

Active Sensing for Dynamic, Non-holonomic, Robust Visual Servoing

Avik De*, Karl S. Bayer[†] and Daniel E. Koditschek*

Abstract—We consider the problem of visually servoing a legged vehicle with unicycle-like nonholonomic constraints subject to second-order fore-aft dynamics in its horizontal-plane. We target applications to rugged environments characterized by complex terrain likely to perturb significantly the robot’s nominal dynamics. At the same time, it is crucial that the camera avoid “obstacle” poses where absolute localization would be compromised by even partial loss of landmark visibility. Hence, we seek a controller whose robustness against disturbances and obstacle avoidance capabilities can be assured by a strict global Lyapunov function. Since the nonholonomic constraints preclude smooth point stabilizability we introduce an extra degree of sensory freedom, affixing the camera to an actuated panning axis mounted on the robot’s back. Smooth stabilizability to the robot-orientation-indifferent goal cycle no longer precluded, we construct a controller and strict global Lyapunov function with the desired properties. We implement several versions of the scheme on a RHex robot maneuvering over slippery ground and document its successful empirical performance.

I. INTRODUCTION

In GPS-denied settings, visual servoing—closed-loop position control by reference to fixed visual landmarks [1]—offers an attractive approach to self-localization, particularly over complex terrain where broken [2], unstable [3], and flowing [4] substrates preclude odometry. Advances in computer vision, processing power, and algorithmic insight [5] lend ever more speed and reliability to the extraction and tracking of natural features from successive camera images. We by-pass that problem in favor of a structured vision solution [6], focusing on the servo problem only.

In that context, range and field-of-view limits pose a fundamental problem: getting too close or far from a near-field scene will quickly degrade its efficacy as a visual landmark. Thus, our first major focus, following [7] is visual-obstacle avoidance: we require that our closed-loop vector fields not only guarantee convergence to a goal set but avoid straying out of the landmark’s visibility set along the way.

A second major focus places our paper within the voluminous and ever-growing literature on active sensing. It is well-known [8], [9], [10], [11] that simultaneous, coupled control of an actuated sensory subsystem in coordination with the robot’s body dynamics can remediate the stabilizability of underactuated base platforms. Following suit by recourse to a panning camera, our unicycle-like base, the dynamical legged XRHex [12] re-implementation of RHex [13], raises



Fig. 1. Our experimental apparatus with XRHex [12] and visual fiducials [6] as a landmark (Section III) using the proposed controller (Section II) to move to a goal location while avoiding visual obstacles.

two related new considerations. First, we seek to run the robot over considerably more complex terrain than in [14] and at considerably higher speeds known to necessitate a second-order unicycle model [15] for reasonable fidelity to horizontal plane dynamics. Second, our preoccupation with unstable terrain raises the prospect of significant perturbations, well beyond the infidelities of any unicycle model. The paper’s central contribution is to show how the active sensor’s relaxation of the goal set from a point into a cycle attractor affords a controller with a strict global Lyapunov function and its accompanying guarantee of both obstacle avoidance and ISS [16] defense against perturbations.

A. Related Literature

Fully-actuated first-order visual servo control is largely a solved problem [1] and, in particular, smooth stabilizability opens the door to the full suite of navigation function (NF) methods [7], including the lift to second-order dynamics [17] that can preserve all the first-order guarantees—not merely convergence, but obstacle avoidance [18], and robustness against disturbances as well [19]. Recent work incorporating visual servo methods into global landmark-based localization, navigation and mapping [20], [21] suggests the powerful role that effectively stabilized visual servo loops can achieve in complex task settings.

When smooth state-feedback point stabilizability fails, e.g. for nonholonomically constrained kinematic carts [22], an important and longstanding remediation is the introduction of systematic time-variation into the feedback signal [23], restoring exponential asymptotic point stability and the disturbance robustness it confers (cf. [24]). While this approach has been successfully applied in the broad context of visual servoing including the explicit construction of a strict Lyapunov function [10], [21], it is not clear to us how to extend that method to the present setting requiring

*Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA. {avik,kod}@seas.upenn.edu.

[†]Mechanical Engineering, Columbia University, New York, NY, USA. ksb2153@columbia.edu.

This work was supported by AFOSR MURI FA9550-10-1-0567.

obstacle avoidance and disturbance rejection in the context of a second-order plant. The alternative of introducing an explicitly nonsmooth (hybrid) point stabilizing law [14] can avoid obstacles but does not extend straightforwardly to second-order dynamics and cannot confer the robustifying ISS properties [16] (i.e., no single strict Lyapunov function can exist) in view of Brockett’s condition [22].

Instead, we take recourse to a third approach that also enjoys a long tradition in the visual servoing literature: the relaxation of the required task from point to submanifold stabilization via the introduction of an actively panned camera [10], [25], [21]. For example, in [25] this is arranged by locating the camera on a single d.o.f. boom actuated so as to rotate around the robot’s base. Now, placing the camera at a specified point in $SE(2)$ (the group of planar rigid transformations) corresponds via the inverse image of the camera map to an entire cycle (the circle, S^1) in the robot’s four d.o.f. configuration space. In [25], [21] this additional freedom is exploited by composing three distinct task controllers to (i) localize the camera frame, (ii) avoid both image plane occlusions, and (iii) avoid obstacles encountered in the corporeal world. No formal analysis of the composition of these controllers is furnished, whereas with an eye toward the robustness our application requires, we aim for a controller that admits a strict global Lyapunov function. Because it is not presently clear how to encompass all three of the desiderata within the control Lyapunov framework, we confine attention to the first two (convergence of the camera frame and the positive invariance of the visible set) with the benefit of introducing (and achieving) that alternative third goal of robustness.

