Dynamic Inverse Models in Human-Cyber-Physical Systems

Ryan M. Robinson[a], Dexter R.R. Scobee[b], Samuel A. Burden[c], and S. Shankar Sastry[b]

[a]Army Research Laboratory, 2800 Powder Mill Rd., Adelphi, MD USA
[b]Department of Electrical Engineering and Computer Sciences at the University of California, Berkeley, CA USA
[c]Department of Electrical Engineering at the University of Washington, Seattle, WA USA

ABSTRACT

Human interaction with the physical world is increasingly mediated by automation. This interaction is characterized by dynamic coupling between robotic (i.e. cyber) and neuromechanical (i.e. human) decision-making agents. Guaranteeing performance of such human-cyber-physical systems will require predictive mathematical models of this dynamic coupling. Toward this end, we propose a rapprochement between robotics and neuromechanics premised on the existence of internal forward and inverse models in the human agent. We hypothesize that, in tele-robotic applications of interest, a human operator learns to invert automation dynamics, directly translating from desired task to required control input. By formulating the model inversion problem in the context of a tracking task for a nonlinear control system in control-affine form, we derive criteria for exponential tracking and show that the resulting dynamic inverse model generally renders a portion of the physical system state (i.e., the internal dynamics) unobservable from the human operator's perspective. Under stability conditions, we show that the human can achieve exponential tracking *without formulating an estimate of the system's state* so long as they possess an accurate model of the system's dynamics. These theoretical results are illustrated using a planar quadrotor example. We then demonstrate that the automation can intervene to improve performance of the tracking task by solving an optimal control problem. Performance is guaranteed to improve under the assumption that the human learns and inverts the dynamic model of the altered system. We conclude with a discussion of practical limitations that may hinder exact dynamic model inversion.

Keywords: human-cyber-physical system (HCPS), internal model, dynamic inverse model, mixed-initiative system, autonomous intervention

1. INTRODUCTION

Human-cyber-physical systems (HCPS) are intelligent networked systems with deeply-integrated human, cyber, and physical elements. In many cases, the human and cyber (autonomous) components jointly control the dynamic physical plant. The joint behavior that emerges from the tight coupling and reciprocal adaptation of these elements often cannot be described by naively combining models of their individual contributions. Although the complementary features of the human and cyber components are intended to improve efficiency and robustness, improper integration of these components may inadvertently degrade task performance and may lead to unintended consequences. For instance, enhanced safety features in motor vehicles such as automated lane-departure warning systems have been associated with *increased* accident rates.[1] Due to the critical function that many HCPS serve (e.g. in energy, healthcare, defense, transportation, emergency response), the safety and resilience of these systems is of importance. As the application space for HCPS grows, it is critical to develop design frameworks with provable performance and safety guarantees.

Existing frameworks for human-machine interaction[2–4] provide high-level guidelines for the design, analysis, and evaluation of HCPS; for instance, Parasuraman and Sheridan[3] developed a four-stage model of automation to

Send correspondence to R.M.R
R.M.R: E-mail: ryan.robinson14.ctr@mail.mil, Telephone: 1 301 394 2066
D.R.R.S.: E-mail: dscobee@eecs.berkeley.edu
S.A.B.: Email: sburden@uw.edu
S.S.S.: Email: sastry@eecs.berkeley.edu

describe the major functions that automation can contribute toward, and proposed evaluation criteria based on human performance consequences (e.g. workload, situation awareness, complacency). However, such frameworks do not include mathematical models of the HCPS. This precludes the use of control-theoretic techniques to analyze stability, reachability, and other quantitative performance metrics.

Predictive mathematical models that directly characterize the coupling of human and cyber-physical subsystems may provide a general means for deriving stability and performance guarantees; yet given the richness of human behavior, modeling the human contribution in a generalized task is seemingly intractable. While it may be sufficient to derive models specific to a particular system/task, such as quasi-linear models of piloting behavior,[5] we seek a mathematical framework that can be used to design and control nonlinear systems performing tasks more general than regulation. Toward this end, we are inspired by concepts in the field of neuromechanical motor control.

