

Feedback Synthesis for Controllable Underactuated Systems using Sequential Second Order Actions

Giorgos Mamakoukas

Malcolm A. MacIver

Todd D. Murphey

Abstract—This paper derives nonlinear feedback control synthesis for general control affine systems using second-order actions—the needle variations of optimal control—as the basis for choosing each control response to the current state. A second result of the paper is that the method provably exploits the nonlinear controllability of a system by virtue of an explicit dependence of the second-order needle variation on the Lie bracket between vector fields. As a result, each control decision necessarily decreases the objective when the system is nonlinearly controllable using first-order Lie brackets. Simulation results using a differential drive cart, an underactuated kinematic vehicle in three dimensions, and an underactuated dynamic model of an underwater vehicle demonstrate that the method finds control solutions when the first-order analysis is singular. Moreover, the simulated examples demonstrate superior convergence when compared to synthesis based on first-order needle variations. Lastly, the underactuated dynamic underwater vehicle model demonstrates the convergence even in the presence of a velocity field.

I. INTRODUCTION

With many important applications in aerial or underwater missions, systems are underactuated either by design—in order to reduce actuator weight, expenses or energy consumption—or as a result of technical failures. In both cases, it is important to develop control policies that can exploit the nonlinearities of the dynamics, are general enough for this broad class of systems, and easily computable. Various approaches to nonlinear control range from steering methods using sinusoid controls [1], sequential actions of Lie bracket sequences [2] and backstepping [3], [4] to perturbation methods [5], sliding mode control (SMC) [6]–[8], intelligent [9], [10] or hybrid [11] control and nonlinear model predictive control (NMPC) methods [12]. These schemes have been successful on well-studied examples including, but not limited to, the rolling disk, the kinematic car, wheeling mobile robots, the Snakeboard, surface vessels, quadrotors, and cranes [13]–[28].

The aforementioned methods are not ideal in dealing with controllable systems. In the case of perturbations, the applied controls assume a future of control decisions that did not take the disturbance history into account; backstepping is generally ineffective in the presence of control limits and NMPC methods are typically computationally expensive. SMC methods suffer from chattering, which results in high energy consumption and instability risks by virtue of exciting unmodeled high-frequency dynamics [29], intelligent control methods are subject to data uncertainties [30], while other methods are often case-specific and will not hold for the level of generality encountered in robotics. We address this

limitation by using needle variations to compute feedback laws for general nonlinear systems affine in control, discussed next.

A. Needle Variations Advantages to Optimal Control

In this paper, we investigate using needle variation methods to find optimal control for nonlinear controllable systems. Needle variations consider the sensitivity of the cost function to infinitesimal application of controls and synthesize actions that reduce the objective [31], [32]. Such control synthesis methods have the advantage of efficiency in terms of computational effort, making them appropriate for online feedback (similar to other model predictive control methods, such as iLQG [33], but with the advantage—as shown here—of having provable formal properties over the entire state space). For time evolving objectives, as in the case of trajectory tracking tasks, controls calculated from other methods (such as sinusoids or Lie brackets for nonholonomic integrators) may be rendered ineffective as the target continuously moves to different states. In such cases, needle variation controls have the advantage of computing actions that directly reduce the cost, without depending on future control decisions. However, needle variation methods, to the best of our knowledge, have not yet considered higher than first-order sensitivities of the cost function.

We demonstrate analytically later in Section II that, by considering second-order needle variations, we obtain variations that explicitly depend on the Lie brackets between vector fields and, as a consequence, the higher-order nonlinearities in the system. Later, in Section III, we show that, for classically studied systems, such as the differential drive cart, this amounts to being able to guarantee that the control approach is *globally* certain to provide descent at every state, despite the conditions of Brockett’s theorem [34] on nonexistence of smooth feedback laws for such systems.

B. Paper Contribution and Structure

This paper derives the second-order sensitivity of the cost function with respect to infinitesimal duration of inserted control, which we will refer to interchangeably as the second-order mode insertion gradient or mode insertion Hessian (MIH). We relate the MIH expression to controllability analysis by revealing its underlying Lie bracket structure and present a method of using second-order needle variation actions to expand the set of states for which individual actions that guarantee descent of an objective function can be computed. Finally, we compute an analytical solution of controls that

uses the first two orders of needle variations. Due to length constraints, the details of some proofs are shortened to an acceptable length, allowing us to include examples demonstrating the method.

The content is structured as follows. In Section II, we prove that second-order needle variations guarantee control solutions for systems that are nonlinearly controllable using first-order Lie brackets. In Section III, we present an analytical control synthesis method that uses second-order needle actions. In Section IV, we implement the proposed synthesis method and present simulation results on a controllable, underactuated model of a 2D differential drive vehicle, a 3D controllable, underactuated kinematic rigid body and a 3D underactuated dynamic model of an underwater vehicle.

II. NEEDLE VARIATION CONTROLS BASED ON NON-LINEAR CONTROLLABILITY

In this section, we relate the controllability of systems to first- and second-order needle variation actions. After presenting the MIH expression, we reveal its dependence on Lie bracket terms between vector fields. Using this connection, we tie the descent property of needle variation actions to the controllability of a system and prove that second-order needle variation controls can produce control solutions for a wider set of the configuration state space than first-order needle variation methods. As a result, we are able to constructively compute, via an analytic solution, control formulas that are guaranteed to provide descent, provided that the system is controllable with first-order Lie brackets. Generalization to higher-order Lie brackets appears to have the same structure, but that analysis is postponed to future work.

A. Second-Order Mode Insertion Gradient

Consider a system with state $x : \mathbb{R} \mapsto \mathbb{R}^{N \times 1}$ and control $u : \mathbb{R} \mapsto \mathbb{R}^{M \times 1}$ with control-affine dynamics of the form

$$f(t, x(t), u(t)) = g(t, x(t)) + h(t, x(t))u(t), \quad (1)$$

where $g(t, x(t))$ is the drift vector field. Further consider a time period $[t_o, t_f]$ and control modes described by

$$\dot{x}(t) = \begin{cases} f_1(x(t), v), & t_0 \leq t < \tau \\ f_2(x(t), u), & \tau - \frac{\lambda}{2} \leq t < \tau + \frac{\lambda}{2} \\ f_1(x(t), v), & \tau + \frac{\lambda}{2} \leq t \leq t_f, \end{cases} \quad (2)$$

where f_1 and f_2 are the dynamics associated with *default* and *inserted* control v and u , respectively. Parameters λ and τ are the duration of the inserted dynamics f_2 and the switching time between the two modes. Dynamics of the form (2) are typically used in optimal control of hybrid systems to optimize the time scheduling of a-priori known modes [35]. Here, we use such dynamics to obtain a new control mode u that will optimally perturb the trajectory of any type of system with a needle action [36]. Given a cost function J of the form

$$J(x(t)) = \int_{t_o}^{t_f} l_1(x(t))dt + m(x(t_f)), \quad (3)$$

where $l_1(x(t))$ is the running cost and $m(x(t))$ the terminal cost, the mode insertion gradient (MIG) is

$$\frac{dJ}{d\lambda_+} = \rho^T (f_2 - f_1). \quad (4)$$

For brevity, the dependencies of variables are dropped. Although not presented here because of the length of the derivation and its similarity to [37], a similar analysis shows that, for dynamics that do not directly depend on the control duration, the mode insertion Hessian (MIH) is given by

