

Data-Assisted Vision-Based Hybrid Control for Robust Stabilization with Obstacle Avoidance via Learning of Perception Maps

Alejandro Murillo-González, Jorge I. Poveda,

Abstract—We study the problem of target stabilization with robust obstacle avoidance in robots and vehicles that have access only to vision-based sensors for the purpose of real-time localization. This problem is particularly challenging due to the topological obstructions induced by the obstacle, which preclude the existence of smooth feedback controllers able to achieve simultaneous stabilization and robust obstacle avoidance. To overcome this issue, we develop a vision-based hybrid controller that switches between two different feedback laws depending on the current position of the vehicle using a hysteresis mechanism and a data-assisted supervisor. The main innovation of the paper is the incorporation of suitable *perception maps* into the hybrid controller. These maps can be learned from data obtained from cameras in the vehicles and trained via convolutional neural networks (CNN). Under suitable assumptions on this perception map, we establish theoretical guarantees for the trajectories of the vehicle in terms of convergence and obstacle avoidance. Moreover, the proposed vision-based hybrid controller is numerically tested under different scenarios, including noisy data, sensors with failures, and cameras with occlusions.

I. INTRODUCTION

During recent years, there has been an increasing number of works on systems that integrate high-dimensional inputs, such as images, into feedback control loops. For example, several successful end-to-end approaches have employed reinforcement learning (RL), including [1], where the state-space construction is automated by learning a state representation directly from camera images. Also, in [2] the authors introduced deep Q-networks into a control command, achieving approximate human-level performance. Other works have used deep generative models to synthesize controllers with inputs coming from an embedding space of high-dimensional data, which does not necessarily correspond to an interpretable space (e.g., joint coordinates of the robot). Some examples in this direction include [3]–[7]. In other works, such as [5], [8] and [9], the authors integrated state predictions via robust control tools to handle approximation errors.

While significant progress has been made in different communities during the last years, most of the results in the literature have focused on applications where the feedback control law $\alpha(\cdot)$ leads to closed-loop systems of the form

$$\dot{x} = f(x + e_1, u) + e_2, \quad u = \alpha(x + e_3) + e_4, \quad (1)$$

where f and α are continuous functions, and the signals e_i , for $i \in \{1, 2, 3, 4\}$, model measurement noise, implementation errors, or approximation inaccuracies induced by learning mechanisms such as linear parametric approximations, neural networks, multi-time scale techniques, etc. For these perturbed dynamical systems, stability and robustness results are well established, and they can be characterized via practical or input-to-state stability results [10], [11].

On the other hand, many robust control problems cannot be solved via smooth dynamical systems of the form (1). Typical examples include robust global stabilization problems on smooth compact manifolds [12], global stabilization of a disconnected set [13], the asymptotic stabilization of vehicles with geometric constraints [11], the robust control of switched systems [14], and the *robust stabilization of targets in obstacle avoidance problems* [15], to name just a few. In the latter problem, the objective is to robustly stabilize a target point in spaces with global obstacles, i.e., the operational space, or “free world”, is a strict subset of $\mathbb{R}^n$. In this setting, global stabilization using smooth feedback of the form (1) is precluded by the fact that the domain of attraction of an asymptotically stable vector field (i.e., the operational space) must be diffeomorphic to the Euclidean space, a condition that is not satisfied under global obstacles [11]. Given that discontinuous controllers have also been shown to suffer from fundamental robustness limitations [12], most works have focused on achieving local or almost global convergence results [16]–[20], which exclude from the basin of attraction a particular set of measure zero. On the other hand, the impossibility result for smooth controllers has also triggered an active line of research on the development of hybrid control techniques able to achieve robust *global* stabilization and obstacle avoidance, e.g., [13], [15], [21], [22], [23]. However, unlike smooth dynamical systems of the form (1), establishing suitable robustness guarantees for hybrid controllers is far from trivial, which motivates current research on the integration and analysis of learning-based mechanisms into these types of systems. In particular, to the best knowledge of the authors, the systematic integration of hybrid control and data-assisted vision-based mechanisms for robust stabilization and obstacle avoidance has remained an open problem.

Contributions: In this work, we develop a *vision-based* hybrid controller for robust and resilient obstacle avoidance in mobile robots. We show that, unlike standard smooth feedback controllers, the proposed hybrid algorithm can overcome arbitrarily small and potentially adversarial disturbances, noisy states, sensor failures, as well as camera

Alejandro Murillo-González is with the Department of Mathematical Sciences at Universidad EAFIT, Colombia. Email {amurillo@eafit.edu.co}

Jorge I. Poveda is with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO, 80309. Email:{jorge.poveda@colorado.edu}.

occlusions. Our approach synergistically leverages robust hybrid control theory [14] and recent results in perception-based control [6], [8], which have studied the incorporation of *perception maps* learned from data to predict the states and dynamics of the system. The proposed hybrid controllers are suitable for vehicles with *vision* sensors, such as cameras, that have access to historical data in order to learn a suitable perception map via convolutional neural networks (CNNs). Our main results provide theoretical guarantees, as well as extensive numerical validations in different scenarios.

