the dynamics around a point, the information gathered from all the controls combined allows the controller to reconstruct full local dynamics.
- The primary question still left partly unexplored by this paper concerns the choice of a good function G . We discussed some details on this choice in Section IV. However, that was merely a preliminary exploration. It would be useful to obtain results which provide a formal relationship between a solution to Problem 5 (for a particular function G) and solutions to Problem 1, Problem 3, and other non-myopic control problems. The addition of side knowledge into the decision process by choosing an appropriate function G was briefly discussed in Section IV, and information coming from physical laws was successfully used in the design of a

function G in Section VIII-A. However, no formal work on the use of side information in goodness function design has been presented in this paper.

On the side of lower-level considerations for the choice of function G , we note that in Algorithm 10, we assume that $\operatorname{argmax}_u G(\phi, \tilde{v}(u))$, or at least one element of it, can be found, and that this operation can be done momentarily. However, the validity of this assumption depends on the structure of the function G : if, for instance, G is affine in the second variable, as in Remark 7, the solution of the above optimization problem is trivial. If, however, G is highly nonlinear, determining $\operatorname{argmax}_u G(\phi, \tilde{v}(u))$ can become difficult. Thus, it might make sense to choose G to be, e.g., convex, affine, or quadratic in some sense.

Additionally, we note that the results of Section VI-B are all made under the assumption that G has a Lipschitz constant. While this assumption often makes sense, Example 8 introduces a goodness function which is not continuous in $\phi(t)$. If G does not have a Lipschitz constant, the performance guarantees of Section VI-B would have to be reworked using some other measure of “tameness” of G — for instance, a possible approach may be to use locally bounded variations.

X. CONCLUSION

This paper proposes a method for control of an a priori unknown control-affine system in order to solve reach-avoid-type problems. Because of obvious limitations when controlling a system with unknown dynamics, real-time design of an optimal control law is impossible. Instead, we propose a notion of myopic control, where the optimal control law is one which, at a given time, seems to best satisfy the control specifications. We propose an algorithm that finds a near-optimal solution to the myopic control problem. The algorithm works by simultaneously learning the local system dynamics by applying a series of small controls and using those dynamics to find an optimal control input for a given instance of time, and can produce a control law arbitrarily close to the optimal one. Because the proposed algorithm does not use previously known system dynamics, but learns them on the way, it is highly robust to noise. We provide extensive discussion and simulation results, in particular, for a particular example of control of an aircraft that underwent a catastrophic failure, and, finally, describe future avenues of research based on the proposed algorithm.

APPENDIX A

EXISTENCE OF A MYOPICALLY-OPTIMAL CONTROL LAW

Let $\tilde{\mathcal{S}} : \mathcal{F} \rightarrow 2^{\mathcal{U}}$ be a set-valued map defined by $\tilde{\mathcal{S}}(\phi|_{[0,t]}) = \operatorname{argmax}_{u \in \mathcal{U}} G(\phi|_{[0,t]}, v_{\phi(t)}(u))$. Then, Problem 5 has a solution if and only if the following problem has a solution:

$$\begin{aligned} \dot{x}(t) &\in \left\{ f(x(t)) + \sum_{i=1}^m g_i(x(t)) u_i \mid u \in \tilde{\mathcal{S}}(x|_{[0,t]}) \right\}, \\ x(0) &= x_0. \end{aligned} \tag{15}$$

If we define a set-valued map $S : \mathcal{F} \rightarrow 2^{\mathbb{R}^n}$ as the right-hand side of the differential inclusion in (15), the existence of solution to Problem 5 can thus be interpreted in terms of the existence of a solution to

$$\begin{aligned} \dot{x}(t) &\in S(x|_{[0,t]}), \\ x(0) &= x_0, \end{aligned} \tag{16}$$

that is, a delay differential inclusion in the sense of [52].