$$\frac{d^2J}{d\lambda_+^2} = (f_2 - f_1)^T \Omega (f_2 - f_1) + \rho^T (D_x f_2 \cdot f_2 + D_x f_1 \cdot f_1 - 2D_x f_1 \cdot f_2) - D_x l_1 \cdot (f_2 - f_1), \quad (5)$$

where $\rho : \mathbb{R} \mapsto \mathbb{R}^{N \times 1}$ and $\Omega : \mathbb{R} \mapsto \mathbb{R}^{N \times N}$ are the first- and second-order adjoint states (costates). These quantities are calculated from the default trajectory and are given by

$$\begin{aligned} \dot{\rho} &= -D_x l_1^T - D_x f_1^T \rho \\ \dot{\Omega} &= -D_x f_1^T \Omega - \Omega D_x f_1 - D_x^2 l_1 - \sum_{i=1}^N \rho_i D_x^2 f^i, \end{aligned}$$

that are subject to $\rho(t_f) = D_x m(x(t_f))^T$ and $\Omega(t_f) = D_x^2 m(x(t_f))^T$. The superscript i in the dynamics f refers to the i^{th} element of the vector and is used to avoid confusion against default and inserted dynamics f_1 and f_2 , respectively.

B. Dependence of Second Order Needle Variations on Lie Bracket Structure

The Lie bracket of two vectors $f(x)$, and $g(x)$ is

$$[f, g](x) = \frac{\partial g}{\partial x} f(x) - \frac{\partial f}{\partial x} g(x),$$

which generates a control vector that points in the direction of the net infinitesimal change in state x created by infinitesimal noncommutative flow $\phi_\epsilon^f \circ \phi_\epsilon^g \circ \phi_\epsilon^{-f} \circ \phi_\epsilon^{-g} \circ x_0$, where ϕ_ϵ^f is the flow along a vector field f for time ϵ [2], [38]. Lie brackets are most commonly used for their connection to controllability [39], [40], but here they will show up in the expression describing the second-order needle variation.

We relate second-order needle variation actions to Lie brackets in order to provide controls that are conditional on the nonlinear controllability of a system. Let $h_i : \mathbb{R} \mapsto \mathbb{R}^{N \times 1}$ denote the column control vectors that compose $h : \mathbb{R} \mapsto \mathbb{R}^{N \times M}$ in (1) and $u_i \in \mathbb{R}$ be the individual control inputs. Then, we can express dynamics as

$$f = g + \sum_i^M h_i u_i.$$

and, for default control $v = 0$, we can re-write the MIH as

$$\begin{aligned} \frac{d^2J}{d\lambda_+^2} &= \left(\sum_{i=1}^M h_i u_i \right)^T \Omega \sum_{j=1}^M h_j u_j + \rho^T \left(\sum_{i=1}^M (D_x h_i u_i) \cdot g \right. \\ &\quad \left. - D_x g \cdot (h_i u_i) + \sum_{i=1}^M D_x h_i u_i \sum_{i=1}^M h_i u_i \right) - D_x l_1 \sum_{i=1}^M h_i u_i. \end{aligned}$$

Splitting the sum expression into diagonal ($i = j$) and off-diagonal ($i \neq j$) elements, and by adding and subtracting $2 \sum_i^M \sum_{j=1}^{i-1} (D_x h_i u_i)(h_j u_j)$, we can write

$$\begin{aligned} \sum_{i=1}^M D_x h_i u_i \sum_{i=1}^M h_i u_i &= \sum_i \sum_{j=1}^{i-1} [h_i, h_j] u_i u_j \\ &+ 2 \sum_i \sum_{j=1}^{i-1} (D_x h_i u_i)(h_j u_j) \\ &+ \sum_{i=j=1}^M (D_x h_i u_i)(h_i u_i). \end{aligned}$$

Then, we can express the MIH as

$$\begin{aligned} \frac{d^2 J}{d\lambda_+^2} &= \sum_{i=1}^M \sum_{j=1}^M u_i u_j h_i^T \Omega h_j + \rho^T \left(\sum_{i=2}^M \sum_{j=1}^{i-1} [h_i, h_j] u_i u_j \right. \\ &+ 2 \sum_{i=2}^M \sum_{j=1}^{i-1} (D_x h_i) h_j u_i u_j + \sum_{i=1}^M (D_x h_i) h_i u_i u_i \\ &\left. + \sum_{i=1}^M [g, h_i] u_i \right) - D_x l \left(\sum_{i=1}^M h_i u_i \right). \end{aligned}$$

The expression contains Lie bracket terms of the control vectors that appear in the system dynamics, indicating that second-order needle variations consider higher-order nonlinearities. By associating the MIH to Lie brackets, we next prove that second-order needle variation actions can guarantee decrease of the objective for certain types of controllable systems.

C. Existence of Control Solutions with First- and Second-Order Mode Insertion Gradients

In this section, we prove that the first two orders of the mode insertion gradient can be used to guarantee controls that reduce objectives of the form (3) for systems that are controllable with first-order Lie brackets. The analysis is applicable to optimization problems that satisfy the following assumptions.

Assumption 1. The vector elements of dynamics f_1 and f_2 are real, bounded, C^2 in x , and C^0 in u and t .

Assumption 2. The incremental cost $l_1(x)$ is real, bounded, and C^2 in x . The terminal cost $m(x(t_f))$ is real and twice differentiable with respect to $x(t_f)$.

Assumption 3. Default and inserted controls v and u are real, bounded, and C^0 in t .

Under Assumptions 1-3, the MIG and MIH expressions are well-defined. Then, as we show next, there are control actions that can improve any objective that is not a local optimizer.

Definition 1. A local optimizer of the cost function (3) is given by a set (x^*, u^*) if and only if the set describes a trajectory that corresponds to an objective function $J(x^*(t))$ for which $D_x J(x^*(t)) = 0$.

Proposition 1. Consider a set (x, v) that describes the state and default control of (2). If $(x, v) \neq (x^*, v^*)$, then the first-order adjoint ρ is a non-zero vector.

Proof: Using (3),

$$\begin{aligned} x \neq x^* &\Rightarrow D_x J(x(t)) \neq 0 \\ &\Rightarrow \int_{t_0}^{t_f} D_x l_1(x(t)) dt + D_x m(x(t_f)) \neq 0 \\ &\Rightarrow \int_{t_0}^{t_f} D_x l_1(x(t)) dt \neq 0 \text{ OR } D_x m(x(t_f)) \neq 0 \\ &\Rightarrow D_x l_1(x(t)) \neq 0 \text{ OR } D_x m(x(t_f)) \neq 0 \\ &\Rightarrow \dot{\rho} \neq 0 \vee \rho(t_f) \neq 0. \end{aligned}$$

Therefore, if $x \neq x^*$, then $\exists t \in [t_0, t_f]$ such that $\rho \neq 0$. ■

Proposition 2. Consider dynamics given by (2) and a pair of state and control $(x, v) \neq (x^*, v^*)$ such that $\frac{dJ}{d\lambda_+} = 0 \forall u \in \mathbb{R}^M$ and $\forall t \in [t_0, t_f]$. Then, the first-order adjoint ρ is orthogonal to all control vectors h_i .