Human interaction with the physical world has been investigated for over a century in the field of neuromechanical motor control. One popular paradigm posits existence of internal "forward" and "inverse" models that reduce behavior to a predictable map between input and output. Internal models can be described by both their forward and inverse formulations. A forward model M is a hypothetical computational network within the sensorimotor system that predicts the output y from motor command input u.[6] Conversely, an inverse model M^{-1} specifies the input(s) u that produce a desired output y. Tracking tasks—for instance, trajectory tracking with the human hand—can be performed by directly inverting the dynamics of the human body to provide required motor commands.[7] Although motor control and motor learning studies overwhelmingly focus on sensorimotor control of the human body itself rather than control of a cyber-physical system, valuable insight into human behavior in the presence of automation can be gained by understanding the principles behind human motor control.

Experimental evidence supports the use of forward and inverse models for motor control.[7–11] Previous studies suggest that weighted combinations of motor primitives (or muscle synergies)[12–14] can be learned, stored as internal models, and later recalled in order to transform desired limb trajectories into motor commands.[15] Shadmehr and Mussa-Ivaldi[7] performed multi-joint arm motion experiments with and without a disturbing force field. Their experiments demonstrated (i) recovery of kinematics after adaptation to the force field, indicating there was a kinematic plan (desired trajectory) independent of dynamical conditions, (ii) opposite kinematics upon sudden removal of a force field (after prolonged exposure), suggesting that the dynamics related to the force field had been inverted, and (iii) aftereffects from adaptation in workspace regions where no exposure to the field had taken place. Other studies have directly shown an inverse model relationship being generated for stable force fields (in parallel with impedance control during learning and unstable situations).[16,17] There is physiological evidence that internal models are learned and stored in the cerebellum[18–20] and executed in coordination with motor cortex.[21]

We hypothesize that human interaction with complex automated systems is also characterized by internal forward and inverse models. If a human can learn internal models of "simple" external tasks such as limb movement and tool use,[22,23] it is likely that this capability extends to more complex dynamical systems.

Hypothesis: *humans learn and invert dynamic models to control cyber-physical systems.*

In the present work, we introduce a mathematical framework to describe the result of this *dynamic inverse modeling* process for a trajectory tracking task, as would be employed in robotic teleoperation. We assume that the human is able to exactly learn the input-output relationship of a dynamic automated plant and is able to construct a state-dependent *dynamic inverse model* of the plant to implement the desired behavior. Using standard techniques from control theory, we demonstrate that implementation of the dynamic inverse model is equivalent to feedback linearization (also known as input-output linearization) of a nonlinear dynamic plant.

While the dynamic inverse model generally renders the physical system state *unobservable* from the human operator's perspective, we show that the human operator can achieve exponential tracking *without formulating an estimate of the system's state* so long as he/she possesses an accurate model of the system's dynamics. Additionally, we propose a concept for autonomous *interventions*, whereby a set of parameters is optimized and autonomously implemented to minimize a cost function. Assuming that the human then learns to exactly invert the new dynamic model, performance as defined by the cost function is guaranteed to improve while maintaining trajectory tracking performance.

Section 2 describes a stabilizing feedback linearization strategy (mathematically representing dynamic model inversion) for nonlinear control systems in control-affine form, and derives conditions for exponential tracking using feedback. Section 3 highlights the ambiguities inherent in inverting generalized nonlinear static and dynamic models, justifying our focus on a particular class of nonlinear systems (those with finite *strict relative degree*). Section 4 describes how the human operator might implement a dynamic inverse model when a portion of the states are unobservable, as well as our concept for autonomous interventions. Lastly, Section 5 discusses how human characteristics (e.g. approximate modeling, disturbance rejection, motor delay, input saturation, and physiological state) may affect dynamic-inverse-model-based control.