Differential inclusion (16) differs from the setting covered by the bulk of research in delay differential inclusions (e.g., [53], [54], [55]), as the delay (i.e., the length of time interval on which the right side depends) in (16) is not fixed. Nonetheless, some sufficient conditions for the existence of a solution to (16) have been previously identified. In particular, Theorem 21 presents an adaptation of the central result of [56]. As the work in [56] only deals with $\mathcal{X} = \mathbb{R}^n$, we make that assumption in the remainder of this section. In the context of Theorem 21, we define the norm on $C([0, t], \mathbb{R}^n)$ by $\|\phi\| = \max_{\tau \in [0, t]} \|\phi(\tau)\|$. We additionally remind the reader of the following definitions.

Definition 19. *Let \mathcal{A}, \mathcal{B} be topological spaces, and let $f : \mathcal{A} \rightarrow 2^{\mathcal{B}}$. Map f is upper hemicontinuous at a point $a \in \mathcal{A}$ if for any open set $O_{\mathcal{B}} \supseteq f(a)$ in \mathcal{B} , there exists an open set $O_{\mathcal{A}} \ni a$ in \mathcal{A} such that $f(a') \subseteq O_{\mathcal{B}}$ for all $a' \in O_{\mathcal{A}}$. Map f is upper hemicontinuous if it is upper hemicontinuous at all points $a \in \mathcal{A}$.*

Definition 20. *Let \mathcal{B} be a topological space with a Lebesgue measure, and $I \subseteq \mathbb{R}$ an interval. A map $f : I \rightarrow 2^{\mathcal{B}}$ is Lebesgue measurable if for every closed set $C \subseteq \mathcal{B}$, the set $\{t \in I \mid f(t) \cap C \neq \emptyset\} \subseteq I$ is Lebesgue measurable.*

We are now ready to present a sufficient condition for the existence of a solution to Problem 5.

Theorem 21. *Let $\gamma > 0$, and let $\mathcal{X} = \mathbb{R}^n$. Assume that all of the following hold:*

- (i) $S(\phi|_{[0,t]}) \subseteq \mathbb{R}^n$ is a convex set for all $t \in [0, \gamma]$ and all continuous maps $\phi|_{[0,t]} : [0, t] \rightarrow \mathcal{X}$.
- (ii) For any fixed $t \in [0, \gamma]$, the restriction of the map S to $C([0, t], \mathcal{X})$ is upper hemicontinuous.
- (iii) For any fixed $\phi \in C([0, \gamma], \mathcal{X})$, the map $S_{\phi} : [0, \gamma] \rightarrow 2^{\mathbb{R}^n}$ defined by $S_{\phi}(t) = S(\phi|_{[0,t]})$ is Lebesgue-measurable.
- (iv) There exists $k > 0$ such that, for every absolutely continuous function $\phi \in C([0, \gamma], \mathcal{X})$, every $t \in [0, \gamma]$ such that $d\phi(t)/dt$ exists, and every $y \in S(\phi|_{[0,t]})$, it holds that $y \cdot d\phi(t)/dt \leq k(1 + \|\phi|_{[0,t]}\|^2)$.

Then, (16) has a solution on the interval $[0, \gamma]$.

A slightly more permissive, but significantly less elegant, sufficient condition for the existence of a solution to (16) has been identified in [52]; we direct the reader to the details of that paper for more information. We forgo the proof of Theorem 21, as it follows immediately from the main result in [56]. We note that the statement of Theorem 21 is somewhat simpler than the original theorem statement of [56], as one of the original conditions is automatically satisfied because \mathcal{U} is assumed to be bounded.

Remark 22. *It can be shown that the goodness function from the example of the Van der Pol oscillator in Section VIII-B does not satisfy the sufficient condition of Theorem 21. Nevertheless, as can be seen from Section VIII-B, this does not prevent Algorithm 10 from constructing a control law arbitrarily close to the myopically-optimal one.*

APPENDIX B
PROOFS OF PERFORMANCE BOUNDS

Proof of Lemma 12.