Proof: The linear combination of the elements of a vector x is always zero if and only if x is the zero vector. Given that, rewrite (4) as

$$\begin{aligned} \frac{dJ}{d\lambda_+} = 0 &\Rightarrow \rho^T \sum_i^M h_i (u_i - v_i) = 0 \\ &\Rightarrow \sum_i^M k_i w_i = 0 \forall w_i, \end{aligned}$$

where $w_i = (u_i - v_i)$ and $k_i = \rho^T h_i \in \mathbb{R}$. The linear combination of the elements of k is zero for any w_i , which means k must be the zero vector. By Proposition 1, $\rho \neq 0$ for a non-optimizer pair of state and control and, as a result, $\rho^T h_i = 0 \forall i \in [1, M]$. ■

Proposition 3. Consider dynamics given by (2) and a pair of state and control $(x, v) \neq (x^*, v^*)$ such that $\frac{dJ}{d\lambda_+} = 0 \forall u \in \mathbb{R}^M$ and $\forall t \in [t_0, t_f]$. Further assume that the control vectors h_i and their Lie Bracket terms $[h_i, h_j]$ span the state space \mathbb{R}^N . Then, there exist i and j such that $\rho^T [h_i, h_j] \neq 0$.

Proof: A set of vectors $S = (v_1, \dots, v_M)$ is linearly independent if and only if every vector $r \in \text{span}(S)$ can be uniquely written as a linear combination of (v_1, \dots, v_M) . The control vectors and their Lie Brackets span the \mathbb{R}^N space. On that assumption, it follows that any N -dimensional vector can be expressed in terms of the control vectors and their Lie Brackets. The first-order adjoint is an N -dimensional vector, which is non-zero for a non-optimizer pair of x, v by Proposition 1. Therefore, it can be expressed as

$$\rho = c_1 h_1 + \dots + c_M h_M + \sum_{i \neq j}^M [h_i, h_j] \neq 0. \quad (6)$$

Given that $\frac{dJ}{d\lambda_+} = 0$, and by Proposition 2, ρ is orthogonal to all control vectors h_i (which also implies that the control

vectors h_i do not span \mathbb{R}^N). Then, left-multiplying (6) by ρ^T yields

$$\rho^T \rho = \sum_{i \neq j}^M \rho^T [h_i, h_j] \neq 0.$$

It follows that there is at least one Lie bracket term $[h_i, h_j]$ that is not orthogonal to the costate ρ . ■

Proposition 4. *Consider dynamics given by (2) and a trajectory described by state and control $(x, v) \neq (x^*, v^*)$, then there are always control solutions $u \in \mathbb{R}^M$ such that $\frac{dJ}{d\lambda_+} \leq 0$ for some $t \in [t_o, t_f]$.*

Proof: Using dynamics of the form in (1), the expression of the mode insertion gradient can be written as

$$\frac{dJ}{d\lambda_+} = \rho^T (f_2 - f_1) = \rho^T (h(u - v)).$$

By Proposition 1, $\rho \neq 0$ for a non-optimizer trajectory. Given controls u and v that generate a positive mode insertion gradient, there always exist control u' such that the mode insertion gradient is negative, i.e. $u' - v = -(u - v)$. The mode insertion gradient is zero for all $u \in \mathbb{R}^M$ if and only if the costate vector is orthogonal to each control vector h_i ¹. ■

First-order needle variation methods are singular when the mode insertion gradient is zero. When that is true, the second-order mode insertion gradient is guaranteed to be negative for systems that are controllable with first-order Lie Brackets, which in turn implies that a control solution can be found with second-order needle variation methods.

Proposition 5. *Consider dynamics given by (2) and a trajectory described by state and control $(x, v) \neq (x^*, v^*)$ such that $\frac{dJ}{d\lambda_+} = 0$ for all $u \in \mathbb{R}^M$ and $t \in [t_o, t_f]$. If the control vectors h_i and the Lie brackets $[h_i, h_j]$ and $[g, h_i]$ span the state space (\mathbb{R}^N), then there always exist control solutions $u \in \mathbb{R}^M$ such that $\frac{d^2J}{d\lambda_+^2} < 0$.*

Proof: Let $k \in [1, M]$ be an index chosen such that $[h_i, h_k]$ for some $i \in [1, M] \setminus \{k\}$ ² is a vector that is linearly independent of all control vectors $h_i \forall i \in [1, M]$. The proof then considers controls such that $u_j = v_i \forall j, i \neq k$ and $v_k = 0$ and expresses the MIH expression (5) as

$$\frac{d^2J}{d\lambda_+^2} = u^T \mathcal{G} u - u_k ((D_x l_1) h_k - \rho^T [g, h_k]),$$

where $\mathcal{G}_{ij} = 0 \forall i, j \in [1, M] \setminus \{k\}$, $\mathcal{G}_{ik} = \mathcal{G}_{ki} = \frac{1}{2} [h_i, h_k]$, and $\mathcal{G}_{kk} = h_k^T \Omega h_k + \rho^T D_x h_k \cdot h_k$. The matrix \mathcal{G} is shown to be either indefinite or negative semidefinite if there exists a Lie bracket term $[h_i, h_k]$ such that $\rho^T [h_i, h_k] \neq 0$. If, on the other hand, $\rho^T [h_i, h_k] = 0$, by reasoning that is similar to

¹If the control vectors span the state space \mathbb{R}^N , the costate vector $\rho \in \mathbb{R}^N$ cannot be orthogonal to each of them. Therefore, for first-order controllable (fully actuated) systems, there always exist controls for which the cost can be reduced to first order.

²The notation \setminus indicates that the element k is subtracted from the set $[1, M]$.

Proposition 3, there is at least a Lie bracket $[g, h_k] \neq 0$ and the MIH expression reduces to a quadratic in u_k . In either case, it then becomes straightforward to show that there exist controls for which the MIH expression is negative. ■

Theorem 1. *Consider dynamics given by (2) and a trajectory described by state and control $(x, v) \neq (x^*, v^*)$. If the control vectors h_i and the Lie brackets $[h_i, h_j]$ and $[g, h_i]$ span the state space (\mathbb{R}^N), then there always exists a control vector $u \in \mathbb{R}^M$ and a duration λ such that the cost function (3) can be reduced.*

Proof: The local change of the cost function (3) due to inserted control u of duration λ can be approximated with a Taylor series expansion

$$J(\lambda) - J(0) \approx \lambda \frac{dJ}{d\lambda_+} + \frac{\lambda^2}{2} \frac{d^2J}{d\lambda_+^2}.$$

By Propositions 4 and 5, either 1) $\frac{dJ}{d\lambda_+} < 0$ or 2) $\frac{dJ}{d\lambda_+} = 0$ and $\frac{d^2J}{d\lambda_+^2} < 0$. Therefore, there always exist controls that reduce the cost function (3) to first or second order. ■

III. CONTROL SYNTHESIS

In this section, we present an analytical solution of first- and second-order needle variation controls that reduce the cost function (3) to second order. We then describe the algorithmic steps of the feedback scheme used in the simulation results of this paper.