II. PRELIMINARIES

Given a compact set $\mathcal{A} \subset \mathbb{R}^n$ and an arbitrary vector $z \in \mathbb{R}^n$, we use $|z|_{\mathcal{A}} := \min_{s \in \mathcal{A}} \|z - s\|_2$ to denote the minimum distance from z to $\mathcal{A}$. We use (x, y) to denote the concatenation of the vectors x and y . A set-valued mapping $M : \mathbb{R}^p \rightrightarrows \mathbb{R}^n$ is said to be: a) outer semicontinuous (OSC) at z if for each sequence $\{z_i, s_i\} \rightarrow (z, s) \in \mathbb{R}^p \times \mathbb{R}^n$ satisfying $s_i \in M(z_i)$ for all $i \in \mathbb{Z}_{\geq 0}$, we have $s \in M(z)$; b) locally bounded at z if there exists an open neighborhood $N_z \subset \mathbb{R}^p$ of z such that $M(N_z)$ is bounded. We use $r\mathbb{B}$ to denote a closed ball in the Euclidean space, of radius $r > 0$, and centered at the origin, and we use $\{p\} + r\mathbb{B}$ to denote the union of all the points p_i that satisfy $|p - p_i| \leq r$. Given a set B , we use $\bar{B}$ and $\text{bd}(B)$ to denote the closure, and the boundary, respectively, and we use $\text{int}(B)$ to denote its interior. Given a single-valued or set-valued map f , we use $\text{dom}(f)$ to denote its domain.

In this paper, we will use the formalism of hybrid dynamical systems [14] for the synthesis and analysis of robust vision-based control systems. Specifically, a hybrid dynamical system (HDS) with state $z \in \mathbb{R}^n$ is represented by its *data* $\mathcal{H} := \{C, F, D, G\}$, and the dynamics

$$z \in C, \quad \dot{z} \in F(z), \quad (2a)$$

$$z \in D, \quad z^+ \in G(z), \quad (2b)$$

where the set-valued mappings $F : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ and $G : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$, called the flow map and the jump map, respectively, describe the evolution of the state z when it belongs to the flow set C , and the jump set D , respectively. Solutions to (2) are defined on *hybrid time domains*, which, under mild assumptions on the data $\mathcal{H}$, permits the use of graphical convergence notions to establish sequential compactness results for the solutions of (2), e.g., the graphical limit of a sequence of solutions is also a solution. These sequential compactness results play an important role in the robustness analysis of dynamical systems. For a precise definition of hybrid time-domains and solutions to HDS of the form (2) we refer the reader to [14, Ch. 2].

To guarantee suitable robustness properties, we will impose the following Basic Conditions on the data $\mathcal{H}$.

Definition 1: The HDS (2) is said to satisfy the Basic Conditions if: (a) the sets $C \subset \text{dom}(F)$ and $D \subset \text{dom}(G)$ are closed; (b) F is convex-valued, outer-semicontinuous, and locally bounded relative to C ; (c) G is outer-semicontinuous and locally bounded relative to D . $\square$

Note that when F is a (single-valued) continuous function, item (b) of Definition 1 is automatically satisfied.

III. THE OBSTACLE AVOIDANCE PROBLEM: ROBUSTNESS LIMITATIONS IN SMOOTH VISION-BASED CONTROL

In this paper, we are interested in the synthesis and analysis of robust feedback controllers able to autonomously steer a vehicle from any initial position $p_0 \in \mathbb{R}^2$ to a final target $p_T \in \mathbb{R}^2$, by using real-time data provided by a *visual sensor* as feedback. Typical examples include cameras and high-dimensional data generated by the fusion of multiple noisy sensors. To illustrate our controllers, we will consider simple velocity actuated vehicle dynamics, given by an integrator evolving on the plane, of the form

$$\dot{x} = u_x, \quad \dot{y} = u_y, \quad \theta = h(x, y), \quad (3)$$

where (x, y) are the coordinates in the Cartesian plane, and θ corresponds to real-time data generated by h , which can be seen as a map that produces images as functions of the vehicle's position. The main goal is to design a feedback law (u_x, u_y) such that the trajectories of the vehicle avoid an obstacle $\mathcal{N} \subset \mathbb{R}^2$ contained in a sphere of constant radius, and also converge to an arbitrarily small neighborhood of the target destination $p_T \in \mathbb{R}^2$. Such types of navigation problems have been extensively studied in the literature via different approaches, including planning and tracking algorithms [24], [25], triangular partitions [26], and barrier functions [27], to name just a few. In contrast to these settings, in this paper, we are interested in real-time feedback-based controllers where planning and navigation are simultaneously executed, and where robustness guarantees can be provided under arbitrarily small disturbances.

Remark 1: Even though, for simplicity, in this paper we focus on simple velocity actuated dynamics of the form (3), our results can be easily extended to more complex models, including nonlinear and nonholonomic dynamics, by using a multi-time scale approach, where a low-level controller (smooth, or hybrid, if needed) stabilizes the vehicle with respect to an external reference, see [15, Sec. VI]. $\square$

Gradient Flows, Anti-Potentials, and Perception Maps

One of the most popular approaches for the solution of navigation problems in mobile robots is based on implementing navigation functions $\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$, and gradient-based feedback laws of the form

$$u_x = k_x \frac{\partial \phi(x, y)}{\partial x}, \quad u_y = k_y \frac{\partial \phi(x, y)}{\partial y}. \quad (4)$$

In this setting, a continuously differentiable function ϕ is usually designed to have a maximizer at the desired target point, while also having minimizers at the location of the obstacles. In this way, the control law (4) incorporates attractive terms (to converge to the target point) and repulsive terms (to avoid the obstacles); see [16]–[20]. Note that the closed-loop dynamics (3)–(4) can be written as $\dot{p} = k \nabla \phi(\theta)$, with $p = (x, y)^\top$, evolving in the set $\mathbb{R}^2 \setminus \mathcal{N}$, where for simplicity we used $k_x = k_y = k \in \mathbb{R}_{>0}$.