The paper makes one further new contribution to the active visual servoing problem. The prior literature generally separates the camera’s optical center from its panning axis, thereby conferring full rank upon the Jacobian of the camera map. But we observe that there are many practical situations which mitigate against the enlarged workspace envelope associated with a “panning arm” of finite length.¹ At zero length, the camera map Jacobian (the interaction matrix of [25]) becomes singular and stabilizability of even the relaxed task specification, the circle, S^1 , requires passage to a second-order plant model (in a manner analogous to the method of approximate input-output linearization [26]). More generally, this Jacobian will be very poorly conditioned when the servo target is much farther away than the length of the camera-wielding manipulator (for instance, in an aerial surveillance vehicle), and we presume that successful physical implementation will require a model with singular camera map Jacobian, as we assume here.

B. Organization and Contributions of the Paper

The primary contribution of this paper is the development of a new, smooth, state-feedback-based visual servo con-

¹ For instance, we are able to protect the camera against the dangers of shocks and energetic impacts with the environment as RHex scrambles through its treacherous environment by mounting it so as to allow only pure rotation around the optical axis.

troller using the second-order lift of a family of navigation functions along with a formal proof that the controller achieves the three objectives: convergence (stabilization to any landmark configuration in the visibility set); obstacle avoidance (repulsion from the boundary of the visibility set); and robustness against disturbances (the guarantee that bounded-input disturbances result in bounded servo errors along with explicit computation of controller gain values relative to disturbance magnitudes guaranteeing the preservation of obstacle avoidance). These results are achieved by the analysis of a global Lyapunov function that is strict (i.e. positive definite with negative definite derivative) relative to a relaxed goal set (the circle, S^1 , arising as the inverse image of the specified camera frame under the camera map), and of uniform height along the zero section over the boundary of the visibility set.

The second contribution of the paper is an empirical implementation of various versions (the provably correct version as well as a higher performing version for which the correctness proof is still in progress) of this controller on the XRHex robot [12], [13], providing experimental evidence of its efficacy with emphasis on its strong disturbance rejection characteristics.

II. STABILIZING A UNICYCLE TO S^1

The most basic model of the horizontal-plane dynamics of a RHex with configuration $(p, \theta) \in \mathbb{R}^2 \times S^1 \approx SE(2)$ is that of a kinematic unicycle,

$$\begin{bmatrix} \dot{p} \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} \cos(\theta) & 0 \\ \sin(\theta) & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} =: B(\theta) \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}, \quad (1)$$

where u_1 is the forward speed and u_2 is the turning rate. In instantiations of this model where the body inertia or the speed levels are low enough that viscous effects dominate inertial effects, it is reasonable (and standard practice [24]) to assume that (u_1, u_2) can be directly controlled.

However, at the target speeds we wish to reach, a first-order model is a poor fit to the dynamics of RHex [15]; we will need to respect the robot’s significant inertial lags in this operating regime and impose a number of regularity conditions on any feedback controller $u = k(q)$ we introduce. Adding an integrator and assuming our controlled input is $v_1 = \dot{u}_1$ imposes the inertial model (and automatically ensures that $u(t)$ is C^1). The framework we propose here can further guarantee that (a) v_1 is a smooth function of the state (so $u_1(t)$ is C^2), and (b) $\|v_1\| = \|\dot{u}_1\|$ is bounded.

A. Second-order Control Scheme

We append the panning d.o.f. as a configuration variable $\alpha \in S^1$, defined in world coordinates for convenience, and assume that it is a fully actuated double integrator

$$\ddot{\alpha} = v_3. \quad (2)$$

Now $q := (p, \theta, \alpha) \in \mathcal{Q} := \mathbb{R}^2 \times T^2$ is the full configuration of the system, while $z := (p, \alpha) \in \mathcal{Z} := \mathbb{R}^2 \times S^1$ are the “task coordinates” that we are interested in.

In some of the relevant prior unicycle servoing literature [10], the Jacobian relating the control inputs to (in effect) \dot{z} is nonsingular due to the lever arm of the panning boom, i.e. the camera's position and orientation in any desired world frame can now be smoothly approached from any infinitesimal direction by appropriately coordinating the body and eye degrees of freedom. However, with our reluctance to commit (and be hostage) to the conditioning problems for a distant target, the camera optical center is modeled to be coincident with the panning axis, and a bit more work must be done to exploit the new actuated d.o.f. to achieve the desired stabilization to a cycle.

At this point, we assume² that the task can be expressed with respect to a smooth non-degenerate NF, φ , (possessing a unique local minimum at a point and attaining the exact height of unity at the obstacle boundaries) [27] over the task space \mathcal{Z} , with the additional properties of convexity³ and bounded gradient. For convenience, we will refer to the task positions "below" the φ -unity height as the visible set

$$\mathcal{V} := \{z \in \mathcal{Z} : 0 \leq \varphi(z) \leq 1\}$$

and detail the application-specific interpretation in Section III. Let $f := -D\varphi$ define the smooth, stable and uniformly bounded reference dynamics, $\dot{r} = f(r)$ for $r \in \mathcal{Z}$.

1) *Plant model*: Let $b = B^\perp$. Collecting and differentiating once the first two components of (1) and appending (2) yields the system dynamics,

$$\begin{bmatrix} B^T \\ b^T \end{bmatrix} \ddot{z} = \begin{bmatrix} v_1 \\ v_3 \\ u_1 u_2 \end{bmatrix}, \quad (3)$$

$$\dot{\theta} = u_2, \quad (4)$$

where $u_1 = B_1^T \dot{z}$ is the forward velocity of the unicycle. Note that $I = BB^T + bb^T$.