2. CONTROL-THEORETIC FORMALISM FOR DYNAMIC INVERSE MODELS

The scientific concept of dynamic inverse modeling found in neuromechanical motor control studies can be described from the perspective of mathematical control theory. Here, we provide mathematical formulations for forward and inverse models applicable to a class of nonlinear control systems relevant for modeling a broad range of applications. We focus on the model inversion problem for a *tracking* task.

First, a nonlinear state-dependent forward model is defined, and a standard coordinate transformation (Eq. 12) is applied to the model, enabling the development of a generalized linearizing feedback law (analogous to a dynamic inverse model); we discuss the derivation of this transformation in a tutorial manner. We then describe additional conditions on the input and the physical system that ensure the time-varying feedback prescribed by the inverse model achieves exponential tracking. This section discusses the derivation and properties of the dynamic inverse model for single-input, single-output (SISO) systems. An extension to square multi-input, multi-output (MIMO) systems is straightforward along the lines of Chapter 9 in (Sastry, 1999).[24]

2.1 Transformation of Forward Model

Consider a forward model that can be represented as a SISO nonlinear system in control-affine form:

$$\begin{aligned} \dot{x} &= f(x) + g(x)u \\ y &= h(x)\,, \end{aligned} \tag{1}$$

where[*] $x \in \mathbb{R}^n$, $f, g \in C^r(\mathbb{R}^n, \mathbb{R}^n)$, and $h \in C^r(\mathbb{R}^n, \mathbb{R})$. The forward model of this system is a map $M : X \times U \to Y$ transforming a given input trajectory $u : \mathbb{R} \to \mathbb{R}$ to an output trajectory $y : \mathbb{R} \to \mathbb{R}$ by solving the ordinary differential equation governing the dynamics of the state variable x.

A change of coordinates, which simplifies the development of a generalized control law, is partially obtained by taking consecutive time derivatives of the output y. We preface details of the coordinate transformation with a differentiation of y and a definition of strict relative degree. Differentiating y from Eq. 1 with respect to time, one obtains

$$\begin{aligned} \dot{y} &= \frac{\partial h}{\partial x} f(x) + \frac{\partial h}{\partial x} g(x) u \\ &:= L_f h(x) + L_g h(x) u\,. \end{aligned} \tag{2}$$

$L_f h(x) : \mathbb{R}^n \to \mathbb{R}$ and $L_g h(x) : \mathbb{R}^n \to \mathbb{R}$ denote the *Lie derivatives* of h with respect to f and g. We can understand $L_f h(x)$ as a function giving the rate of change of h along the *flow* of the vector field f; in other words, it describes how the internal dynamics f manifest themselves in the time derivative of the output. Likewise, $L_g h(x)$ is the rate of change of h along the flow of the vector field g. If $L_g h(x)$ is bounded away from zero for all $x \in \mathbb{R}^n$, the linearizing feedback law given by

[*] $C^r(A, B)$ denotes the set of r-times continuously differentiable functions from domain A to codomain B.

$$u = \frac{1}{L_g h(x)}(-L_f h(x) + v) \tag{3}$$

yields the *first-order linear system* from the artificial input v to the output y:

$$\dot{y} = v \tag{4}$$

From Eq. 3, it is clear that there exist functions $a(x)$ and $b(x)$ such that the feedback law $u = a(x) + b(x)v$ linearizes the system near the equilibrium x_0.