Fig. 1: Closeness between trajectory of the robot and the predicted states by a learned perception map via convolutional neural networks (CNN).

To study controllers based on vision-based sensors, and similar to [8], we will assume the existence of a perception map ℓ that generates imperfect predictions of the state of the vehicle using the images θ , namely, $\ell(\theta) = Mp + e$, where $M \in \mathbb{R}^{2 \times 2}$ is a constant matrix, and $e \in \mathbb{R}^2$ is the approximation error. Using this perception map to close the loop between the camera and the vehicle, the resulting dynamics become

$$\dot{p} = k \nabla \phi(\ell(\theta)), \quad p \in \mathbb{R}^2 \setminus \mathcal{N}. \quad (5)$$

To learn the perception map ℓ , in this paper we will use data-driven techniques that make use of a sequence of labeled training data $\mathcal{T} = \{p_i, \theta_i\}_{i=1}^N$, to be used in traditional supervised learning methods (e.g., CNNs), and which is selected to satisfy the following assumption.

Assumption 1: For each compact set $K \subset \mathbb{R}^2$, and each pair $L, \varepsilon > 0$, there exists a function ℓ learned with training data $\mathcal{T} = \{p_i, \theta_i\}_{i=1}^N$, such that $K \subset \text{int}(\mathcal{S}_\varepsilon^L)$, where $\mathcal{S}_\varepsilon^L :=$

$$\bigcup_{(p_d, \theta_d) \in \mathcal{T}} \{p \in \{p_d\} + r\mathbb{B} : |\ell(\theta_d) - Mp_d| + L|p - p_d| \leq \varepsilon\}.$$

In words, Assumption 1 guarantees the existence of sufficient data to learn a suitable perception map that can cover any compact set K of interest. This assumption is standard in the literature of perception-based control, e.g., [8]. It allows establishing the following lemma, which will be instrumental for the characterization of the approximation error of the perception map learned from the data. The proof follows by a straightforward application of the triangle inequality.

Lemma 1: Let $F(p) := \ell \circ h(p) - Mp$, and suppose that Assumption 1 holds and $p \mapsto F(p)$ is L -Lipschitz. Then, $|\ell(\theta) - Mp| \leq \varepsilon$, for all (p, θ) such that $p \in \mathcal{S}_\varepsilon^L$. $\square$

To illustrate Lemma 1, Figure 1 shows a trajectory of a vehicle, as well as the predicted states by a perception map ℓ that satisfies $|\ell(\theta) - Mp| \leq \varepsilon$ on compact sets, with M being the identity matrix. The perception map was learned by using a convolutional neural network (CNN). As observed, the predictions of the perception map remain in an ε -neighborhood of the actual trajectory. Therefore, throughout the rest of this document, we will take $M = I$.

Note that if the learned perception map ℓ satisfies the conditions of Lemma 1, then the closed-loop system (5)




Fig. 2: An obstacle avoidance problem with target $\mathcal{G}$ and obstacle $\mathcal{N}$, and the “sensitive” set $\mathcal{K}$.

behaves as the following perturbed dynamical system

$$\dot{p} = k \nabla \phi(p + e), \quad |e| \leq \varepsilon, \quad \forall p \in \mathcal{S}_\varepsilon^L \setminus \mathcal{N}, \quad (6)$$

which has the form of (1). Stability and convergence properties of perturbed systems of the form (6) have been extensively studied in the control’s literature [11]. Indeed, as discussed in [11], [13], and [15], for the obstacle avoidance problem the disturbance e can have a dramatic effect on the trajectories of the vehicle. To illustrate this fact, consider Figure 2, where a vehicle, denoted with a white square, aims to converge to the target, denoted with a red circle while avoiding the obstacle $\mathcal{N}$ denoted with a white circle. Note that, to arrive at the target, the vehicle must choose a trajectory that goes above the obstacle or below the obstacle. Let $\mathcal{K}_1$ denote the set of initial conditions for which the closed-loop system (6) converges to the region $\mathcal{P}$ from above, and let $\mathcal{K}_2$ denote the initial conditions for which the closed-loop system (6) converges to the region $\mathcal{P}$ from below. It then follows that there must exist a set $\mathcal{K}$ where the vehicle must make a binary decision. Mathematically, for the obstacle avoidance problem, this behavior is captured by the following assumption; see also [13], [15]:

Assumption 2: There exists $T > 0$ such that for each $\rho > 0$ and each $\tilde{p}_0 \in \mathcal{K}$, where $\mathcal{K} := \bar{\mathcal{K}}_1 \cap \bar{\mathcal{K}}_2$, there exist points $\tilde{p}_1(0), \tilde{p}_2(0) \in \{\tilde{p}_0\} + \rho\mathbb{B}$, for which there exist solutions $\tilde{p}_1$ and $\tilde{p}_2$ of (6) with $e = 0$, satisfying $\tilde{p}_1(t) \in \mathcal{K}_1 \setminus \mathcal{K}$ and $\tilde{p}_2(t) \in \mathcal{K}_2 \setminus \mathcal{K}$ for all $t \in [0, T]$. $\square$