2) *Feedback Controller*: The incorporation of an actuated panning d.o.f. effectively affords us a partially decoupled plant model whereby the base orientation θ excites the task variables of interest, z purely through a feedforward cascade⁴. The task state, z , is still non-holonomically constrained ($b^T \dot{z} \equiv 0$), but our control idea here is to align θ so that there is no component of the reference acting in the constrained motion direction, i.e. we try to (a) minimize $|b^T f|$, thereby mitigating harmful consequences of underactuation as long as (b) $e := \dot{z} - f$ is small.

a) *Orientation Controller*: Let $\Pi := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$, let \angle be the function that returns the angle of a non-zero vector on the plane, and define the saturating angle function,

$$\angle_s(p) := \frac{\|p\|^2}{\|p\|^2 + \varepsilon_z^2} \angle p, \quad (5)$$

²This assumption is ensured by our construction in Section III-A.

³Convexity—a convenient property available to our particular model space, \mathcal{V} , to be introduced in Section III—is not essential to the controller. Its chief virtue here is to afford an immediate extension of the ISS framework [16] to goal sets which are compact manifolds rather than points [28] without the need for more intricate reasoning about the existence and nature of Control-Lyapunov Functions relative to non-goal components of the invariant set that would otherwise be required [29].

⁴However, in (11), that we will find it necessary to feed the task variables back into the θ dynamics, thereby imposing a coupling in the closed-loop.

which returns 0 when $p = 0$, and is smooth. Let $\phi := \angle_s(\Pi f)$ and $\theta_r := \theta - \phi$. Define the orientation error,

$$\eta_0 = 1 - \cos \theta_r. \quad (6)$$

Let us define $s := \sin \theta_r$, and pick

$$u_2 = \dot{\phi} - k_0 s \implies \dot{\eta}_0 = -k_0 s^2 = -k_0 \eta_0 (2 - \eta_0). \quad (7)$$

Lastly, note that the parasitic component of the reference which we wish to minimize satisfies (proof in Appendix I)

$$|b^T f| \leq \|\Pi f\| (s + \delta), \quad (8)$$

where the magnitude of $\delta(z) = \varepsilon_\perp o(\|\Pi f\|)$ (an artifact of the saturation in \angle_s) can be controlled directly by ε_\perp .

In Section IV, we present and use a modified orientation controller with a multistability property, which demonstrates empirically superior performance.

b) *Graph Error Controller*: To control the velocity error e , we adopt established methods [17] of lifting first-order reference dynamics to second-order plants by using the graph error as a "kinetic energy," and we briefly summarize them here. In our notation, let

$$d(z, \dot{z}) := \dot{f} - D\varphi - k_1 e \quad (9)$$

be the desired graph error control, where $e := \dot{z} - f$. Note that $B^T d$ can be matched exactly with available control inputs, i.e. we can make $B^T (\ddot{z} - d) \equiv 0$.

3) *Closed-loop dynamics*: Summarizing from the previous section, our feedback controller sets

$$\begin{bmatrix} v_1 \\ v_3 \\ u_2 \end{bmatrix} = \text{FB}(z, \dot{z}, \theta) := \begin{bmatrix} B^T d \\ \dot{\phi} - k_0 s \end{bmatrix}. \quad (10)$$

Inserting this into the plant dynamics, we obtain the closed-loop dynamics

$$\dot{\theta} = g_0(z, \dot{z}, \theta) := D_z \phi \dot{z} - k_0 s(z, \theta) \quad (11)$$

$$\ddot{z} = g_1(z, \dot{z}, \theta) := BB^T d + \dot{\theta} b B_1^T \dot{z}, \quad (12)$$

where B_1 is defined to be the first column of B .

The state space takes the form of a cross product, $S^1 \times T\mathcal{Z}$, and we will typically use the coordinates (θ, z, \dot{z}) for computation purposes. The following propositions express the stability properties of the proposed control scheme.

Proposition 1 (Stability). *Denote by $\mathcal{G} := \varphi^{-1}(0) \times \{0\} \in T\mathcal{Z}$ the task goal set, comprising the zero section over the desired position in the task variables. Under the closed-loop dynamics (11), (12), the base-indifferent goal set, $S^1 \times \mathcal{G}$ is asymptotically stable and its basin of attraction includes all of the visibility set, $\mathcal{V} \subset \mathcal{Z}$.*

Proof. Please refer to Appendix II for the definitions and detailed derivations of the computations in this proof.

We will show that $\eta := \Upsilon_2 \eta_0 + \eta_1$, where $\eta_1 := \varphi + \frac{1}{2} e^T e$, is a strict global Lyapunov function for the goal set $S^1 \times \mathcal{G}$. Clearly, η is positive definite with respect to this goal set. It now suffices to show that $\dot{\eta}$ is negative definite along the motions of the closed-loop dynamics. We have,

$$\dot{\eta}_1 = \dot{\varphi} + e^T \dot{e} = D\varphi^T f - k \|e\|^2 - f^T b b^T (\ddot{z} - d), \quad (13)$$

where the last summand can be thought of as “noise” owed to our underactuation. However, $|f^T b|$ is controlled by our orientation controller, for which we have already presented the Lyapunov function η_0 (6).

As shown in Appendix II, the “noise” terms are subquadratic in the velocity error $\|e\|$ or $|s|$, with coefficients that are bounded functions of z (recall $z \in \mathcal{V}$, a compact domain), only perturbed by our saturation by $\varepsilon_{\mathcal{L}}$ of Section II-A.2.a. The derivative of η is

$$\dot{\eta} \leq -\tilde{k}_1(\|e\| - a_8)^2 - \frac{\|f\|^2}{k_0}(\tilde{k}_0 - \Upsilon_3) + \varepsilon_{\mathcal{L}}o(\|f\|^2), \quad (14)$$

where Υ_3 is bounded and so \tilde{k}_0 can simply be made large enough to make the last parenthesized term positive, and the magnitude of the last term can be made arbitrarily small. Moreover, $\|f\| = 0 \implies a_8 = 0$.

The sum of the first two terms above is negative semidefinite, and the sum is zero only if both of the following hold:

- 1) $\|f\| = 0 \implies z \in \varphi^{-1}(0)$ by our assumptions on φ , and that $a_8 = 0$,
- 2) $\|e\| = 0 \implies \dot{z} = f(z)$, which, along with 1), in turn implies $\dot{z} = 0$.