In the instance that $L_g h(x) \equiv 0$, meaning that $\forall x \in U, L_g h(x) = 0$, it is necessary to differentiate again:

$$\begin{aligned}\ddot{y} &= \frac{\partial L_f h}{\partial x} f(x) + \frac{\partial L_f h}{\partial x} g(x) u \\ &:= L_f^2 h(x) + L_g L_f h(x) u\,.\end{aligned} \tag{5}$$

In Eq. 5 above, $L_f^2 h(x) \equiv L_f(L_f h)(x)$ and $L_g L_f h(x) \equiv L_g(L_f h(x))$. If $L_g L_f h(x)$ is bounded away from zero for all $x \in U$, the control law given by

$$u = \frac{1}{L_g L_f h(x)}(-L_f^2 h(x) + v) \tag{6}$$

yields the *second-order linear system* from input v to output y:

$$\ddot{y} = v \tag{7}$$

The differentiation procedure may terminate at some finite γ, which is defined as the *strict relative degree* of the nonlinear system. Formally, the SISO nonlinear system Eq. 1 is said to have *strict relative degree* γ at $x_0 \in U$ if:

$$\begin{aligned}&L_g L_f^i h(x) \equiv 0 \qquad \forall x \in U\;, i = 0, ..., \gamma - 2\,, \\ &L_g L_f^{\gamma-1} h(x_0) \neq 0\,.\end{aligned} \tag{8}$$

In this generalized case, the control law given by:

$$u = \frac{1}{L_g L_f^{\gamma-1} h(x)}(-L_f^{\gamma} h(x) + v) \tag{9}$$

yields the linear γth order system from input v to output y:

$$y^{(\gamma)} = v \tag{10}$$

With these details, we now define a change of coordinates following the development in (Sastry, 1999).[24] For an nth order system with strict relative degree γ, the new coordinates are given by:

$$\begin{aligned}\phi_1(x) &= h(x) \\ \phi_2(x) &= L_f h(x) \\ &\vdots \\ \phi_\gamma(x) &= L_f^{\gamma-1} h(x)\,.\end{aligned} \tag{11}$$

The $\phi_i(x)$ are the coordinates described by y and the first $\gamma-1$ time derivatives of y. It is a consequence of the definition of relative degree that the first $\gamma-1$ time derivatives of y do not depend on u. The above system of equations qualifies as a *partial* change of coordinates for the system (since $\gamma \leq n$). If Eq. 1 has *strict relative degree* $\gamma \in \mathbb{N}$ at $x_0 \in \mathbb{R}^n$,[24] then there exists (by Frobenius' Theorem) a local diffeomorphism $\Phi \in C^r(X, \Phi(X))$ of the form

$$\Phi(x) = \begin{bmatrix} h(x) \\ L_f h(x) \\ \vdots \\ L_f^{\gamma-1} h(x) \\ \zeta(x) \end{bmatrix} = \begin{bmatrix} \xi_1 \\ \xi_2 \\ \vdots \\ \xi_\gamma \\ \zeta \end{bmatrix} \tag{12}$$

where $\zeta \in C^r(X, \mathbb{R}^{n-\gamma})$ is a coordinate transformation for the remaining state dimensions not described by ξ, L_f^ℓ denotes the ℓth Lie derivative of h along f, and $X \subset \mathbb{R}^n$ is a neighborhood containing x_0. ζ represent the remaining $n-\gamma$ coordinates. This allows us to transform Eq. 1 via Φ:

$$\begin{bmatrix} \dot{\xi}_1 \\ \dot{\xi}_2 \\ \vdots \\ \dot{\xi}_{\gamma-1} \\ \dot{\xi}_\gamma \\ \dot{\zeta} \end{bmatrix} = \begin{bmatrix} \xi_2 \\ \xi_3 \\ \vdots \\ \xi_\gamma \\ b(\xi,\zeta) + a(\xi,\zeta)u \\ q(\xi,\zeta) \end{bmatrix} \tag{13}$$

where $b(\xi,\zeta) = L_f^\gamma h(\Phi^{-1}(\xi,\zeta))$, $a(\xi,\zeta) = L_g L_f^{\gamma-1} h(\Phi^{-1}(\xi,\zeta))$ and $q(\xi,\zeta)$ represents $L_f\zeta$ in (ξ,ζ) coordinates. The system description Eq. 13 is known as the *normal form* for the system described in Eq. 1. The re-parameterized system above remains a forward model of the nonlinear system, mapping the input u to output $y = \xi_1$. Note the lack of input terms in the differential equations for ζ; this implies the input cannot directly influence these dynamics.