Under Assumption 2, the next proposition establishes zero margins of robustness against small adversarial perturbations $t \mapsto e(t)$ in the closed-loop system (6). The proposition follows by [28, Thm. 6.5] or [15, Prop. 1]:

Proposition 1: Suppose that Assumption 2 holds. Then for each $\varepsilon, \rho', \rho'' > 0$, and every $\tilde{p}_0 \in \mathcal{K} + \varepsilon\mathbb{B}$ such that $\tilde{p}_0 + \rho'\mathbb{B} \subset \mathbb{R}^2 \setminus \mathcal{N}$ and $\tilde{p}_0 + \rho''\mathbb{B} \subset (\mathcal{K}_1 \cup \mathcal{K}_2)$ there exist a piecewise constant function $e : \text{dom}(e) \rightarrow \varepsilon\mathbb{B}$ and a (Carathéodory) solution $\tilde{p} : \text{dom}(\tilde{p}) \rightarrow \mathbb{R}^2 \setminus \mathcal{N}$ to (6) such that $\tilde{p}(t) \in (\mathcal{K} + \varepsilon\mathbb{B}) \cap (\mathcal{K}_1 \cup \mathcal{K}_2) \cap (\tilde{p}_0 + \rho'\mathbb{B})$, for all $t \in [0, T']$ for some $T' \in (T^*, \infty]$, where $\text{dom } \tilde{p} = \text{dom } \tilde{e}$, $T^* = \min\{\rho', \rho''\}m^{-1}$, and $m = \sup\{1 + |k \nabla J(\eta)| : \eta \in p_0 + \max\{\rho', \rho''\}\mathbb{B}\}$. If T' is finite, then $\lim_{t \rightarrow T'} \tilde{p}(t) \notin (\mathcal{K}_1 \cup \mathcal{K}_2) \cup (\tilde{p}(0) + \rho'\mathbb{B})$. $\square$

The result of Proposition 1 has important implications for vision-based controllers based on perception maps, operating under topological obstructions such as obstacles. Namely, it establishes the existence of a set of points $\mathcal{K} \subset \mathbb{R}^2$ where *arbitrarily small* approximations e on the learned perception map ℓ can have a dramatic effect on the stability properties of the controller. Given that, in general, the error in the perception map ℓ can only be guaranteed to be bounded (see Lemma 1), Proposition 1 establishes that no robust controller based on smooth vector fields (e.g., based on navigation functions) exists for the solution of obstacle avoidance problems with inexact perception maps. Indeed, for navigation functions that combine attractive fields and repulsive fields, the set $\mathcal{K}$ will contain the spurious critical points of the navigation function ϕ , which includes any saddle-point¹. In this case, it is even possible to design adversarial disturbances $t \mapsto e(t)$ in (6) able to stabilize a spurious equilibrium [15, Ex. 1].

IV. ROBUST VISION-BASED HYBRID CONTROL

To synthesize a hybrid controller that overcomes the limitations of smooth feedback laws, we first characterize a class of admissible obstacles. We recall that $p_T \in \mathbb{R}^2$ denotes the target point of the robot.

Assumption 3: There exists $\rho \in \mathbb{R}_{>0}$ and $\varepsilon \in \mathbb{R}_{>0}$ such that the obstacle $\mathcal{N} \subset \mathbb{R}^2$ satisfies $\mathcal{N} \subset p_0 + \rho\mathbb{B}$ and $(p_0 + 2\rho\sqrt{2}\mathbb{B}) \cap (\{p_T\} + \varepsilon\mathbb{B}) = \emptyset$, where $p_0 = [x_0, y_0]^T \in \mathbb{R}^2$. $\square$

In words, Assumption 3 considers obstacles that are contained in spheres located sufficiently far away from the target point. Next, to achieve robust obstacle avoidance, we will design a switched *perception-based* controller that implements different potential fields in different sub-regions of the operational space of the vehicle. By using this switching approach, we will be able to rule out the emergence of problematic sets $\mathcal{K}$ such as the one shown in Figure 2. We note that our approach differs from existing works on hybrid control [13], [15], [21] due to the use of perception maps employed by the vehicles to estimate their positions in real-time. However, we also stress that our methodology can be naturally extended to other hybrid controllers that are well-posed in the sense of [14, Ch.7].

A. Synthesis of the Controller

To design the covering of the operational space, for each $p_0 \in \mathbb{R}^2$ and $\rho > 0$, define the set $\mathcal{B}_{p_0, \rho} := \{p \in \mathbb{R}^2 : \|p - p_0\| \leq 2\rho\sqrt{2}\}$, which satisfies $\{p_0\} + \rho\mathbb{B} \subset \mathcal{B}_{p_0, \rho} \subset \{p_0\} + 2\rho\sqrt{2}\mathbb{B}$. As in the standard state-based hybrid control [13], [15], we define the sets:

$$\begin{aligned} L_{1a} &:= \{p \in \mathbb{R}^2 : y < -x + y_0 + x_0 - 2\rho\sqrt{2}\}, \\ L_{1b} &:= \{p \in \mathbb{R}^2 : y < x + y_0 + x_0 + 2\rho\sqrt{2}\}, \\ L_{2a} &:= \{p \in \mathbb{R}^2 : y > x + y_0 + x_0 - 2\rho\sqrt{2}\}, \\ L_{2b} &:= \{p \in \mathbb{R}^2 : y > -x + y_0 + x_0 + 2\rho\sqrt{2}\}, \end{aligned}$$

¹The existence of such saddle points in navigation functions with attractive and repulsive fields was established in [16].