(i) We note that, for any control law u ,

$$\begin{aligned} \|\phi_u(t_1, x^0) - \phi_u(t_2, x^0)\| &= \left\| \int_{t_1}^{t_2} f(\phi_u(t, x^0)) + \sum_{i=1}^m g_i(\phi_u(t, x^0))u_i dt \right\| \\ &\leq \int_{t_1}^{t_2} \left\| f(\phi_u(t, x^0)) + \sum_{i=1}^m g_i(\phi_u(t, x^0))u_i \right\| dt \leq \int_{t_1}^{t_2} M_0(m+1) dt = M_0(m+1)|t_2 - t_1|, \end{aligned}$$

where the last inequality follows from the triangle inequality and because we took $\mathcal{U} = [-1, 1]^m$.

The claim for x^j 's follows immediately from the definition: $x^j = \phi_u(j\varepsilon, x^0)$, where u is the control law from Algorithm 10.

(ii) Let $j \in \{0, \dots, m\}$. Then,

$$\begin{aligned} \left\| \frac{x^{j+1} - x^j}{\varepsilon} - v_{x^{j+1}}(u^j) \right\| &\leq \frac{1}{\varepsilon} \left\| \int_0^\varepsilon \left(f(\phi_{u^j}(t, x^j)) + \sum_{i=1}^m g_i(\phi_{u^j}(t, x^j))u_i^j - v_{x^{j+1}}(u^j) \right) dt \right\| \\ &\leq \frac{1}{\varepsilon} \int_0^\varepsilon \left(\|f(\phi_{u^j}(t, x^j)) - f(x^{j+1})\| + \sum_{i=1}^m \left\| g_i(\phi_{u^j}(t, x^j))u_i^j - \sum_{i=1}^m g_i(x^{j+1})u_i^j \right\| \right) dt \\ &\leq \frac{1}{\varepsilon} \int_0^\varepsilon (m+1)M_1 \|\phi_{u^j}(t, x^j) - x^{j+1}\| dt \leq \frac{1}{\varepsilon} \int_0^\varepsilon M_0M_1(m+1)^2 t dt = \frac{M_0M_1(m+1)^2\varepsilon}{2}, \end{aligned}$$

where we use the definition of $v_{x^{j+1}}(u^j)$ and the fact that Lipschitz constants of f, g_i are no larger than M_1 .

The last inequality follows from (i).

(iii) Let $j \in \{0, \dots, m\}$. We note that $\|v_{x^{j+1}}(u^j) - v_x(x^j)\| \leq \|f(x^{j+1}) - f(x^j)\| + \sum_{i=1}^m \|(g_i(x^{j+1}) - g_i(x^j))u_i^j\| \leq (m+1)M_1\|x^{j+1} - x^j\| \leq M_0M_1(m+1)^3\varepsilon$, where the last two inequalities follow from the bounds on Lipschitz constants of f, g_i , and from (i), respectively. □

Proof of Theorem 13. From $u = \sum \lambda_j u^j = \sum \lambda_j (u^* + \Delta u^j)$ and $\sum \lambda_j = 1$ we obtain $u = u^* + \lambda_1 \Delta u^1 + \dots + \lambda_m \Delta u^m$. By definition of Δu^j from line 5 of Algorithm 10, we obtain $(u - u^*)_j = \pm \lambda_j \delta$. Hence, $|\pm \lambda_j \delta| < \|u - u^*\|$ for all $j \geq 1$. Since $u, u^* \in \mathcal{U} = [-1, 1]^m$, we obtain $|\lambda_j| \leq 2\sqrt{m}/\delta$ for all $j \geq 1$. By using $\sum \lambda_j = 1$ and triangle inequality, we get $|\lambda_0| \leq 1 + 2m\sqrt{m}/\delta$.