A. Analytical Solution for Second Order Actions

For underactuated systems, there are states at which ρ is orthogonal to the control vectors h_i (see Proposition 4). At these states, control calculations based only on first-order sensitivities fail, while controls based on second-order information still decrease the objective provided that the control vectors and their Lie brackets span the state space (see Theorem 1). We use this property to compute an analytical synthesis method that expands the set of states for which individual actions that guarantee descent of an objective function can be computed.

Consider the Taylor series expansion of the cost around control duration λ . Given the expressions of the first- and second-order mode insertion gradients, we can write the cost function (3) as a Taylor series expansion around the infinitesimal duration λ of inserted control u :

$$J(\lambda) \approx J(0) + \lambda \frac{dJ}{d\lambda_+} + \frac{\lambda^2}{2} \frac{d^2J}{d\lambda_+^2}.$$

The first- and second-order mode insertion gradients used in the expression are functions of the inserted control $u(t)$ in (2). For a fixed λ , we can minimize the function using Newton's Method to update the control actions. Control solutions that minimize the Taylor expansion of the cost will have the form

$$u^*(t) = \underset{u}{\operatorname{argmin}} J(0) + \lambda \frac{dJ}{d\lambda_+} + \frac{\lambda^2}{2} \frac{d^2J}{d\lambda_+^2} + \frac{1}{2} \|u\|_R^2, \quad (7)$$

where the MIH has both linear and quadratic terms in $u(t)$. The time dependence of the control u is purposefully used

here to emphasize that control solutions are functions of time t . Using the Gâteaux derivative, we computed the minimizer of (7) to be

$$u^*(t) = \left[\frac{\lambda^2}{2} \Gamma + R \right]^{-1} \left[\frac{\lambda^2}{2} \Delta + \lambda(-h^T \rho) \right], \quad (8)$$

where $\Delta : \mathbb{R} \mapsto \mathbb{R}^{M \times 1}$ and $\Gamma : \mathbb{R} \mapsto \mathbb{R}^{M \times M}$ are respectively the first- and second-order derivatives of $d^2 J / d\lambda_+^2$ with respect to the control u and are given by

$$\begin{aligned} \Delta &\triangleq \left[h^T (\Omega^T + \Omega) h + 2h^T \cdot \left(\sum_{k=1}^n (D_x h_k) \rho_k \right)^T \right] v \\ &\quad + (D_x g \cdot h)^T \rho - \left(\sum_{k=1}^n (D_x h_k) \rho_k \right) \cdot g + h^T D_x l^T \\ \Gamma &\triangleq [h^T (\Omega^T + \Omega) h + h^T \cdot \left(\sum_{k=1}^n (D_x h_k) \rho_k \right)^T + \sum_{k=1}^n (D_x h_k) \rho_k \cdot h]^T. \end{aligned}$$

The parameter R denotes a metric on control effort.

The existence of control solutions in (8) depend on the inversion of the Hessian $H = \frac{\lambda^2}{2} \Gamma + R$. To ensure H is positive definite, we implement a spectral decomposition on the Hessian $H = V D V^{-1}$, where matrices V and D contain the eigenvectors and eigenvalues of H , respectively. We replace all elements of the diagonal matrix D that are smaller than ϵ with ϵ to obtain \bar{D} and replace H with $\bar{H} = V \bar{D} V^{-1}$ in (8). We prefer the spectral decomposition approach to the Levenberg-Marquardt method ($\bar{H} = H + \kappa I \succ 0$), because the latter affects all eigenvalues of the Hessian and further distorts the second-order information. At saddle points, we set the control equal to the eigenvector of H that corresponds to the most negative eigenvalue in order to descend along the direction of most negative curvature [41]–[44].

This synthesis technique provides controls at time t that guarantee to reduce the cost function (3) for systems that are controllable using first-order Lie brackets. Control solutions are computed solely by forward simulating the state over a time horizon T and backward simulating the first- and second-order costates ρ and Ω . As we see next, this leads to a very natural, and easily implementable, algorithm for applying cost-based feedback.

B. Algorithmic Description of Control Synthesis Method

The proposed second-order analytical controls presented in (8) are implemented in a series of steps shown in Algorithm 1. We compare first- and second-order needle variation actions by implementing different controls in Step 2 of Algorithm 1. For the first-order case, we implement controls that are the solution to a minimization problem of the first-order sensitivity of the cost function (3) and the control effort

$$\begin{aligned} u^*(t) &= \min_u \frac{1}{2} \left(\frac{dJ_1}{d\lambda_+} - \alpha_d \right)^2 + \frac{1}{2} \|u\|_R^2 \\ &= (\Lambda + R^T)^{-1} (\Lambda v + h^T \rho \alpha_d), \end{aligned} \quad (9)$$

where $\Lambda \triangleq h^T \rho \rho^T h$ and $\alpha_d \in \mathbb{R}^-$ expresses the desired value of the mode insertion gradient term (see, for example,

Algorithm 1

1. Simulate states and costates with default dynamics f_1 over a time horizon T
 2. Compute optimal needle variation controls
 3. Saturate controls
 4. Find the insertion time that corresponds to the most negative mode insertion gradient
 5. Use a line search to find control duration that ensures reduction of the cost function (3)
-

[45]). Typically, $\alpha_d = \gamma J_o$, where J_o is the cost function (3) computed using default dynamics f_1 . For second-order needle variation actions, we compute controls using (8).

C. Comparison to Alternative Optimization Approaches

Algorithm 1 differs from controllers that compute control sequences over the entire time horizon in order to locally minimize the cost function. Rather, the proposed scheme utilizes the time-evolving sensitivity of the objective to infinitesimal switched dynamics and searches in a one-dimensional space for a finite duration of a single action that will optimally improve the cost. It does so using a closed-form expression and, as a result, it avoids the expensive iterative computational search in high-dimensional spaces, while it may still get closer to the optimizer with one iterate.

First-order needle variation solution are shown in (9) to exist globally, demonstrate a larger region of attraction and have a less complicated representation on Lie Groups [46]. These traits naturally transfer to second-order needle controls (8) that also contain the first-order information that is present in (9). In addition, as this paper demonstrates, the suggested second-order needle variation controller has formal guarantees of descent for systems that are controllable with first-order Lie brackets, which—to the best of our knowledge—is not provided by any alternative method. Given these benefits, the authors propose second-order needle variation actions as a complement to existing approaches for time-sensitive robotic applications that may be subject to large initial error, Euler angle singularities, or fast-evolving (and uncertain) objectives.

Next, we implement Algorithm 1 using first or second-order needle variation controls (shown in (9) and (8), respectively) to compare them in terms of convergence success on various underactuated systems.

IV. SIMULATION RESULTS

The proposed synthesis method is implemented on three underactuated examples, the differential drive cart, a 3D kinematic rigid body and a dynamic model of an underwater vehicle. The kinematic systems of a 2D differential drive and a 3D rigid body are controllable using first-order Lie brackets of the vector fields and help verify Theorem 1. The underactuated dynamic model of a 3D rigid body serves to compare controls in (8) and (9) in a more sophisticated environment. In all simulation results, we start with default control $v = 0$ and

an objective function of the form

$$J(x(t)) = \frac{1}{2} \int_{t_0}^{t_f} \|\dot{\vec{x}}(t) - \dot{\vec{x}}_d(t)\|_Q^2 dt + \frac{1}{2} \|\vec{x}(t_f) - \vec{x}_d(t_f)\|_{P_1}^2,$$

where \vec{x}_d is the desired state-trajectory, and $Q = Q^T \geq 0$, $P_1 = P_1^T \geq 0$ are metrics on state error.