Fig. 3: Covering of the operational space of the vehicle. Top: $q = 1$. Bottom: $q = 2$.

as well as the unions $\mathcal{O}_1 := L_{1a} \cup L_{1b}$, $\mathcal{O}_2 := L_{2a} \cup L_{2b}$, and $\mathcal{O} := \mathcal{O}_1 \cup \mathcal{O}_2$. In this way, $\mathcal{O} = \mathbb{R}^2 \setminus \mathcal{B}_{p_0, \rho}$ and $\mathcal{N} \cap \mathcal{O} = \{\emptyset\}$. For each of the sets $\mathcal{O}_1$ and $\mathcal{O}_2$, we will design suitable potential functions V_q , $q \in \{1, 2\}$, that can be used in a gradient-based controller of the form (5). The controller will then switch between these two potential functions depending on its current location p generated by a perception map ℓ . Specifically, the potential functions are defined as

$$V_q(p) := \begin{cases} \phi_q(p) - \phi(p) & \forall p \in \mathcal{O}_q \\ \infty & \forall p \notin \mathcal{O}_q, \end{cases} \quad (7)$$

where ϕ and ϕ_q satisfy the next assumption.

Assumption 4: The functions $\{V_q\}_{q \in \{1, 2\}}$ satisfy the following: (a) For each $q \in \{1, 2\}$ there exist functions $\alpha_{1,q}, \alpha_{2,q} \in \mathcal{K}_\infty$, and proper indicators² $\tilde{\omega}_q$ of $\{p_T\}$ on $\mathcal{O}_q$, such that $\alpha_{1,q}(\tilde{\omega}_q(p)) \leq V_q(p) \leq \alpha_{2,q}(\tilde{\omega}_q(p))$, $\forall p \in \mathcal{O}_q$; (b) For each $q \in \{1, 2\}$, we have $\{p^* \in \mathcal{O}_q : \nabla V_q(p^*) = 0\} = p_T$; (c) For each $q \in \{1, 2\}$, the function $V_q(\cdot)$ is continuously differentiable in $\mathcal{O}_q$. $\square$

Remark 2: As shown in [15], the conditions of Assumption 4 can be readily satisfied using different classes of functions ϕ and ϕ_q . For example, they hold when ϕ is given by $\phi = -(x - x_T)^2 - (y - y_T)^2$, and ϕ_q is given by $\phi_q(p) := B(\tilde{d}_q(p))$, where $\tilde{d}_q(p) := |p|_{\mathbb{R}^2 \setminus \mathcal{O}_q}^2$, and $B(s) := (s - \rho)^2 \log(\frac{1}{s})$, if $s \in [0, \rho]$, and $B(s) := 0$, if $s > \rho$, with $\rho \in (0, 1]$ being a tunable parameter selected sufficiently small. Figure 3 shows the geometric structure of both sets $\mathcal{O}_1$ (top plot), and $\mathcal{O}_2$ (bottom plot). The level sets of the

²For a compact set $\mathcal{A}$ contained in an open set $\mathcal{U}$, a continuous function $\tilde{\omega} : \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$ is a proper indicator of $\mathcal{A}$ on $\mathcal{U}$ if $\tilde{\omega}(z) = 0$ if and only if $z \in \mathcal{A}$, and $\tilde{\omega}(z_i) \rightarrow \infty$ when $i \rightarrow \infty$ if either $|z_i| \rightarrow \infty$, or the sequence $\{z_i\}_{i=1}^\infty$ approaches the boundary of $\mathcal{U}$.




Fig. 4: Scheme of the closed-loop system.

functions V_q are also shown in Figure 3. Note that in each of the sets $\mathcal{O}_q$ the potential function V_q has a unique critical point located at the position of the target.

B. Main Results: Stability and Robustness

Using the above construction, we can now formulate the complete perception-based hybrid control system. Let $\chi \in (1, \infty)$ and $\lambda \in (0, \chi - 1)$ be tunable parameters. The closed-loop hybrid system has states $(p, q) \in \mathbb{R}^2 \times Q$, where $Q = \{1, 2\}$. The continuous-time dynamics are given by

$$\dot{p} = -k \nabla V_q(\ell(\theta)), \quad \dot{q} = 0, \quad (8)$$

which are allowed to evolve in the set

$$C_{p,q} := \{(\ell(\theta), q) \in \overline{\mathcal{O}} \times Q : V_q(\ell(\theta)) \leq \chi V_{3-q}(\ell(\theta))\}. \quad (9)$$

The discrete-time dynamics are given by

$$p^+ = p, \quad q^+ = 3 - q, \quad (10)$$

which are allowed to evolve in the set

$$D_{p,q} := \{(\ell(\theta), q) \in \overline{\mathcal{O}} \times Q : V_q(\ell(\theta)) \geq (\chi - \lambda) V_{3-q}(\ell(\theta))\}. \quad (11)$$

Note that in (8), (9), and (11), the position of the vehicle is given by the perception map $\ell(\theta)$ rather than the state p . The term $(\chi - \lambda)$ in (11) guarantees that the intersection of the sets $C_{p,q}$ and $D_{p,q}$ is not empty. Thus, for initial conditions in $C_{p,q} \cap D_{p,q}$ solutions are not unique. The set $C_{p,q}$ characterizes the points where the vehicle implements the controller (8) with constant state q . On the other hand, the set $D_{p,q}$ describes the points in the space where the vehicle toggles the logic state q whenever it approaches the boundary of the respective set $\mathcal{O}_q$. In particular, note that since $\chi > 1$ and $\chi - \lambda > 1$, the robot toggles the potential field V_q whenever its current value exceeds a threshold compared to the other potential fields $V_{q'}$. After each jump, the robot flows again using now the new potential function $V_{q'}$, until a new jump (if at all) is triggered. Note that this switching rule describes a hysteresis property in the feedback controller based on a supervisor mechanism. Figure 4 presents a schematic representation of the controller.