Now, we noted before that $v_x(u) = \sum \lambda_j v_x(u^j)$. Hence, we have $\|v_x(u) - \sum \lambda_j (x^{j+1} - x^j)/\varepsilon\| \leq \sum |\lambda_j| \|v_x(u^j) - (x^{j+1} - x^j)/\varepsilon\|$. Combining parts (ii) and (iii) of Lemma 12, the right hand side of the previous inequality can be bounded by $\sum |\lambda_j| M_0M_1((m+1)^2/2 + (m+1)^3)\varepsilon \leq \sum |\lambda_j| 2M_0M_1(m+1)^3\varepsilon \leq (1 + 4m\sqrt{m}/\delta) 2M_0M_1(m+1)^3\varepsilon$, where the last inequality was obtained by using bounds on λ_j from the previous paragraph. □

Proof of Theorem 15. Let \bar{u} be the optimal control under dynamics (1) at y , i.e.,

$$G(\phi_2|_{[0, T_2]}, v_y(\bar{u})) = \max_{u \in \mathcal{U}} G(\phi_2|_{[0, T_2]}, v_y(u)).$$

Then,

$$\begin{aligned} & \left| G(\phi_2|_{[0, T_2]}, v_y(\bar{u})) - G(\phi_2|_{[0, T_2]}, v_y(u^*)) \right| \\ & \leq \left| G(\phi_2|_{[0, T_2]}, v_y(\bar{u})) - G\left(\phi_1|_{[0, T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i u_i^*\right) \right| + \left| G\left(\phi_1|_{[0, T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i u_i^*\right) - G(\phi_2|_{[0, T_2]}, v_y(u^*)) \right|. \end{aligned} \quad (17)$$

We first claim that

$$\left| G(\phi_2|_{[0, T_2]}, v_y(\bar{u})) - G\left(\phi_1|_{[0, T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i u_i^*\right) \right| \leq Ld(\phi_1|_{[0, T_1]}, \phi_2|_{[0, T_2]}) + L(m+1)M_1\|x-y\| + L\nu. \quad (18)$$

We note that $\|v_y(\bar{u}) - \tilde{f} - \sum_{i=1}^m \tilde{g}_i \bar{u}_i\| \leq \|v_y(\bar{u}) - v_x(\bar{u})\| + \|v_x(\bar{u}) - \tilde{f} - \sum_{i=1}^m \tilde{g}_i \bar{u}_i\| \leq (m+1)M_1\|y-x\| + \nu$.

Hence,

$$\begin{aligned} & \left| G(\phi_2|_{[0, T_2]}, v_y(\bar{u})) - G\left(\phi_1|_{[0, T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i \bar{u}_i\right) \right| \\ & \leq L\left(d(\phi_1|_{[0, T_1]}, \phi_2|_{[0, T_2]}) + \left\| v_y(\bar{u}) - \tilde{f} - \sum_{i=1}^m \tilde{g}_i \bar{u}_i \right\| \right) \\ & \leq Ld(\phi_1|_{[0, T_1]}, \phi_2|_{[0, T_2]}) + LM_1(m+1)\|x-y\| + L\nu. \end{aligned} \quad (19)$$

Analogously to (19),

$$\left| G\left(\phi_1|_{[0, T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i u_i^*\right) - G(\phi_2|_{[0, T_2]}, v_y(u^*)) \right| \leq Ld(\phi_1|_{[0, T_1]}, \phi_2|_{[0, T_2]}) + L(m+1)M_1\|x-y\| + L\nu. \quad (20)$$

Now, if (18) was incorrect, then, by combining its negation with either (19) or (20), we obtain that $G(\phi_1|_{[0, T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i \bar{u}_i) > G(\phi_1|_{[0, T_1]}, \tilde{f} + \sum_{i=1}^m \tilde{g}_i u_i^*)$ or $G(\phi_2|_{[0, T_2]}, v_y(u^*)) > G(\phi_2|_{[0, T_2]}, v_y(\bar{u}))$, which are both contradictions with the definitions of u^* and \bar{u} , respectively. Thus, (18) holds, and from (17), (18), and (20), we obtain the theorem claim. \square