A. 2D Kinematic Differential Drive

The differential drive system demonstrates that controls shown in (9) that are based only on the first-order sensitivity of the cost function (3) can be insufficient for controllable systems, contrary to controls shown in (8) that guarantee decrease of the objective for systems that are controllable using first-order Lie brackets (see Theorem 1). The system states are its coordinates and orientation, given by $s = [x, y, \theta]^T$, with kinematic ($g = 0$) dynamics

$$f = r \begin{bmatrix} \cos(\theta) & \cos(\theta) \\ \sin(\theta) & \sin(\theta) \\ \frac{1}{L} & -\frac{1}{L} \end{bmatrix} \begin{bmatrix} u_R \\ u_L \end{bmatrix},$$

where $r = 3.6$ cm, $L = 25.8$ cm denote the wheel radius and the distance between them, and u_R, u_L are the right and left wheel control angular velocities, respectively (these parameter values match the specifications of the iRobot Roomba).

The control vectors h_1, h_2 and their Lie bracket term $[h_1, h_2] = 2 \frac{r^2}{L^2} [-\sin(\theta), -\cos(\theta)]^T$ span the state space (\mathbb{R}^3). Therefore, from Theorem 1, there always exist controls that reduce the cost to first or second order.

Fig. 1 demonstrates how different first- and second-order needle variation actions perform on reaching a nearby target. Actions based on first-order needle variations (9) do not generate solutions that turn the vehicle, but rather drive it straight until the orthogonal displacement between the system and the target location is minimized. Actions based on second-order needle variations (8), on the other hand, converge successfully.

We also present a Monte Carlo simulation that compares convergence success using first- and second-order needle variations controls shown in (9) and (8), respectively. We sampled over initial coordinates $x_0, y_0 \in [-1500, 1500]$ mm using a uniform distribution and keeping only samples for which the initial distance from the origin exceeded $L/5$; $\theta_0 = 0$ for all samples. Successful samples were within $L/5$ from the origin with an angle $\theta < \pi/12$ within 60 seconds using feedback sampling rate of 4 Hz. Results were generated using $Q = \text{diag}(10, 10, 1000)$, $P_1 = \text{diag}(0, 0, 0)$, $T = 0.5$ s, $R = \text{diag}(100, 100)$ for (9), $R = \text{diag}(0.1, 0.1)$ for (8), $\gamma = -15$, $\lambda = 0.1$ and saturation limits on the angular velocities of each wheel $\pm 150/36$ mm/s³. As shown in Fig. 2, the system always converges to the target using second-order needle variation actions, matching the theory.

³The metric on control effort is necessarily smaller for (8), due to parameter λ . The parameter was chosen carefully to ensure that control solutions from (8) and (9) were comparable in magnitude.

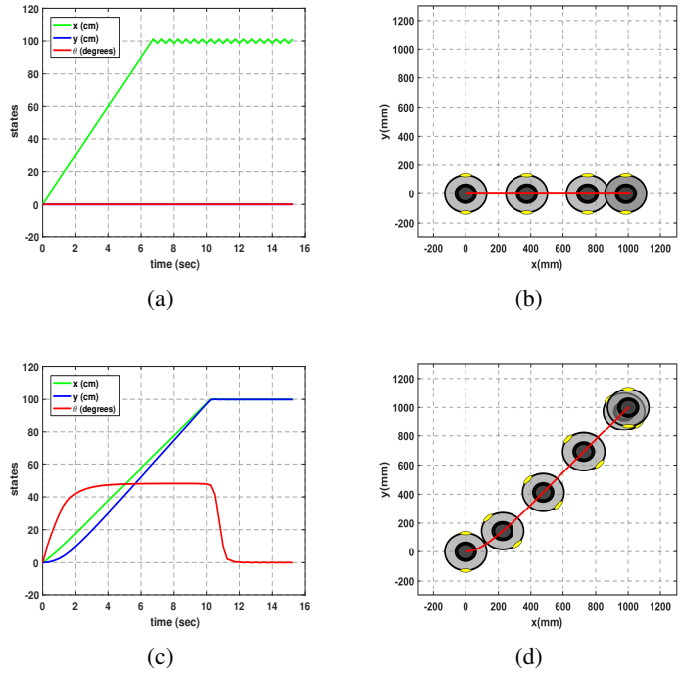


Fig. 1: Differential drive using first- (top) and second-order (bottom) needle variation actions. Snapshots of the system are shown at $t = 0, 2.5, 5, 7.5, 10,$ and 12.5 sec. The target state is $[x_d, y_d, \theta_d] = [1000 \text{ mm}, 1000 \text{ mm}, 0]$.

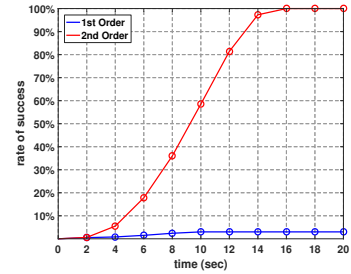


Fig. 2: Convergence success rates of first- (9) and second-order (8) needle variation controls for the kinematic differential drive model. Simulation runs: 1000.

B. 3D Kinematic Rigid Body

The underactuated kinematic rigid body is a three dimensional example of a system that is controllable with first-order Lie brackets. To avoid singularities in the state space, the orientation of the system is expressed in quaternions [47], [48]. The states are $s = [x, y, z, q_0, q_1, q_2, q_3]$, where $b = [x, y, z]$ are the world-frame coordinates and $q = [q_0, q_1, q_2, q_3]$ are unit quaternions. Dynamics $f = [\dot{b}, \dot{q}]^T$ are given by

$$\dot{b} = R_q v, \quad (10)$$

$$\dot{q} = \frac{1}{2} \begin{bmatrix} -q_1 & -q_2 & -q_3 \\ q_0 & -q_3 & q_2 \\ q_3 & q_0 & -q_1 \\ -q_2 & q_1 & q_0 \end{bmatrix} \omega, \quad (11)$$

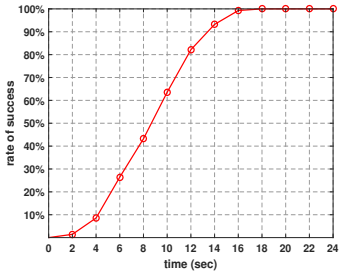


Fig. 3: Convergence success rates of second-order needle variation controls (8) for the underactuated kinematic vehicle. First-order actions (9) do not affect the y -coordinate of the rigid body and therefore never converge. Simulation runs: 280.

where v and ω are the body frame linear and angular velocities, respectively [49]. The rotation matrix for quaternions is

$$R_q = \begin{bmatrix} q_0^2 + q_1^2 - q_2^2 - q_3^2 & 2(q_1q_2 - q_0q_3) & 2(q_1q_3 + q_0q_2) \\ 2(q_1q_2 + q_0q_3) & q_0^2 - q_1^2 + q_2^2 - q_3^2 & 2(q_2q_3 - q_0q_1) \\ 2(q_1q_3 - q_0q_2) & 2(q_2q_3 + q_0q_1) & q_0^2 - q_1^2 - q_2^2 + q_3^2 \end{bmatrix}.$$

The system is kinematic: $v = F$ and $\omega = T$, where $F = (F_1, F_2, F_3)$ and $T = (T_1, T_2, T_3)$ describe respectively the surge, sway, and heave input forces, and the roll, pitch, and yaw input torques. We render the rigid body underactuated by removing the sway and yaw control authorities ($F_2 = T_3 = 0$).