Since $\varepsilon_{\mathcal{L}}$ is arbitrarily small, η is nonpositive, vanishing only when $(z, \dot{z}) \in \mathcal{G}$, and the result follows from Lasalle’s Invariance Principle [30]. \square

Proposition 2 (Obstacle Avoidance). *If the initial (rest) configuration satisfies $\eta(0) \leq 1$, the closed-loop system avoids φ -modeled obstacles by never leaving the visible states, $T\mathcal{V}$.*

Proof. Proposition 1 ensures that $\dot{\eta} \leq 0$, and consequently $\eta(t) < 1$ for $t > 0$. Using the definition of η , $\varphi \leq \eta < 1$, and from the definition of φ , we get the desired conclusion. \square

In addition to the properties above, the existence of the strict (with respect to the compact goal set, $S^1 \times \mathcal{G}$ [28]) Lyapunov function η implies that the closed-loop system (11), (12) can handle bounded disturbance vectors injected into the dynamics (3), (4) and converge to “small” (instead of 0, in the disturbance-free situation) asymptotic error. Disturbance vectors could be unmodeled perturbations from terrain interactions (see Section IV), or a time-varying reference system $f(z, t)$ with bounded $D_t f$. We emphasize here that (for e.g.) a hybrid controller that cannot have a strict Lyapunov function could not provide this guarantee.

Lastly, we note that (although not considered here) a second-order steering version of the θ -dynamics (a parsimonious model for RHex at higher speeds [15]) could be handled by a very similar control methodology, by altering the orientation controller in Section II-A.2.a.

III. VISUAL SERVOING WITH NAVIGATION FUNCTIONS

Since our configurations are strictly in the horizontal plane, a constellation of three identifiable collinear feature points provides a diffeomorphic “sensor map” (section III-A in [7]). We note here that our methods and results extend directly to any scenario where there is a diffeomorphism from




Fig. 2. A simulation comparing the task-space behavior of the NF construction in motivating prior work [7] (left) and the proposed construction in III-A with the smooth orientation controller of Section II-A.2.a (middle), and the non-smooth orientation controller of Section III-B (right). The beacon constellation is drawn as the triple of black dots, the contour lines are the level sets of φ at the slice $z_3 = e_3^T \varphi^{-1}(0)$ (where e_3 is the third basis vector), and the red dot on the right is the goal location. The rightmost two simulations are deliberately underdamped, to demonstrate the “preferred direction” side-effect of the smooth orientation controller.

the sensorium to the robot task coordinates, not restricted to planar robots or indeed to visual sensing.

Define the sensor space as $\mathcal{Y} = [-1, 1]^3$, where the bounds are imposed by the limited field-of-view of the camera. Adapted from the definition of the “camera map” in [7] (which we shall refer to as c_C), we define a sensor map $c : \text{SE}(2) \rightarrow \mathcal{Y}$ such that $y = c(z)$ as $c = c_C \circ \mathcal{R}$, where $\mathcal{R} : \text{SE}(2) \rightarrow \text{SE}(2)$ is the rigid body transformation

$$\mathcal{R}(p, \alpha) = \begin{bmatrix} -p_1 \cos \alpha - p_2 \sin \alpha \\ p_1 \sin \alpha - p_2 \cos \alpha \\ -\alpha \end{bmatrix}. \quad (15)$$

In the “visible set” (configurations z that face the camera and are such that $z_2 < 0$), $\det(Dc) \neq 0$. Let us define (for notational brevity) $J(y) := Dc^{-1}|_{c^{-1}(y)}$. For our implementation, we assume that our sensor measures (y, \dot{y}) . Using the same notation as [7], the set $\mathcal{J} \subset \mathcal{Y}$ is the compact domain of interest in the image space, and $\mathcal{V} = c^{-1}(\mathcal{J})$ is the corresponding (also compact) region in the task space.

The implementation of visual servo control is a straightforward adaptation of our controller synthesis in Section II. Instead of specifying reference dynamics in the task coordinates, we specify them directly in image coordinates⁵ (we construct φ explicitly in Section III-A).

We wish to emphasize here that all the calculations we perform are done completely in the image-space. Instead of algebraically solving for $z = c^{-1}(y)$ and using this in z -reference dynamics (this method is sensitive to camera calibration errors, which affect c^{-1}), we are solving these image-task equations dynamically. Of course, the stability and obstacle avoidance results of Propositions 1 and 2 hold in whatever coordinates the closed-loop system is executed.

It should be noted that $J(y)$ is non-trivial to compute. The approximation used in [7] was accurate enough for initial conditions close to the goal, however we wish to perform large-scale motions which preclude local approximations. The world x -coordinates of our beacon (three collinear visual fiducials, each pair equidistant) are assumed to be

⁵It should be mentioned that the nonlinear change of coordinates c introduces a “Coriolis-like” term, because when pulled back to z -coordinates, $\ddot{z} = \tilde{J}\dot{y} + J\ddot{y}$. Empirically, we found that the \tilde{J} terms were small in magnitude in our energetic regime, and ignored them in our experiments.

$[-1, 0, 1]$ without loss of generality. If $v := (y_1 y_2 + y_2 y_3 - 2y_1 y_3, -y_1 + 2y_2 - y_3)$ and $v^\perp := \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} v$, then

$$c^{-1}(y) = - \left[\frac{1}{\|v\|} [v \ v^\perp] \cdot \begin{bmatrix} y_2(y_3 - y_1) \\ y_3 - y_1 \end{bmatrix} \right]. \quad (16)$$

We can use this analytical form for c^{-1} to compute $J(y)$.