Proof of Corollary 17. Using the same notation as in the proof of Theorem 15, we obtain $|G(\phi_2|_{[0, T_2]}, v_y(\bar{u})) - G(\phi_2|_{[0, T_2]}, v_y(u^* + \tilde{u}))| \leq |G(\phi_2|_{[0, T_2]}, v_y(\bar{u})) - G(\phi_2|_{[0, T_2]}, v_y(u^*))| + |G(\phi_2|_{[0, T_2]}, v_y(u^*)) - G(\phi_2|_{[0, T_2]}, v_y(u^* + \tilde{u}))|$. From Theorem 15, we know that the first summand on the right hand side of the above inequality is bounded by $2Ld(\phi_1|_{[0, T_1]}, \phi_2|_{[0, T_2]}) + 2LM_1(m+1)\|x-y\| + 2L\nu$. By the definition of $v_y(\cdot)$ and the Lipschitz constant on G , the second summand is bounded by $L\|\sum_{i=1}^m \tilde{g}_i(y)\tilde{u}_i\| \leq LM_0m\delta < LM_0(m+1)\delta$. We thus obtain the desired bound. \square

Proof of Theorem 18. Let us take any $t \geq (m+1)\varepsilon$. Time t is certainly contained in one iteration of the **repeat** loop in lines 3–14 of Algorithm 10, and since $t \geq (m+1)\varepsilon$, this **repeat** loop is not the first. Let x^0, \dots, x^{m+1} be the x^j 's used in that iteration of the **repeat** loop. From Theorem 13, we note that the previous iteration resulted in a control input $u^* \in \mathcal{U}$ which is near-optimal at x^0 , that is, u^* is optimal for dynamics $u \rightarrow \tilde{f} + \sum \tilde{g}_i u_i$, with $\|v_{x^0}(u) - (\tilde{f} + \sum \tilde{g}_i u_i)\| \leq 2M_0M_1(m+1)^3(1 + 4m^{3/2}/\delta)\varepsilon$ for all $u \in \mathcal{U}$.

Now, we note that, at time t , $u^+(t) = u^* + \Delta u^j$, for some $j \in \{0, \dots, m\}$ and $\|\Delta u^j\| \leq \delta$. Thus, we can apply Corollary 17, with $\nu = 2M_0M_1(m+1)^3(1+4m^{3/2}/\delta)\varepsilon$. We hence obtain

$$\begin{aligned} & \left| \max_u G(\phi_{u^+}(\cdot, x_0)|_{[0,t]}, v_x(u)) - G(\phi_{u^+}(\cdot, x_0)|_{[0,t]}, v_x(u^+)) \right| \\ & \leq 2Ld(\phi_{u^+}(\cdot, x_0)|_{[0,t]}, \phi_{u^+}(\cdot, x_0)|_{[0,t_0]}) + 2LM_1(m+1)\|x - x^0\| \\ & \quad + 4LM_0M_1(m+1)^3(1+4m^{3/2}/\delta)\varepsilon + LM_0(m+1)\delta, \end{aligned} \quad (21)$$

where t_0 is the time of the beginning of the current iteration of the repeat loop, i.e., the time when the system was at x^0 .