The four control vectors span a four dimensional space. First order Lie bracket terms add two more dimensions to span the state space (\mathbb{R}^6) (the fact that there are seven states in the model of the system is an artifact that is inherent in quaternion representation; it does not affect controllability, given that there is also a constraint that the norms of quaternions must sum up to one). The vectors $h_1, h_2, [h_2, h_3]$ span \mathbb{R}^3 associated with the world frame coordinate dynamics \dot{x}, \dot{y} , and \dot{z} . Similarly, vectors h_3, h_4 , and $[h_4, h_3]$ also span \mathbb{R}^3 . Thereby, control vectors and their first-order Lie brackets span the state space and, from Theorem 1, optimal actions shown in (8) will always reduce the cost function (3).

To verify the theory, we present the convergence success of the system on 3D motion (see Fig. 3). Using Monte Carlo sampling with uniform distribution, initial locations were randomly generated such that $x_0, y_0, z_0 \in [-50, 50]$ cm keeping only samples for which the initial distance from the origin exceeded 6 cm. We regarded as a convergence success each trial in which the rigid body was within 6 cm to the origin by the end of 60 seconds at any orientation. Results were generated at a sampling rate of 20 Hz using $Q = 0$, $P_1 = \text{diag}(100, 200, 100, 0, 0, 0, 0)$, $T = 1.0$ s, $\gamma = -50000$, $\lambda = 10^{-3}$, $R = 10^{-6} \text{diag}(1, 1, 100, 100)$ for (8), and $R = \text{diag}(10, 10, 1000, 1000)$ for controls in (9). Controls were saturated at ± 10 cm/s for the linear velocities and ± 10 rad/s for the angular ones. As shown in Fig. 3, and as expected, all locomotion trials were successful.

C. Underactuated Dynamic 3D Fish

We represent the three dimensional rigid body with states $s = [b, q, v, \omega]^T$, where $b = [x, y, z]$ are the world-frame coordinates, $q = [q_0, q_1, q_2, q_3]$ are the quaternions that

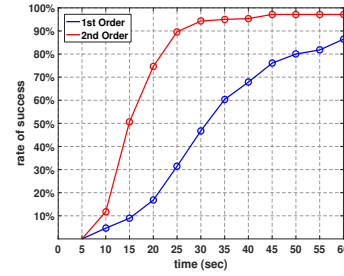


Fig. 4: Convergence success rates of first- and second-order needle variation controls (9) and (8), respectively) for the underactuated dynamic vehicle model. Simulation runs: 280

describe the world-frame orientation, and $v = [v_x, v_y, v_z]$ and $\omega = [\omega_x, \omega_y, \omega_z]$ are the body-frame linear and angular velocities. The rigid body dynamics are given by \dot{b} and \dot{q} shown in (10) and (11) and

$$\begin{aligned} M\dot{v} &= Mv \times \omega + F, \\ J\dot{\omega} &= J\omega \times \omega + T, \end{aligned}$$

where the effective mass and moment of inertia of the rigid body are given by $M = \text{diag}(6.04, 17.31, 8.39)$ g and $J = \text{diag}(1.57, 27.78, 54.11)$ g-cm², respectively. This example is inspired by work in [45], [50] and the parameters used for the effective mass and moment of inertia of a rigid body correspond to measurements of a fish. The control inputs are $F_2 = T_3 = 0$ and $F_3 \geq 0$.

The control vectors only span a four dimensional space and, since they are state-independent, their Lie brackets are zero vectors. However, the Lie brackets containing the drift vector field g (that also appear in the MIH expression) add from one to four (depending on the states) independent vectors such that control solutions in (8) guarantee decrease of the cost function (3) for a wider set of states than controls in (9).

Simulation results based on Monte Carlo sampling are shown in Fig. 4. Initial coordinates x_0, y_0, z_0 were generated using a uniform distribution in $[-100, 100]$ cm, discarding samples for which the initial distance to the origin was less than 15 cm. Successful trials where the ones for which, within a simulation window of 60 seconds, the system approached within 5 cm to the origin (at any orientation) and whose magnitude of the linear velocities was, at the same time, less than 5 cm/s. Results were generated at a sampling rate of 20 Hz using $T=1.5$ s, $P_1=0$, $Q = \frac{1}{200} \text{diag}(10^3, 10^3, 10^3, 0, 0, 0, 0, 1, 1, 1, 2 \cdot 10^3, 10^3, 10^3)$, $\gamma = -5$, $R = \text{diag}(10^3, 10^3, 10^6, 10^6)$ for (9), $R = \frac{1}{2} \text{diag}(10^{-6}, 10^{-6}, 10^{-3}, 10^{-3})$ for (8), and $\lambda = 10^{-4}$. The same control saturations ($F_1 \in [-1, 1]$ mN, $F_3 \in [0, 1]$ mN, $T_1 \in [-0.1, 0.1]$ μ N-m, and $T_2 \in [-0.1, 0.1]$ μ N-m) were used for all simulations of the dynamic 3D fish. As shown in Fig. 4, controls computed using second-order needle variations converge faster than those based on first-order needle variations, and 97% of trials converge within 60 seconds.

Both methods converge over time to the desired location; as the dynamic model of the rigid body tumbles around and

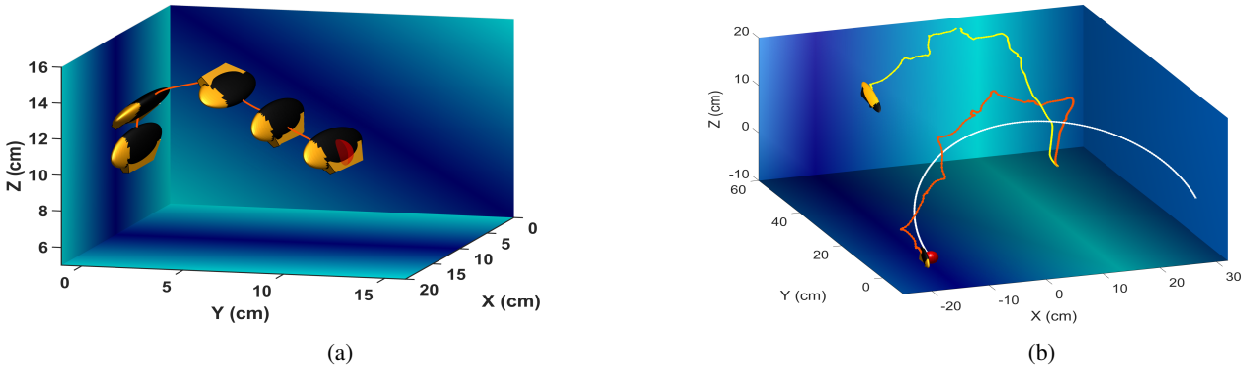


Fig. 5: Figure 5a shows snapshots of a parallel displacement maneuver using an underactuated dynamic vehicle model with second-order controls given by (8); first-order solutions (9) are singular throughout the simulation. Figure 5b shows tracking performance of the same system in the presence of +10 cm/s \hat{y} fluid drift. The yellow system corresponds to first-order needle variation actions; the red one to second order. The target trajectory (red ball) is indicated with white traces over a 10-second simulation. Animation of these results is available at <https://vimeo.com/219628387>.

its orientation changes, possible descent directions of the cost function (3) change and the control is able to push the system to the target. Controls for the first-order needle variation case (9) are singular for a wider set of states than second-order needle variation controls (8) and, for this reason, they benefit more from tumbling. In a 3D parallel locomotion task, only second-order variation controls (8) manage to provide control solutions through successive heave and roll inputs, whereas controls based on first-order sensitivities (9) fail (see Fig. 5a).