A. Navigation Function Construction

Here we construct a NF [27] for servoing to a point while avoiding image-space obstacles [7]. Let us define

- the goal function $\gamma = \gamma_z \circ c^{-1} : \mathcal{Y} \rightarrow \mathbb{R}_+$, where $\gamma_z(z) = \frac{1}{2} \|z - z_0\|^2$, and z_0 is the goal point in the task coordinates⁶, and
- the image-space obstacle function $\beta : \mathcal{Y} \rightarrow \mathbb{R}_+$ identically to the construction in section III-A of [7].

Let $\sigma : \mathbb{R}_+ \rightarrow [0, 1]$ be the squashing diffeomorphism $\sigma(x) := x/(1+x)$, and $\varphi_0 := \sigma \circ \gamma$ be the ‘‘obstacle-free’’ NF. We are essentially seeking a way to compose the constituent functions, γ, β , in a manner that simultaneously preserves the properties of attraction to the goal and repulsion from obstacles. These functions are topologically coupled, precluding methods like cross-product composition [31]; to our knowledge the only formally well established construction in the literature is still from Koditschek and Rimon [27],

$$\varphi_{KR} = (\sigma(\gamma^k/\beta))^{1/k}. \quad (17)$$

This construction is easy to compute and provably correct over a number of free-space models [27] including the present one [32]. Furthermore, it enjoys the additional property of convexity over the (convex) space \mathcal{Y} [32] of interest in this application, hence it supports the hypotheses of Proposition 1, and we will rely upon it for the formally instantiated construction of φ in that Proposition.

However, the high values of k required to achieve convexity can make the associated Jacobian somewhat numerically ill-conditioned in regions of \mathcal{Y} . Hence we propose a new construction with improved numerical properties,

$$\varphi := \sigma \left(\frac{\gamma}{\sigma(\beta/\varepsilon)} \right), \quad (18)$$

which is intuitively designed (for small values of ε) so that $D\varphi \approx D\varphi_0$ when β is large, and $D\varphi$ is dominated by $D\beta$ and repels away from the obstacle when β is small. The tuning parameter ε controls the trade-off between steep (approaching logical) obstacle ‘‘walls’’ and torque requirements.

The new construction has the basic task-encoding properties $\beta^{-1}(0) \subset \varphi^{-1}(1)$, and $\varphi^{-1}(0) = \gamma^{-1}(0)$, and we conjecture that it is a NF that is also convex over the model space, \mathcal{Y} , but this latter property remains to be proven true.

⁶We use the composition with c^{-1} and use the modified reference dynamics $f(y) = -J(y)^{-1}J(y)^{-T}D\varphi$ in our implementation; these modifications from (11), (14) of [7] give rise to relatively more desirable trajectories (metrically) in \mathcal{Z} (cf. Figure 2). Completing the proof of correctness for the new constructions is presently in progress.

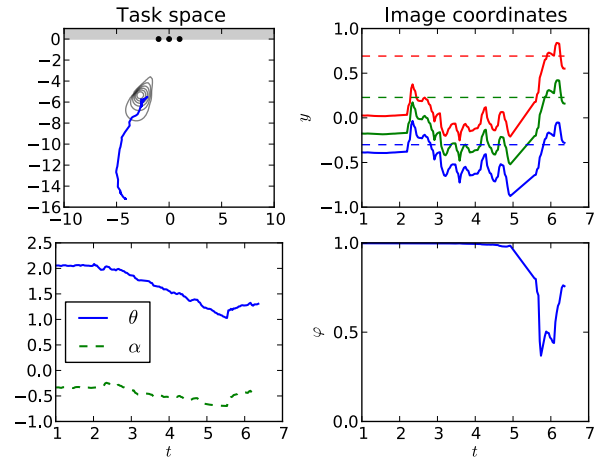


Fig. 3. A single trial of our algorithm running on RHex. The top left plot shows level sets of a slice of φ at the goal α . The bottom left plot shows the benefit of the active sensor: since θ is free from the task coordinates, the robot can orient itself independently of the task-space obstacle, eliminating the need for ‘‘parallel-parking’’ maneuvers.

B. Simulations

Absent a formal treatment, in Figure 2, we demonstrate with simulations some of the empirical benefits of the construction suggested in Section III-A. Additionally, we use our simulations to test relaxations of some of the conditions required for our analytical guarantees in Proposition 1:

- 1) We use a smaller k_0 than the conservative bound.
- 2) We introduce a new ‘‘tuning knob’’ by using separate proportional and damping gains to form a modified version of (9) and drop the feedforward \dot{f} to get

$$d(z, \dot{z}) := -k_p D\varphi - k_d \dot{z}.$$

- 3) We drop the feedforward $\dot{\phi}$ term in (7) and substitute a non-smooth, Lipschitz version⁷ which does not have a preferred ‘‘forward direction’’,

$$u_2^{\text{NS}} = \mp k_0 \sin \theta_r,$$

with the ‘‘-’’ sign chosen when $|\theta_r| < \pi/2$ and vice versa, whose Lyapunov function is $\eta_0^{\text{NS}} = 1 \mp \cos \theta_r$. Figure 2 shows that the preferred orientation of (7) causes the simulated robot to reverse its heading after the initial overshoot.

These relaxations invalidate the conditions in the proof of Proposition 1, but Figure 2 provides evidence that the conditions are not *necessary*, which is encouraging for robot implementation in the face of restrictions imposed by computational resource scarcity and desire for higher performance.