The following theorem is the main result of this paper.

Theorem 1: Let $\delta > 0$ and $K_0 \subset C_{p,q} \cup D_{p,q}$, where K_0 is compact. Suppose that $p \mapsto F(p)$ is L -Lipschitz continuous, where F is defined in Lemma 1. Then, there exists a perception map ℓ and training data $\mathcal{T} = \{p_i, \theta_i\}_{i=1}^N$ such

that every trajectory of the vehicle generated by the hybrid system (8)-(11), with initial condition in K_0 , is complete and converges to a δ -neighborhood of the target point p_T while avoiding the obstacle $\mathcal{N}$. $\square$

Proof: Under Assumption 1, and using Lemma 1, for any compact set $K \subset \mathbb{R}^2$, and any pair $\varepsilon, L \in \mathbb{R}_{>0}$ there exists a perception map $\ell(\theta)$ satisfying the bound:

$$|\ell(\theta) - p| \leq \varepsilon, \quad (12)$$

for all $p \in \mathcal{S}_\varepsilon^L$. It follows that $\ell(\theta) \in p + \varepsilon \mathbb{B}$, $\forall p \in \mathcal{S}_\varepsilon^L$. Based on this observation, on compact sets, the solutions of (8)-(11) are also solutions of the inflated inclusion with state $z = (p, \theta)$:

$$\dot{z} = \begin{pmatrix} \dot{p} \\ \dot{\theta} \end{pmatrix} \in F(z) := \begin{pmatrix} -k \nabla V_q(p + \varepsilon \mathbb{B}) \\ 0 \end{pmatrix} \quad (13a)$$

$$C_{p,q} = \left\{ z \in \overline{\mathcal{O}} \times Q : V_q(p + \varepsilon \mathbb{B}) \leq \chi V_{3-q}(p + \varepsilon \mathbb{B}) \right\} + \varepsilon \mathbb{B}, \quad (13b)$$

$$z^+ = \begin{pmatrix} p^+ \\ \theta^+ \end{pmatrix} \in G(z) := \begin{pmatrix} p \\ 3 - q \end{pmatrix}, \quad (13c)$$

$$D_{p,q} := \left\{ z \in \overline{\mathcal{O}} \times Q : V_q(p + \varepsilon \mathbb{B}) \geq (\chi - \lambda) V_{3-q}(\ell(\theta)) \right\} + \varepsilon \mathbb{B}. \quad (13d)$$

In turn, every solution of (13) is also a solution of an inflated hybrid system, given by

$$z \in C_\varepsilon, \quad \dot{z} \in F_\varepsilon(z), \quad (14a)$$

$$z \in D_\varepsilon, \quad z^+ \in G_\varepsilon(z), \quad (14b)$$

where the data $(C_\varepsilon, F_\varepsilon, D_\varepsilon, G_\varepsilon)$ is defined as

$$C_\varepsilon := \{z \in \mathbb{R}^n : (z + \varepsilon \mathbb{B}) \cap C \neq \emptyset\},$$

$$F_\varepsilon(z) := \overline{\text{co}} F((z + \varepsilon \mathbb{B}) \cap C) + \varepsilon \mathbb{B}$$

$$D_\varepsilon := \{z \in \mathbb{R}^n : (z + \varepsilon \mathbb{B}) \cap D \neq \emptyset\},$$

$$G_\varepsilon(z) := \{v \in \mathbb{R}^n : v \in g + \varepsilon \mathbb{B}, g \in G((z + \varepsilon \mathbb{B}) \cap D)\}.$$

Based on this observation, in order to establish a stability property for the closed-loop system (8)-(11), it suffices to establish a stability result for the inflated system (14). The following Lemmas will be instrumental for our results.

Lemma 2: The closed-loop hybrid system (13) with $\varepsilon = 0$ satisfies the Basic Conditions. $\square$

Proof: Follows directly by [14, Thm. 6.8]. $\blacksquare$

Lemma 3: Consider the HDS (13) with $\varepsilon = 0$. Then, under Assumption 4, the set $\{p_T\} \times Q$ is asymptotically stable with basin of attraction given by $\mathcal{O} \times Q$.

Proof: The proof follows the same ideas of [28, Sec. 6] and [15, Thm. 1]. Using Assumption 1, let us define $\tilde{\omega}(z) := \min_{q \in Q} \text{s.t. } p \in \mathcal{O}_q \tilde{\omega}_q(p)$ for each $z \in \mathcal{O}$. We obtain that $\tilde{\omega}$ is a proper indicator of p_T on $\mathcal{O}$. Let

$$\alpha_1(s) := \min_{q \in Q} \alpha_{1,q}(s), \quad \alpha_2(s) := \max_{q \in Q} \alpha_{2,q}(s), \quad (15a)$$

Using Assumption 4-(a), the function V_q satisfies

$$\alpha_1(\tilde{\omega}(p)) \leq V_q(p) \leq \alpha_2(\tilde{\omega}(p)), \quad \forall p \in \mathcal{O}. \quad (16)$$