We note that by (8), $d(\phi_{u^+}(\cdot, x_0)|_{[0,t]}, \phi_{u^+}(\cdot, x_0)|_{[0,t_0]}) = t - t_0 \in [0, (m+1)\varepsilon)$. Additionally, we note that $\|x - x^0\| = \|\phi_{u^+}(t', x^0) - \phi_{u^+}(0, x^0)\|$ for $t' = t - t_0 \in [0, (m+1)\varepsilon)$. Thus, by Part (i) of Lemma 12, we obtain $\|x - x^0\| \leq M_0(m+1)^2\varepsilon$. Hence, from (21), and taking some very liberal bounds, it now follows that $|\max_u G(x, v_x(u)) - G(x, v_x(u^+))| \leq 6L(M_0+1)(M_1+1)(m+1)^3(1+4m^{3/2}/\delta)\varepsilon + LM_0(m+1)\delta$. The claim of the theorem follows by noting that $d\phi_{u^+}(0+, x)/dt = f(x) + \sum g_i(x)u_i = v_x(u)$. \square

ACKNOWLEDGEMENTS

This research was partly funded by grants FA8650-15-C-2546 and FA8650-16-C-2610 from the Air Force Research Laboratory (AFRL), grant W911NF-16-1-0001 from the Defense Advanced Research Projects Agency (DARPA), and grant W911NF-15-1-0592 from the Army Research Office (ARO). The authors thank Mohamadreza Ahmadi for his valuable comments in improving this manuscript, as well as for his suggestion of a Van der Pol oscillator as an example in Section VIII, and Srilakshmi Pattabiraman for noticing a slight mistake in the previous version of this manuscript.

REFERENCES

- [1] M. Ornik, A. Israel, and U. Topcu, "Myopic control of systems with unknown dynamics," 2017, submitted to the 2018 American Control Conference.
- [2] S. Aloni, *Israeli F-15 Eagle Units in Combat*. Osprey Publishing, 2006.
- [3] S. L. Brunton, J. L. Proctor, and J. N. Kutz, "Discovering governing equations from data by sparse identification of nonlinear dynamical systems," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 113, no. 15, pp. 3932–3937, 2016.
- [4] T. S. Cubitt, J. Eisert, and M. M. Wolf, "Extracting dynamical equations from experimental data is NP hard," *Phys. Rev. Lett.*, vol. 108, 2012.
- [5] D. J. A. Hills, A. M. Grütter, and J. J. Hudson, "An algorithm for discovering Lagrangians automatically from data," *PeerJ Comput. Sci.*, vol. 1, 2015.
- [6] M. Schmidt and H. Lipson, "Distilling free-form natural laws from experimental data," *Science*, vol. 324, no. 5923, pp. 81–85, 2009.
- [7] D. P. Solomatine and A. Ostfeld, "Data-driven modelling: some past experiences and new approaches," *J. Hydroinf.*, vol. 10, no. 1, pp. 3–22, 2008.
- [8] M. Ahmadi, A. Israel, and U. Topcu, "Safety assessment based on physically-viable data-driven models," in *56th IEEE Conference on Decision and Control*, 2017, accepted.
- [9] R. P. Copeland and K. S. Rattan, "A fuzzy logic supervisor for PID control of unknown systems," in *IEEE International Symposium on Intelligent Control*, 1994, pp. 22–26.
- [10] S. Khadraoui, H. N. Nounou, M. N. Nounou, A. Datta, and S. P. Bhattacharyya, "Adaptive controller design for unknown systems using measured data," *Asian J. Control*, vol. 18, no. 4, pp. 1453–1466, 2016.