IV. ROBOT EXPERIMENTS

Our experiments were performed on XRHex [12] equipped with a panning camera (Point Grey DragonFly 2). The ‘‘beacon’’ comprised three distinguishable visual fiducials [6]. The provably correct scheme [7] was implemented, but

⁷This function is inspired by the multistable potentials in [33] and has stable attractors at $\theta_r = 0$ and π

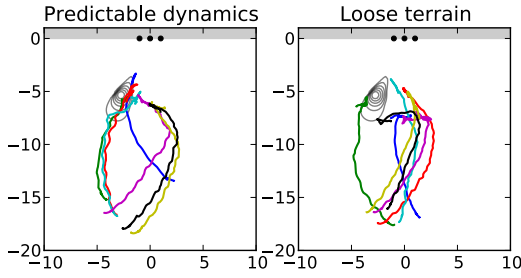


Fig. 4. A visual illustration of performance degradation when the robot dynamics get worse (due to shifting terrain). In some cases on loose terrain the perturbations are large enough to cause failure (Proposition 1 can be used to generate conservative bounds on the magnitude of perturbations allowed relative to the available force and power limits).

suffered from previously alluded-to numerical instabilities (Section III) and metric inconveniences (Section III-B). In the interests of space we will only show data for the newer, higher-performing construction (Section III-A) whose formal proof of correctness is in progress as of the submission. We used three testing environments: (a) indoors with good traction, (b) indoors with shifting terrain⁸, and (c) outdoors. We include results from (c) in the video attachment.

Figure 3 shows a representative trial where the task-space behavior of the robot is plotted. Figure 4 contains results from comparative tests between relatively predictable RHex dynamics, and loose, shifting terrain. Note from the left plot that the robot roughly travels along gradient flows of φ , indicating that lifting a first-order (gradient) reference field, f , to a viable second-order trackable form [17] gets the second-order system to essentially perform gradient descent.

V. CONCLUSIONS AND FUTURE WORK

We develop a theoretical framework targeting robust controller design for a dynamic underactuated system by adding sensor d.o.f.s which enable task-point stabilization accompanied by a strict global Lyapunov function in the task-space—even though the dynamics are still non-holonomically constrained. We apply this framework to the synthesis and implementation of visual servo control on a RHex, documenting the robustness by tests over substrates where the dynamics are unpredictable. Notwithstanding the specific applications focus on a cart (albeit in a regime where a kinematic model is a poor fit), it generalizes easily to truly second-order systems such as UAVs, AUVs, or sagittal-plane hopping robots (the last of which we wish to examine carefully in future work).

We did not consider the problem of world-space obstacle avoidance, though it is straightforward to incorporate modeled obstacles with the mature NF framework [27] that we have adopted (Section III-A). we are intrigued to suspect that an extension of the control methods we presented here might be applicable to that setting.

Lastly, note that while we used a visual system, the framework in Section III is easily extensible to any sensing

⁸This was constructed by placing multiple layers of slippery plastic sheets on the ground such that the inter-layer static friction was far lower than that required to propel the robot, causing the pieces to shift constantly.

TABLE I
LIST OF APPENDIX II SYMBOLS

Name	Definition
$a_1(z)$	$\ b^T Df\ $
$a_2(z, k_1)$	$\ -b^T D\varphi - k_1 \ f\ s\ $
$a_3(z, k_1)$	$\ B_1 D_z \phi\ $
$a_4(z, k_0)$	$\ 2D_z \phi f B_1^T - k_0 s B_1^T\ $
$a_5(z)$	$\ B_1^T f D_z \phi f\ $
$a_6(z, k_0)$	$a_1 + a_4$
$a_7(z, k_1)$	$a_2 + a_5$
$a_8(z, k_1)$	$\frac{a_6 \ \Pi f\ (s + \delta \varepsilon_{\angle})}{2k_1}$
\tilde{k}_1	$k_1 - \ \Pi f\ (s + \delta \varepsilon_{\angle}) a_3$ (must be > 0)
\tilde{k}_0	$(\frac{\Upsilon_2}{\ \Pi f\ } - 1) k_0 - \frac{a_6^2}{4k_1}$ (must be > 0)
Υ_0	$\max_{z \in \mathcal{V}} \frac{\ \Pi f\ ^2}{4k_1}$
Υ_1	$\max_{z \in \mathcal{V}} \ \Pi f\ a_3(z)$
Υ_2	$\max_{z \in \mathcal{V}} \ \Pi f\ ^2$
$\Upsilon_3(k_1)$	$\max_{z \in \mathcal{V}} \frac{a_7^2}{4}$

model which admits a local diffeomorphism between the robot pose and measurements in the sensorium—an avenue we wish to explore in the future.

APPENDIX I DERIVATION OF (8)

We will use a, x, y as dummy variable names in this section. First, note that $b^T f = b^T \Pi f$. Using the definition (5), when $f \neq 0$,

$$\begin{aligned} |b^T f| &= \|\Pi f\| \sin(\theta - \angle(\Pi f)) = \|\Pi f\| \sin(\theta_r + \phi \frac{\phi \varepsilon_{\angle}^2}{\|\Pi f\|^2}) \\ &= \|\Pi f\| (s \cos \frac{\phi \varepsilon_{\angle}^2}{\|\Pi f\|^2} + \cos \theta_r \sin \frac{\phi \varepsilon_{\angle}^2}{\|\Pi f\|^2}) \\ &\leq \|\Pi f\| |s| + \|\Pi f\| \sin \frac{|\phi| \varepsilon_{\angle}^2}{\|\Pi f\|^2}. \end{aligned}$$

The “error term” (last summand) is of the form $m(x) := x \sin(\frac{a}{x^2})$ for $x, a > 0$. It is easy to see that $m(x) \leq x$.

Also, $m(x) \leq \frac{a}{x}$. To see this, let $y = \frac{a}{x^2}$, and note that

$$\lim_{y \rightarrow 0} \sqrt{\frac{a}{y}} \sin y = \sqrt{a} \lim_{y \rightarrow 0} \frac{\sin y}{y} \sqrt{y} = \sqrt{ay},$$

and other than at the limit, $\sin(y) \leq y$ for $y \geq 0$, and so this is an upper bound: $m(x) \leq \frac{a}{x}$.