$\dot{\zeta} = q(\xi,\zeta)$ are the system's internal (unobservable) dynamics associated with the feedback linearization. When the control input is such that the output y is maintained at zero (with $\zeta = \zeta_0$), $\dot{\zeta} = q(\xi_0,\zeta)$ are known as the *zero dynamics* of the linearized system. Eq. 1 is *locally exponentially minimum phase* at x_0 if $\mathrm{spec} D_\zeta q(\xi_0,\zeta_0) \in \mathbb{C}_-^o$ (the eigenvalues of the system linearized at (ξ_0,ζ_0) are in the left half of the complex plane); i.e. if ζ_0 is a locally exponentially stable equilibrium of the zero dynamics.

2.2 Trajectory Tracking via Dynamic Model Inversion

As shown in Eq. 10, proper choice of an input u cancels out the internal states and enables direct control of the γth time derivative of y. This has important implications for trajectory tracking tasks. Consider a SISO tracking task, in which the human operator generates inputs with the goal of following a desired trajectory. If the operator learns to invert the dynamics of the system as we have hypothesized, then the dynamic inverse model M^{-1}, which maps the desired output $y_d(t)$ to a desired input u, takes the form of a control law. In the case that the system has strict relative degree γ, then the linearizing feedback law

$$u = \frac{1}{a(\xi,\zeta)}(-b(\xi,\zeta) + v) \tag{14}$$

yields $\frac{d^\gamma}{dt^\gamma}\xi_1 = \xi_1^{(\gamma)} = v$. This control law effectively cancels out the nonlinear plant dynamics described by functions $a(\xi,\zeta)$ and $b(\xi,\zeta)$. Note that $y = \xi_1$ and $\frac{d^\gamma}{dt^\gamma}\xi_1 = v$, hence this feedback renders the states ζ unobservable from the output. Note there is ambiguity in how this control law can be applied (by a human operator) when the states ζ are unobservable; this will be addressed in Section 4.1.

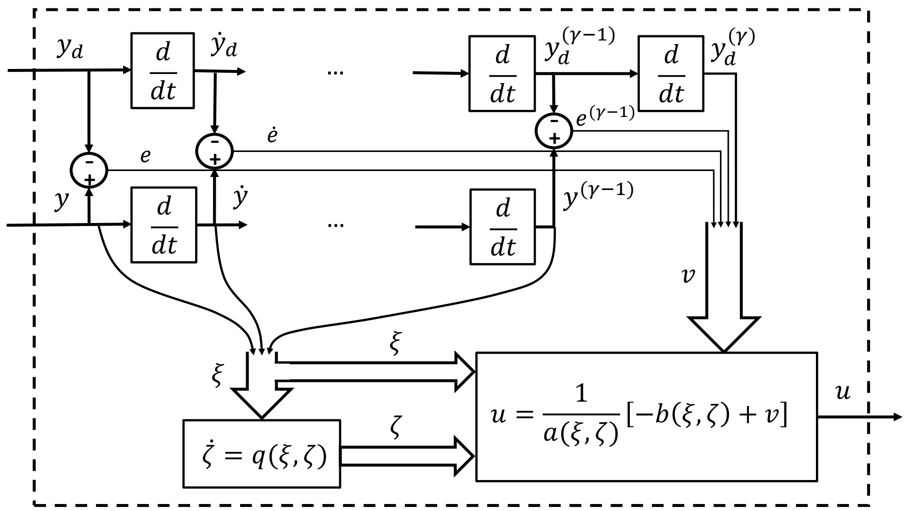

Figure 1: Dynamic-inverse-model-based control for exponential tracking, adapted from (Sastry, 1999).[24] The output y and its derivatives are transformed into coordinates ξ and combined with a model of the internal dynamics $\dot{\zeta} = q(\xi, \zeta)$ to obtain the states ζ. The input signal is a combination of the state-dependent inverse model and error feedback.