As controls in (8) are non-singular for a wider subset of the configuration state space than the first-order solutions in (9), they will provide more actions over a period of time and keep the system closer to a time-varying target. Fig. 5b demonstrates the superior trajectory tracking behavior of controls based on (8) in the presence of +10 cm/s \hat{y} fluid drift. The trajectory of the target is given by $[x, y, z] = [20 + 10 \cos(\frac{1}{5}) \cos(\frac{3t}{10}), 20 + 10 \cos(\frac{1}{5}) \sin(\frac{3t}{10}), 10 \sin(\frac{2t}{5})]$, with $T=2$ s, $\lambda=0.01$, $Q = \text{diag}(10, 10, 10, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0.1)$, $\gamma = -50000$, $P_1 = \text{diag}(10, 10, 10, 0, 0, 0, 0, 0, 0, 0, 0, 0)$, $R = \text{diag}(10^3, 10^3, 10^6, 10^6)$ for (9), and $R = \text{diag}(10, 10, 10^4, 10^4)$ for (8). The simulation runs in real time using a C++ implementation on a laptop with Intel® Core™ i5-6300HQ CPU @2.30GHz and 8GB RAM. The drift is known for both first- and second-order systems and accounted for in their dynamics in the form of $\dot{b} = \dot{b} + \dot{b}_{\text{drift}}$, where \dot{b}_{drift} is a vector that points in the direction of the fluid flow. Simulation results demonstrate superior tracking of second-order needle variation controls that manage to stay with the target, whereas, in the meantime, the system that corresponds to first-order needle variation controls is being drifted away by the flow.

We also tested convergence success of the +10 cm/s \hat{y} drift case. Initial conditions x, y, z were sampled uniformly from a 30 cm radius from the origin, discarding samples for which the initial distance was less than 5 cm. We consider samples to be successful if, during 60 seconds of simulation, they approached the origin within 5 cm. Out of 500 samples, controls based on second-order variations converged 91% of the time (with average convergence time

of 5.87 s), compared to 89% for first-order actions (with average convergence time of 9.3 s). Simulation parameters were $T=1$ s, $\lambda=10^{-4}$, $Q=10^{-3} \text{diag}(10, 10, 10, 0, 0, 0, 0, 0, 1, 1, 1, 1)$, $P_1 = \text{diag}(100, 100, 100, 0, 0, 0, 0, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}, 0, 0)$, $\gamma = -25000$, $R = \text{diag}(0.1, 0.1, 10^4, 10^4)$ for (9), and $R = \frac{1}{2} \text{diag}(10^{-5}, 10^{-5}, 1, 1)$ for (8).

V. CONCLUSION

This paper presents a needle variation control synthesis method for nonlinearly controllable systems that can be expressed in control affine form. Control solutions provably exploit the nonlinear controllability of a system and, contrary to other nonlinear feedback schemes, have formal guarantees to decrease the objective. By optimally perturbing the system with needle actions, the proposed algorithm avoids the expensive iterative computation of controls over the entire horizon that other NMPC methods use and is able to run in real time.

Simulation results on three underactuated systems compare first- and second-order needle variation controls and demonstrate the superior convergence success rate of the proposed feedback synthesis. Because second-order needle variation actions are non-singular for a wider set of the state space than controls based on first-order sensitivity, they are also more suitable for time-evolving objectives, as demonstrated by the trajectory tracking examples in this paper. Second order needle variation controls are also calculated at little computational cost and preserve control effort. These traits, demonstrated in the simulation examples of this paper, render feedback synthesis based on second- and higher-order needle variation methods a promising alternative feedback scheme for underactuated and nonlinearly controllable systems.

ACKNOWLEDGMENTS

This work was supported by the Office of Naval Research under grant ONR N00014-14-1-0594. Any opinions, findings, and conclusions or recommendations expressed here are those of the authors and do not necessarily reflect the views of the Office of Naval Research.