Putting these together,

$$m(x) \leq \min\{x, \frac{a}{x}\} = x \min\{1, \frac{a}{x^2}\}.$$

The maximum value is at $x = \sqrt{a}$, where $m(\sqrt{a}) \leq \sqrt{a}$.

Returning to $b^T f$ with $a = |\phi| \varepsilon_{\angle}^2$, $x = \|\Pi f\|$,

$$|b^T f| \leq \|\Pi f\| (|s| + \frac{\sqrt{2\pi} \varepsilon_{\angle}}{\|\Pi f\|^2}) =: \|\Pi f\| (|s| + \varepsilon_{\angle} \delta(z)),$$

where we define δ for notational brevity, and note that $\|\Pi f\| \varepsilon_{\angle} \delta(z) \leq 1$ from the bound obtained above.

APPENDIX II PROOF OF PROPOSITION 1

a) *Derivation of (13):* From the definition of η_1 ,

$$\begin{aligned} \dot{\eta}_1 &= \dot{\varphi} + e^T \dot{e} = \dot{\varphi} + e^T (\ddot{z} - \ddot{f}) \\ &= \dot{\varphi} + e^T (\ddot{z} - d) + e^T (d - \ddot{f}) \\ &= \dot{\varphi} + e^T B B^T (\ddot{z} - d) + e^T b b^T (\ddot{z} - d) + e^T (-D\varphi - k_1 e) \end{aligned}$$

where we use the fact that $I = BB^T + bb^T$. For the next steps, recall that (10) ensures $B^T(\ddot{z} - d) = 0$, and that $b^T\dot{z} = 0$, resulting in

$$\dot{\eta}_1 = D\varphi^T f - k_1\|e\|^2 - f^T bb^T(\ddot{z} - d).$$

b) *Bounding the ‘‘Actuation Noise’’ Terms:* Note that

$$\begin{aligned} b^T d &= b^T(\dot{f} - D\varphi + k_1 f) \leq a_1\|e\| + a_2, \text{ and} \\ b^T \ddot{z} &= u_1 u_2 = u_1(D_z \phi(e + f) - k_0 s) \\ &= (B_1^T e + B_1^T f)(D_z \phi(e + f)k_0 s) \\ &\leq a_3\|e\|^2 + a_4\|e\| + a_5 + k_0|s|\|\Pi f\|. \end{aligned}$$

Adding together all the terms,

$$\begin{aligned} \dot{\eta}_1 &\leq D\varphi^T f - k_1\|e\|^2 + \|\Pi f\|(|s| + \delta\varepsilon_\perp)(a_3\|e\|^2 + a_6\|e\| \\ &\quad + a_7 + k_0|s|\|\Pi f\|) \\ &\leq D\varphi^T f - \tilde{k}_1(\|e\| - a_8)^2 + \|\Pi f\||s|(a_7 + k_0|s|\|\Pi f\|) \\ &\quad + \frac{a_6^2\|\Pi f\|^2|s|^2}{4k_1} + \varepsilon_\perp\|\Pi f\|\delta(z)a_9(z), \end{aligned}$$

where $a_9(z)$ contains bounded z -terms which are not explicitly defined here.

c) *Combining η_0 and η_1 :* Adding the two summand derivative terms,

$$\begin{aligned} \dot{\eta}_0 + \Upsilon_2 \dot{\eta}_1 &\leq D\varphi^T f - \tilde{k}_1(\|e\| - a_8)^2 + (k_0 + \frac{a_6^2}{4k_1})\|\Pi f\|^2|s|^2 \\ &\quad + \|\Pi f\||s|a_7 + \varepsilon_\perp\|\Pi f\|\delta(z)a_9(z) - \Upsilon_2 k_0 s^2 \\ &= -\|f\|^2 - \tilde{k}_1(\|e\| - a_8)^2 - \tilde{k}_0(|s| - \frac{\|f\|a_7}{2k_0})^2 \\ &\quad + \frac{\|\Pi f\|^2 a_7^2}{4k_0} + \varepsilon_\perp\|f\|\delta(z)a_9(z) \\ &\leq -\tilde{k}_1(\|e\| - a_8)^2 - \frac{\|f\|^2}{k_0}(\tilde{k}_0 - \Upsilon_3) + \varepsilon_\perp\|f\|o(z), \end{aligned}$$

where we catch all the bounded functions of z in $o(z)$. Now note that $\delta(0) = 0$, and a linearization yields that for small $\|z\|$, $\delta(z) \sim o(\|z\|)$. Non-degeneracy of φ implies that close to the goal, $\|z\| \sim o(\|f\|)$, and including this, the last term is also $o(\|f\|^2)$ and proportional to ε_\perp .