Given a desired output trajectory y_d, and assuming the initial state can be prepared precisely, the choice of $v = \frac{d^\gamma}{dt^\gamma} y_d$ yields $\frac{d^\gamma}{dt^\gamma}\xi_1 = \frac{d^\gamma}{dt^\gamma} y_d$ when Eq. 14 is combined with Eq. 13, and therefore $y(t) = \xi_1(t) = y_d(t)$.

To implement trajectory tracking for *arbitrary* initial conditions, one could define the tracking error $e(t, x) = y(x) - y_d(t)$, where $y(x) = h(x) = \xi_1$, and choose v as

$$v(t, x) = y_d^{(\gamma)} - \alpha_{\gamma-1} e^{(\gamma-1)}(x) - ... - \alpha_0 e\,, \tag{15}$$

The feedback law defined by Eqs. 14 and 15 results in exponential tracking for y and its first γ derivatives so long as (i) the polynomial $s^\gamma + \alpha_{\gamma-1}s^{\gamma-1} + ... + \alpha_1 s + \alpha_0$ is Hurwitz, and (ii) the unobserved states ζ remain close enough to ζ_0 to ensure $(\xi, \zeta) \in X$ for all time $t \geq t_0$ (c.f. Remark 1 to Theorem 9.14 in Sastry, 1999[24]). Thus we are led to consider conditions that will ensure ζ remains near ζ_0 under this feedback.

Contraction theory provides an elegant analysis for the behavior of our coupled system, as cascades of contractive systems are again contractive (Sontag, 2010,[25] Theorem 3). Applying the feedback in Eqs. 14 and 15, we obtain a cascaded system governing the closed-loop dynamics

$$\begin{bmatrix} \dot{\xi} \\ \dot{\zeta} \end{bmatrix} = \begin{bmatrix} A(\xi - \eta(t)) \\ q(\xi, \zeta) \end{bmatrix} =: F(t, \xi, \zeta) \tag{16}$$

where the vector $\eta(t) \in \mathbb{R}^\gamma$ is defined by

$$\eta(t) = \begin{bmatrix} y_d(t) & \dot{y_d}(t) & \dots & y_d^{(\gamma-1)}(t) \end{bmatrix}^T, \tag{17}$$

$F \in C^1(\mathbb{R} \times \mathbb{R}^\gamma \times \mathbb{R}^{n-\gamma}, \mathbb{R}^\gamma \times \mathbb{R}^{n-\gamma})$ is the closed-loop vector field, and the matrix $A \in \mathbb{R}^{\gamma\times\gamma}$ is in controllable canonical form,

$$A = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ -\alpha_0 & -\alpha_1 & -\alpha_2 & \dots & -\alpha_{\gamma-1} \end{bmatrix} \tag{18}$$

and hence spec$A \subset \mathbb{C}^o_-$ by construction. If we assume further that Eq. 1 is *locally exponentially minimum phase* at x_0, then spec$D_\zeta q(\xi_0, \zeta_0) \subset \mathbb{C}^o_-$. Computing the Jacobian derivative of F with respect to ξ and ζ we obtain

$$DF(\xi, \zeta) = \begin{bmatrix} A & 0 \\ D_\xi q(\xi, \zeta) & D_\zeta q(\xi, \zeta) \end{bmatrix}. \tag{19}$$

Since F is in C^1 there exists an open ball $B \in \mathbb{R}^n$ containing (ξ_0, ζ_0) for which:

- $\dot{\xi} = A(\xi - \eta(t))$ is *infinitesimally contracting*[25] (since spec $A \subset \mathbb{C}^o_-$);
- $\dot{\zeta} = q(\xi, \zeta)$ is *infinitesimally contracting when ξ is viewed as a parameter*[25] (since spec $D_\zeta q(\xi_0, \zeta_0) \subset \mathbb{C}^o_-$ and DF is continuous); and
- $D_\xi q(\xi_0, \zeta_0)$ is bounded (since DF is continuous).