- [11] J. Na and G. Herrmann, "Online adaptive approximate optimal tracking control with simplified dual approximation structure for continuous-time unknown nonlinear systems," *IEEE/CAA J. Autom. Sin.*, vol. 1, no. 4, pp. 412–422, 2014.
- [12] H.-H. Yeh, S. S. Banda, and P. J. Lynch, "Control of unknown systems via deconvolution," *Dyn. Control*, vol. 13, no. 3, pp. 416–423, 1990.
- [13] Z. Zhou, R. Takei, H. Huang, and C. J. Tomlin, "A general, open-loop formulation for reach-avoid games," in *51st IEEE Conference on Decision and Control*, 2012, pp. 6501–6506.
- [14] D. Joyce, "On manifolds with corners," in *Advances in Geometric Analysis*, S. Janeczko, J. Li, and D. H. Phong, Eds. International Press, 2012, pp. 225–258.
- [15] J. T. Bosworth, "Linearized aerodynamic and control law models of the X-29A airplane and comparison with flight data," National Aeronautics and Space Administration, Tech. Rep. NASA TM-4356, 1992.
- [16] E. L. Duke, R. F. Antoniewicz, and K. D. Krambeer, "Derivation and definition of a linear aircraft model," National Aeronautics and Space Administration, Tech. Rep. NASA RP-1207, 1988.
- [17] B. L. Stevens, F. L. Lewis, and E. N. Johnson, *Aircraft Control and Simulation: Dynamics, Controls Design, and Autonomous Systems*. Wiley, 2016.
- [18] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [19] Y. Naveh, P. Z. Bar-Yoseph, and Y. Halevi, "Nonlinear modeling and control of a unicycle," *Dyn. Control*, vol. 9, no. 4, pp. 279–296, 1999.
- [20] A. Mokhtari, A. Benallegue, and Y. Orlov, "Exact linearization and sliding mode observer for a quadrotor unmanned aerial vehicle," *Int. J. Rob. Autom.*, vol. 21, no. 1, pp. 39–49, 2006.
- [21] W. L. Garrard and J. M. Jordan, "Design of nonlinear automatic flight control systems," *Automatica*, vol. 13, no. 5, pp. 497–505, 1977.
- [22] F. Bullo and A. D. Lewis, *Geometric Control of Mechanical Systems*. Springer, 2004.
- [23] A. Isidori, *Nonlinear Control Systems*. Springer, 1995.
- [24] E. D. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer, 1998.
- [25] F. Fadaie and M. E. Broucke, "A viability problem for control affine systems with application to collision avoidance," in *45th IEEE Conference on Decision and Control*, 2006, pp. 5998–6003.
- [26] R. C. Dorf and R. H. Bishop, *Modern Control Systems*. Prentice Hall, 2011.
- [27] H. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*. Wiley, 1972.
- [28] P. M. Esfahani, D. Chatterjee, and J. Lygeros, "The stochastic reach-avoid problem and set characterization for diffusions," *Automatica*, vol. 70, pp. 43–56, 2016.
- [29] W. S. Levine, *The Control Systems Handbook, Second Edition: Control System Advanced Methods*. CRC Press, 2011.
- [30] W. B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley, 2011.
- [31] R. D. Rimey and C. M. Brown, "Control of selective perception using Bayes nets and decision theory," *Int. J. Comput. Vision*, vol. 12, no. 2, pp. 173–207, 1994.
- [32] S. Zhang and A. J. Yu, "Forgetful Bayes and myopic planning: Human learning and decision-making in a bandit setting," in *Neural Information Processing Systems Conference*, 2013, pp. 2607–2615.
- [33] I. Alvarez and S. Martin, "Geometric robustness of viability kernels and resilience basins," in *Viability and Resilience of Complex Systems*, G. Deffuant and N. Gilbert, Eds. Springer, 2011, pp. 193–218.
- [34] Z. Manna and A. Pnueli, *The Temporal Logic of Reactive and Concurrent Systems*. Springer, 1992.
- [35] I. Papusha, J. Fu, U. Topcu, and R. M. Murray, "Automata theory meets approximate dynamic programming: Optimal control with temporal logic constraints," in *55th IEEE Conference on Decision and Control*, 2016, pp. 434–440.
- [36] A. R. Ansari and T. D. Murphey, "Sequential action control: Closed-form optimal control for nonlinear and nonsmooth systems," *IEEE Trans. Rob.*, vol. 32, no. 5, pp. 1196–1214, 2016.
- [37] H.-A. Eckel, S. Scharring, S. Karg, C. Illg, and J. Peter, "Overview of laser ablation micropropulsion research activities at DLR Stuttgart," in *International Symposium on High Power Laser Ablation and Beamed Energy Propulsion*, 2014.
- [38] D. Krejci, F. Mier-Hicks, R. Thomas, T. Haag, and P. Lozano, "Emission characteristics of passively fed electrospray microthrusters with propellant reservoirs," *J. Spacecr. Rockets*, vol. 54, no. 2, pp. 447–458, 2017.
- [39] A. E. Bryson, *Control of Spacecraft and Aircraft*. Princeton University Press, 1994.
- [40] J. Che and D. Chen, "Automatic landing control using H_∞ control and stable inversion," in *40th IEEE Conference on Decision and Control*, 2001, pp. 241–246.