During flows of the hybrid system, the time-derivative of V_q is given by:

$$\dot{V}_q(p) = -k|\nabla V_q(p)|^2 < 0, \quad (17)$$

for all $(p, q) \in C_{p,q} \cap (\mathcal{O} \setminus \{p_T\}) \times \{q\}$. Inequality (17) implies that, for each $q \in \{1, 2\}$, the function $V_q(z)$ decreases outside the target point p_T . On the other hand, jumps in the closed-loop system are allowed only when V_q gets larger or equal than $(\chi - \lambda)V_{3-q}$. Since, by construction $(\chi - \lambda) > 1$, it follows that during jumps V_q satisfies:

$$V_{q+}(p^+) \leq \frac{1}{\chi - \lambda} V_q(p), \quad \forall (p, q) \in D_{p,q}.$$

Therefore, the Lyapunov function decreases during jumps. The hysteresis mechanism rules out Zeno behavior, and the decrease of the Lyapunov function during flows and jumps implies that, for any complete solution of the system, the position p converges to p_T , uniformly on compact sets in the basin of attraction [14, Prop.7.5]. Completeness of solutions follows because: a) the system has no finite escape times; b) solutions cannot stop due to flows leaving the flow set; c) solutions cannot stop due to jumps leaving the union of the flow and jump set. This establishes the stability result. Obstacle avoidance follows by ε -closeness of solutions between the perturbed dynamics (14) and the nominal dynamics corresponding to (13) with $\varepsilon = 0$. $\square$

With Lemma 3 at hand, Theorem 1 follows now by a direct application [14, Thm. 7.21]. $\blacksquare$

To the best knowledge of the authors, Theorem 1 is the first result in the literature that integrates perception-based maps and hybrid controllers with stability and convergence guarantees. In fact, the previous arguments can be trivially extended to guarantee robustness with respect to additional external disturbances, including small measurement noise, sporadic camera failures, and slowly moving targets (provided they remain in a compact set), a setting that emerges in leader-follower systems where the follower tracks the position of the leader. Figure 5 illustrates this scenario by showing the level sets of V_q at three different instants of time. Here, the white triangle denotes the position of the leader, which acts as a target for the follower, denoted with the green triangle.

V. NUMERICAL EXPERIMENTS

We test the perception-based hybrid controller by training a perception map using a convolutional neural network. The model's architecture consists of three sets of Conv-ReLU-MaxPool blocks, with a kernel size of 3×3 and 2×2 , respectively. A dense layer, preceded by a Dropout layer with probability 0.5, takes the flattened output of the last Conv-ReLU-MaxPool block of layers and outputs a vector that imperfectly describes the state information. The model was trained using Keras. The training took place for 5 epochs, with batch size 128 and input images with shape (60, 100, 3). The output is a 1×2 vector describing the predicted (x, y) -position of the agent. The optimizer was Adam with learning rate 0.001. The loss function was the Mean Squared Error (MSE) between the predicted and real

state of the agent. Mean Absolute Error (MAE) was also used as a validation metric. To test the controller, we consider a leader robot (denoted with a white square) aiming to converge to the static target $\mathcal{G}$ while avoiding the obstacle $\mathcal{N}$. We also consider a follower robot (denoted with a green square), which tracks the leader. Both employ the hybrid controller, and the follower robot uses the learned perception map to approximate the leader's state. Figure 7 shows the trajectories obtained in this scenario under two different initial conditions: (-12, 2) in the left plot, and (-37, -17) in the right plot. In the simulations, we added noise to state measurements and we also included sensor failures (e.g., the camera does not transmit data at each time t with a certain probability). As observed, the hybrid controller provides suitable robustness properties. On the other hand, in Figure 6 we tested the generalization capabilities of the perception map. Here, we considered images containing occlusions (simulating, e.g., clouds). The left plot shows the performance of the hybrid controller when using a perception map trained on images *without* occlusions. It can be seen that the controller successfully handles the increment in prediction error due to inputs to the vision model being generated from a different process. The right plot of Figure 6 considers a controller trained on data *with* occlusions.

VI. CONCLUSIONS

In this paper, we introduced a perception-based hybrid controller for the robust solution of obstacle avoidance problems that use vision-based sensors for the purpose of feedback. Unlike existing results in the literature, our controller incorporates a perception map learned by using supervised learning methods, which provides a suitable approximation of the position of the vehicle based on images generated by a camera. By leveraging the structural robustness properties of the hybrid controller, and the generalization capabilities of the perception map, we established obstacle avoidance and convergence to the target point. Future research will focus on the theoretical guarantees under multiple obstacles.

REFERENCES

- [1] C. Finn, X. Y. Tan, Y. Duan, T. Darrell, S. Levine, and P. Abbeel, "Deep spatial autoencoders for visuomotor learning," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 512–519.
- [2] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [3] N. Wahlström, T. B. Schön, and M. P. Deisenroth, "From pixels to torques: Policy learning with deep dynamical models," *arXiv preprint arXiv:1502.02251*, 2015.
- [4] M. Watter, J. T. Springenberg, J. Boedecker, and M. Riedmiller, "Embed to control: A locally linear latent dynamics model for control from raw images," *arXiv preprint arXiv:1506.07365*, 2015.
- [5] E. Banijamali, R. Shu, H. Bui, A. Ghodsi *et al.*, "Robust locally-linear controllable embedding," in *International Conference on Artificial Intelligence and Statistics*. PMLR, 2018, pp. 1751–1759.
- [6] A. Lambert, A. Shaban, A. Raj, Z. Liu, and B. Boots, "Deep forward and inverse perceptual models for tracking and prediction," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 675–682.