REFERENCES

- [1] F. Chaumette and S. Hutchinson, ‘‘Visual servo control. i. basic approaches,’’ *IEEE Robotics Automation Magazine*, vol. 13, no. 4, pp. 82–90, 2006.
- [2] R. Murphy, ‘‘Human-robot interaction in rescue robotics,’’ *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, vol. 34, no. 2, pp. 138–153, 2004.
- [3] D. Vergano, ‘‘Nasa’s spirit robot stuck in sand on mars,’’ *USA Today*, Jan. 2010.
- [4] C. Li, P. B. Umbanhowar, H. Komsuoglu, D. E. Koditschek, and D. I. Goldman, ‘‘Sensitive dependence of the motion of a legged robot on granular media,’’ *Proceedings of the National Academy of Sciences*, vol. 106, no. 9, pp. 3029–3034, 2009.
- [5] E. Marchand and F. Chaumette, ‘‘Feature tracking for visual servoing purposes,’’ *Robotics and Autonomous Systems*, vol. 52, no. 1, pp. 53–70, 2005.
- [6] E. Olson, ‘‘AprilTag: a robust and flexible visual fiducial system,’’ in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, 2011, pp. 3400–3407.
- [7] N. Cowan, J. Weingarten, and D. Koditschek, ‘‘Visual servoing via navigation functions,’’ *IEEE Transactions on Robotics and Automation*, vol. 18, no. 4, pp. 521–533, Aug. 2002.
- [8] R. Zhang, P.-S. Tsai, J. Cryer, and M. Shah, ‘‘Shape-from-shading: a survey,’’ *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 8, pp. 690–706, Aug. 1999.
- [9] Y.-M. Wei, L. Kang, B. Yang, and L.-d. Wu, ‘‘Applications of structure from motion: a survey,’’ *Journal of Zhejiang University*, vol. 14, no. 7, pp. 486–494, July 2013.
- [10] D. P. Tsakiris, P. Rives, and C. Samson, ‘‘Extending visual servoing techniques to nonholonomic mobile robots,’’ in *The confluence of vision and control*, ser. Lecture Notes in Control and Information Sciences, D. J. Kriegman, G. D. Hager, and A. S. Morse, Eds. Springer London, Jan. 1998, no. 237, pp. 106–117.
- [11] A. De and D. E. Koditschek, ‘‘Toward dynamical sensor management for reactive wall-following,’’ in *Proceedings of the 2013 IEEE Intl. Conference on Robotics and Automation*, May 2013.
- [12] K. C. Galloway, G. C. Haynes, B. D. Ilhan, A. M. Johnson, R. Knopf, G. Lynch, B. Plotnick, M. White, and D. E. Koditschek, ‘‘X-RHex: A highly mobile hexapedal robot for sensorimotor tasks,’’ University of Pennsylvania, Tech. Rep., 2010.
- [13] U. Saranlı, M. Buehler, and D. E. Koditschek, ‘‘Rhex: A simple and highly mobile hexapod robot,’’ *The International Journal of Robotics Research*, vol. 20, no. 7, p. 616, 2001.
- [14] G. A. D. Lopes and D. E. Koditschek, ‘‘Visual servoing for nonholonomically constrained three degree of freedom kinematic systems,’’ *The International Journal of Robotics Research*, vol. 26, no. 7, pp. 715–736, July 2007.
- [15] S. Skaff, G. Kantor, D. Maiwand, and A. A. Rizzi, *Inertial navigation and visual line following for a dynamical hexapod robot*, 2003, vol. 2, pp. 1808–1813.
- [16] E. Sontag, ‘‘Input to state stability: Basic concepts and results,’’ *Nonlinear and Optimal Control Theory*, pp. 163–220, 2008.
- [17] D. E. Koditschek, ‘‘Adaptive techniques for mechanical systems,’’ in *Proc. 5th. Yale Workshop on Adaptive Systems*, May 1987, pp. 259–265.
- [18] D. C. Conner, A. Rizzi, and H. Choset, ‘‘Construction and automated deployment of local potential functions for global robot control and navigation,’’ Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-03-22, Nov. 2003.
- [19] S. Revzen, B. Ilhan, and D. Koditschek, ‘‘Dynamical trajectory re-planning for uncertain environments,’’ in *2012 IEEE 51st Annual Conference on Decision and Control (CDC)*, 2012, pp. 3476–3483.
- [20] F. Bonin-Font, A. Ortiz, and G. Oliver, ‘‘Visual navigation for mobile robots: A survey,’’ *Journal of Intelligent and Robotic Systems*, vol. 53, no. 3, pp. 263–296, Nov. 2008.
- [21] A. D. Petiteville, S. Hutchinson, V. Cadenat, and M. Courdresses, ‘‘2D visual servoing for a long range navigation in a cluttered environment,’’ in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, 2011, pp. 5677–5682.
- [22] R. Brockett, ‘‘Asymptotic stability and feedback stabilization,’’ *Differential geometric control theory*, vol. 27, pp. 181–191, 1983.
- [23] C. Samson, ‘‘Velocity and torque feedback control of a nonholonomic cart,’’ in *Advanced robot control*. Springer, 1991, pp. 125–151.
- [24] J. Laumond, *Robot motion planning and control*, ser. Lecture notes in control and information sciences. Springer, 1998.
- [25] D. Folio and V. Cadenat, ‘‘A controller to avoid both occlusions and obstacles during a vision-based navigation task in a cluttered environment,’’ in *44th IEEE Conference on Decision and Control*, 2005, pp. 3898–3903.
- [26] J. Hauser, S. Sastry, and P. Kokotovic, ‘‘Nonlinear control via approximate input-output linearization: the ball and beam example,’’ *IEEE Transactions on Automatic Control*, vol. 37, no. 3, pp. 392–398, 1992.
- [27] D. E. Koditschek and E. Rimon, ‘‘Robot navigation functions on manifolds with boundary,’’ *Advances in Applied Mathematics*, vol. 11, no. 4, pp. 412–442, Dec. 1990.
- [28] E. Sontag and Y. Wang, ‘‘On characterizations of input-to-state stability with respect to compact sets,’’ in *Proc. IFAC Non-Linear Control Systems Design Symposium (NOLCOS95), Tahoe City, CA*, 1995, pp. 226–231.
- [29] D. Efimov, ‘‘Global lyapunov analysis of multistable nonlinear systems,’’ *SIAM Journal on Control and Optimization*, vol. 50, no. 5, pp. 3132–3154, Jan. 2012.
- [30] H. K. Khalil, *Nonlinear systems*. Prentice Hall Upper Saddle River, NJ, 2002.
- [31] N. Cowan, ‘‘Navigation functions on cross product spaces,’’ *IEEE*

Transactions on Automatic Control, vol. 52, no. 7, pp. 1297–1302, 2007.

- [32] D. E. Koditschek, “An approach to autonomous robot assembly,” *Robotica*(*Cambridge. Print*), vol. 12, pp. 137–155, 1994.
- [33] G. C. Haynes, A. A. Rizzi, and D. E. Koditschek, “Multistable phase regulation for robust steady and transitional legged gaits,” *The International Journal of Robotics Research*, 2012.