- [41] S. Ganguli, A. Marcos, and G. Balas, "Reconfigurable LPV control design for Boeing 747-100/200 longitudinal axis," in *American Control Conference*, 2002, pp. 3612–3617.
- [42] H. J. Smith, "A flight test investigation of the rolling moments induced on a T-37B airplane in the wake of a B-747 airplane," National Aeronautics and Space Administration, Tech. Rep. TM X-56031, 1975.
- [43] S. S. Mulgund and R. F. Stengel, "Target pitch angle for the microburst escape maneuver," *J. Aircr.*, vol. 30, no. 6, pp. 826–832, 1993.
- [44] M. Vidyasagar, *Nonlinear Systems Analysis*. SIAM, 2002.
- [45] M. S. Dutra, A. C. de Pina Filho, and V. F. Romano, "Modeling of a bipedal locomotor using coupled nonlinear oscillators of Van der Pol," *Biol. Cybern.*, vol. 88, no. 4, pp. 286–292, 2003.
- [46] R. Héliot and B. Espiau, "Online generation of cyclic leg trajectories synchronized with sensor measurement," *Rob. Auton. Syst.*, vol. 56, no. 5, pp. 410–421, 2008.
- [47] P. Veskos and Y. Demiris, "Developmental acquisition of entrainment skills in robot swinging using van der Pol oscillators," in *Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, 2005, pp. 87–93.
- [48] S. Wirkus and R. Rand, "The dynamics of two coupled van der Pol oscillators with delay coupling," *Nonlinear Dyn.*, vol. 30, no. 3, pp. 205–221, 2002.
- [49] E. Moulay and W. Perruquetti, "Stabilization of nonaffine systems: A constructive method for polynomial systems," *IEEE Trans. Autom. Control*, vol. 50, no. 4, pp. 520–526, 2005.
- [50] W. H. Fleming and R. W. Rishel, *Deterministic and Stochastic Optimal Control*. Springer, 1975.
- [51] M. Mariton, *Jump Linear Systems in Automatic Control*. CRC Press, 1990.
- [52] L. Boudjenah, "Existence of solutions to differential inclusions with delayed arguments," *Electr. J. Differ. Equ.*, vol. 2010, no. 175, pp. 1–8, 2010.
- [53] D. A. Carlson, "Asymptotic stability for a class of delay differential inclusions," *J. Differ. Equ.*, vol. 139, no. 2, pp. 219–236, 1997.
- [54] J. Henderson and A. Ouahab, "Fractional functional differential inclusions with finite delay," *Nonlinear Anal. Theory Methods Appl.*, vol. 70, no. 5, pp. 2091–2105, 2009.
- [55] K.-Z. Liu, X.-M. Sun, J. Liu, and A. R. Teel, "Stability theorems for delay differential inclusions," *IEEE Trans. Autom. Control*, vol. 61, no. 10, pp. 3215–3220, 2016.
- [56] B. I. Anan'ev, "Teorema sushchestovovaniya dlya differentsial'nogo vklyucheniya s peremenn'im zapazd'ivaniem [An existence theorem for a differential inclusion with variable lag]," *Differentsial'nye uravneniya*, vol. 11, no. 7, pp. 1155–1158, 1975.