Fig. 5: Time-varying level sets of a moving target converging to the static target x_T . In this simulation, a follower agent (green) tracks the leader (white). The left, center and right columns show the level sets at the times 1, 90 and 180, respectively.




Fig. 6: Snapshots of the trajectories of the robots implementing the hybrid controller, but using different perception maps. In the left figure, the perception map was trained on images *without* occlusions (shown in blue color). In the right plot, the perception map was trained using images *with* occlusion. In both cases, the hybrid controller is successful.




Fig. 7: Performance of the controllers when noise ($\epsilon \sim \mathcal{N}(0, 0.5)$) is added to the states of the robot, and the sensor of the follower has a failure rate of 50% (no photo is taken). The flow and jump sets for this controllers are defined with $\chi = 1.1$ and $\lambda = 0.09$.

- [7] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, "Learning latent dynamics for planning from pixels," in *International Conference on Machine Learning*. PMLR, 2019, pp. 2555–2565.
- [8] S. Dean, N. Matni, B. Recht, and V. Ye, "Robust guarantees for perception-based control," in *Learning for Dynamics and Control*. PMLR, 2020, pp. 350–360.
- [9] V. Petrik, M. Tapaswi, I. Laptev, and J. Sivic, "Learning object manipulation skills via approximate state estimation from real videos," *arXiv preprint arXiv:2011.06813*, 2020.
- [10] H. K. Khalil, *Nonlinear Systems*. Upper Saddle River, NJ: Prentice Hall, 2002.
- [11] E. D. Sontag, "Stability and stabilization: discontinuities and the effect of disturbances," *Nonlinear Analysis, Differential Equations and Control*, pp. 551–598, 1989.
- [12] C. G. Mayhew, *Hybrid Control for Topologically Constrained Systems, Ph.D. Dissertation*. University of California, Santa Barbara, 2010.
- [13] R. Sanfelice, M. Messina, S. Emre Tuna, and A. Teel, "Robust hybrid controllers for continuous-time systems with applications to obstacle avoidance and regulation to disconnected set of points," in *2006 American Control Conference*, 2006.
- [14] R. Goebel, R. G. Sanfelice, and A. R. Teel, *Hybrid Dynamical Systems*. Princeton, NJ, USA: Princeton University Press, 2012.
- [15] J. I. Poveda, M. Benosman, A. R. Teel, and R. G. Sanfelice, "Robust coordinated hybrid source seeking with obstacle avoidance in multi-vehicle autonomous systems," *IEEE Transactions on Automatic Control*, vol. 67, no. 2, pp. 706–721, 2022.
- [16] D. E. Koditschek and E. Rimon, "Robot navigation functions on manifolds with boundary," *Advances in applied mathematics*, vol. 11, no. 4, pp. 412–442, 1990.
- [17] H. Dürr, M. S. Stanković, D. V. Dimarogonas, C. Ebenbauer, and K. H. Johansson, "Obstacle avoidance for an extremum seeking system using a navigation function," in *proc. of American Control Conference*, pp. 4062–4067, 2013.
- [18] S. Paternain, D. Koditschek, and A. Ribeiro, "Navigation functions for convex potentials in a space with convex obstacles," *IEEE Transactions on Automatic Control*, vol. 63, no. 4, pp. 2944–2959, 2018.
- [19] A. Zuyev and V. Grushkovskaya, "Obstacle avoidance problem for driftless nonlinear systems with oscillating controls," *20th IFAC World Congress*, vol. 50, no. 1, pp. 10476–10481, 2017.
- [20] S. Paternain and A. Ribeiro, "Stochastic artificial potentials for online safe navigation," *IEEE Trans. Autom. Contr.*, vol. 65, no. 5, pp. 1985–2000, 2020.
- [21] S. Berkane, A. Bisoffi, and D. V. Dimarogonas, "Obstacle avoidance via hybrid feedback," *IEEE Transactions on Automatic Control*, pp. 1–1, 2021.
- [22] P. Braun, C. M. Kellett, and L. Zaccarian, "Explicit construction of stabilizing robust avoidance controllers for linear systems with drift," *IEEE Trans. Autom. Control.*, vol. 66, no. 2, pp. 595–610, 2021.
- [23] T. Strizic, J. I. Poveda, and A. R. Teel, "Hybrid gradient descent for robust global optimization on the circle," *IEEE 56th Conference on Decision and Control*, pp. 2985–2990, 2017.
- [24] E. Biyik and M. Arcak, "Gradient climbing in formation via extremum seeking and passivity-based coordination rules," *Asian Journal of Control*, vol. 10, no. 2, pp. 201–211, 2008.
- [25] S. Azuma, M. S. Sakar, and G. J. Pappas, "Stochastic source seeking by mobile robots," *IEEE Transactions on Automatic Control*, vol. 57, no. 9, pp. 2308–2321, 2012.
- [26] C. Belta, V. Isler, and G. Pappas, "Discrete abstractions for robot motion planning and control in polygonal environments," *IEEE Trans. on Automatic Control*, vol. 21, no. 5, pp. 864–874, 2005.
- [27] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [28] R. Sanfelice, *Robust Hybrid Control Systems*. University of California at Santa Barbara, 2007.