REFERENCES

- [1] R. M. Murray and S. S. Sastry, "Nonholonomic motion planning: Steering using sinusoids," *IEEE Transactions on Automatic Control*, vol. 38, no. 5, pp. 700–716, 1993.
- [2] R. M. Murray, Z. Li, S. S. Sastry, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. CRC press, 1994, ch. 18.
- [3] P. V. Kokotovic, "The joy of feedback: nonlinear and adaptive," *IEEE Control systems*, vol. 12, no. 3, pp. 7–17, 1992.
- [4] D. Seto and J. Baillieul, "Control problems in super-articulated mechanical systems," *IEEE Transactions on Automatic Control*, vol. 39, no. 12, pp. 2442–2453, 1994.
- [5] J. L. Junkins and R. C. Thompson, "An asymptotic perturbation method for nonlinear optimal control problems," *Journal of Guidance, Control, and Dynamics*, vol. 9, no. 4, pp. 391–396, 1986.
- [6] W. Perruquetti and J.-P. Barbot, *Sliding Mode Control in Engineering*. New York: Marcel Dekker, 2002.
- [7] V. I. Utkin, *Sliding Modes in Control and Optimization*. Springer-Verlag, 1992.
- [8] R. Xu and Ü. Özgüner, "Sliding mode control of a class of underactuated systems," *Automatica*, vol. 44, no. 1, pp. 233–241, 2008.
- [9] S. C. Brown and K. M. Passino, "Intelligent control for an acrobat," *Journal of Intelligent & Robotic Systems*, vol. 18, no. 3, pp. 209–248, 1997.
- [10] C. J. Harris, C. G. Moore, and M. Brown, *Intelligent Control: Aspects of Fuzzy Logic and Neural Nets*. Singapore: World Scientific, 1993, vol. 6.
- [11] R. Fierro, F. L. Lewis, and A. Lowe, "Hybrid control for a class of underactuated mechanical systems," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 29, no. 6, pp. 649–654, 1999.
- [12] F. Allgöwer, R. Findeisen, and Z. K. Nagy, "Nonlinear model predictive control: From theory to application," *Journal of the Chinese Institute of Chemical Engineers*, vol. 35, no. 3, pp. 299–315, 2004.
- [13] F. Bullo, N. E. Leonard, and A. D. Lewis, "Controllability and motion algorithms for underactuated Lagrangian systems on Lie groups," *IEEE Transactions on Automatic Control*, vol. 45, no. 8, pp. 1437–1454, 2000.
- [14] Y. Lin, D. Wu, X. Wang, X. Wang, and S. Gao, "Lift path planning for a nonholonomic crawler crane," *Automation in Construction*, vol. 44, pp. 12–24, 2014.
- [15] J. Escareño, S. Salazar, H. Romero, and R. Lozano, "Trajectory control of a quadrotor subject to 2D wind disturbances," *Journal of Intelligent & Robotic Systems*, vol. 70, no. 1–4, pp. 51–63, 2013.
- [16] M. Reyhanoglu, A. van der Schaft, N. H. McClamroch, and I. Kolmanovsky, "Nonlinear control of a class of underactuated systems," in *Proceedings of the 35th IEEE Conference on Decision and Control*, vol. 2, 1996, pp. 1682–1687.
- [17] Y. Fang, W. Dixon, D. Dawson, and E. Zergeroglu, "Nonlinear coupling control laws for an underactuated overhead crane system," *IEEE/ASME Transactions on Mechatronics*, vol. 8, no. 3, pp. 418–423, 2003.
- [18] G. J. Toussaint, T. Basar, and F. Bullo, "Tracking for nonlinear underactuated surface vessels with generalized forces," in *Proceedings of the IEEE International Conference on Control Applications*, 2000, pp. 355–360.
- [19] H. Bouadi, M. Bouchoucha, and M. Tadjine, "Sliding mode control based on backstepping approach for an UAV type-quadrotor," *World Academy of Science, Engineering and Technology*, vol. 26, no. 5, pp. 22–27, 2007.
- [20] —, "Modelling and stabilizing control laws design based on backstepping for an uav type-quadrotor," *IFAC Proceedings Volumes*, vol. 40, no. 15, pp. 245–250, 2007.
- [21] N. Chen, F. Song, G. Li, X. Sun, and C. Ai, "An adaptive sliding mode backstepping control for the mobile manipulator with nonholonomic constraints," *Communications in Nonlinear Science and Numerical Simulation*, vol. 18, no. 10, pp. 2885–2899, 2013.
- [22] K. Nakazono, K. Ohnishi, H. Kinjo, and T. Yamamoto, "Vibration control of load for rotary crane system using neural network with GA-based training," *Artificial Life and Robotics*, vol. 13, no. 1, pp. 98–101, 2008.
- [23] E. Shammas and M. de Oliveira, "An analytic motion planning solution for the snakeboard," in *Robotics: Science and Systems VII*. MIT Press, 2012, p. 297.
- [24] F. Morbidi and D. Prattichizzo, "Sliding mode formation tracking control of a tractor and trailer-car system," *Robotics: Science and Systems*, vol. 3, pp. 126–133, 2007.
- [25] B. Roy and H. Asada, "Closed loop control of a gravity-assisted underactuated snake robot with application to aircraft wing-box assembly," *Robotics: Science and Systems III*, vol. 25, no. 5, 2007.
- [26] A. Becker and T. Bretl, "Motion planning under bounded uncertainty using ensemble control," in *Robotics: Science and Systems*, 2010.
- [27] I. Kolmanovsky and N. H. McClamroch, "Developments in nonholonomic control problems," *IEEE Control Systems*, vol. 15, no. 6, pp. 20–36, 1995.
- [28] J. D. Boskovic, S.-M. Li, and R. K. Mehra, "Intelligent control of spacecraft in the presence of actuator failures," in *Proceedings of the 38th IEEE Conference on Decision and Control*, vol. 5, 1999, pp. 4472–4477.
- [29] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ: Prentice-Hall, 2002.
- [30] A. M. El-Nagar, M. El-Bardini, and N. M. EL-Rabaia, "Intelligent control for nonlinear inverted pendulum based on interval type-2 fuzzy PD controller," *Alexandria Engineering Journal*, vol. 53, no. 1, pp. 23–32, 2014.
- [31] S. Aseev and V. Veliov, "Needle variations in infinite-horizon optimal control," *Variational and Optimal Control Problems on Unbounded Domains*, vol. 619, pp. 1–17, 2014.
- [32] M. S. Shaikh and P. E. Caines, "On the hybrid optimal control problem: theory and algorithms," vol. 52, no. 9, pp. 1587–1603, 2007.
- [33] E. Todorov and W. Li, "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *Proceedings of the American Control Conference*, 2005, pp. 300–306.
- [34] R. W. Brockett *et al.*, "Asymptotic stability and feedback stabilization," *Differential geometric control theory*, vol. 27, no. 1, pp. 181–191, 1983.
- [35] M. Egerstedt, Y. Wardi, and H. Axelsson, "Transition-time optimization for switched-mode dynamical systems," vol. 51, no. 1, pp. 110–115, 2006.
- [36] A. R. Ansari and T. D. Murphey, "Sequential action control: closed-form optimal control for nonlinear and nonsmooth systems," vol. 32, no. 5, pp. 1196–1214, 2016.
- [37] T. M. Caldwell and T. D. Murphey, "Switching mode generation and optimal estimation with application to skid-steering," *Automatica*, vol. 47, no. 1, pp. 50–64, 2011.
- [38] B. Jakubczyk, "Introduction to geometric nonlinear control; controllability and Lie bracket," *Mathematical Control Theory*, vol. 1, no. 2, pp. 107–168, 2001.
- [39] P. Rashevsky, "About connecting two points of a completely nonholonomic space by admissible curve," *Uch. Zapiski Ped. Inst. Libknechta*, vol. 2, pp. 83–94, 1938.
- [40] W. Chow, "Über Systeme von linearen partiellen Differentialgleichungen erster Ordnung," *Mathematische Annalen*, vol. 117, pp. 98–105, 1940/1941. [Online]. Available: <http://eudml.org/doc/160043>
- [41] W. Murray, "Newton-type methods," *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- [42] R. B. Schnabel and E. Eskow, "A new modified Cholesky factorization," *SIAM Journal on Scientific and Statistical Computing*, vol. 11, no. 6, pp. 1136–1158, 1990.
- [43] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge, U.K.: Cambridge University Press, 2004.
- [44] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer, 2006, ch. 3.
- [45] G. Mamakoukas, M. A. MacIver, and T. D. Murphey, "Sequential action control for models of underactuated underwater vehicles in a planar ideal fluid," in *American Control Conference*, 2016, pp. 4500–4506.
- [46] T. Fan and T. D. Murphey, "Online feedback control for input-saturated robotic systems on Lie groups," in *Robotics: Science and Systems*, 2016.
- [47] D. Titterton and J. L. Weston, *Strapdown Inertial Navigation Technology*, 2nd ed. IET, 2004, vol. 17.
- [48] J. B. Kuipers, *Quaternions and Rotation Sequences*. Princeton, NJ: Princeton Univ. Press, 1999, vol. 66.
- [49] A. E. C. da Cunha, "Benchmark: Quadrotor attitude control," in *ARCH@ CPSWeek*, 2015, pp. 57–72.
- [50] C. M. Postlethwaite, T. M. Psemeneke, J. Selimkhanov, M. Silber, and M. A. MacIver, "Optimal movement in the prey strikes of weakly electric fish: a case study of the interplay of body plan and movement capability," *J. Roy. Soc. Interface*, vol. 6, no. 34, pp. 417–433, 2009.