This enables us to infer (by Sontag, 2010,[25] Theorem 3) that the cascaded system in Eq. 16 is infinitesimally contracting on B with some contraction rate $c > 0$. Therefore, we conclude (by Sontag, 2010,[25] Theorem 1) that for any reference $y_d \in Y$ where

$$Y = \{y \in C^\gamma(\mathbb{R}, \mathbb{R}) | \|y\|_{C^\gamma} < \infty, \forall t \in \mathbb{R} : (\eta(t), \zeta_0) \in C\}, \tag{20}$$

all trajectories initialized in B remain in B for all time and achieve exponential tracking,

$$||h \circ \phi^{t_0, (\xi,\zeta), u_d} - y_d||_{C^r(t)} \le e^{-c(t-t_0)}||(\xi, \zeta) - (\eta(0), \zeta_0)||, \ \forall\, t \ge t_0, \ (\xi, \zeta) \in C, \ y_d \in Y \tag{21}$$

By enforcing the strict relative degree property, there exists a unique input u_d that achieves a desired output y_d:

Property 1: Uniqueness. *As noted in (Sastry, 1999, §9.2.3),*[24] *the input that achieves exact tracking is unique.*

Moreover, since the dynamic inverse model is continuously differentiable, its behavior near the exact tracking trajectory is governed by its linearization. In other words, the local behavior of the dynamic inverse model Eq. 16 is determined by the tracking coefficients $\{\alpha_j\}_{j=0}^{\gamma-1}$ employed in Eq. 15 regardless of the global (possibly nonlinear) feedback law employed.

It remains unclear how a human operator might apply the proposed control law when the states ζ are unobservable. This will be addressed in Section 4.

2.3 Example: Planar Quadrotor Model

Consider a planar model of a quadrotor subject to gravitational acceleration (Figure 2). The quadrotor body of mass m and inertia I has position (z, x) and angle θ from the horizontal. There are two possible control inputs: a force $u_1 = F$ applied at the center of mass of the quadrotor body in the direction perpendicular to the quadrotor body (corresponding to equal positive force generated by each rotor), and a torque $u_2 = \tau$ about the center of mass of the quadrotor body (corresponding to opposite forces generated by the two rotors in the plane). The system dynamics are modeled with an equilibrium force mg applied perpendicular to the quadrotor body, so that $f(0) \equiv 0$.

2.3.1 Vertical Trajectory Tracking

First, we consider a SISO version of the system designed to follow a vertical trajectory $y_d(t) = z_d(t)$ while x and θ are stabilized about the origin. The sole input u is vertical force F. The system can be modeled in control-affine form:

$$\dot{x} = \begin{bmatrix} \dot{z} \\ \dot{x} \\ \dot{\theta} \\ \ddot{z} \\ \ddot{x} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} \dot{x_1} \\ \dot{x_2} \\ \dot{x_3} \\ \dot{x_4} \\ \dot{x_5} \\ \dot{x_6} \end{bmatrix} = \begin{bmatrix} x_4 \\ x_5 \\ x_6 \\ g(\cos x_3 - 1) \\ g\sin x_3 + k_{P,x}x_2 + k_{D,x}x_5 \\ k_{P,\theta}x_3 + k_{D,\theta}x_6 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \cos x_3/m \\ -\sin x_3/m \\ 0 \end{bmatrix} u \tag{22}$$
$$y = h(x) = x_1\,,$$

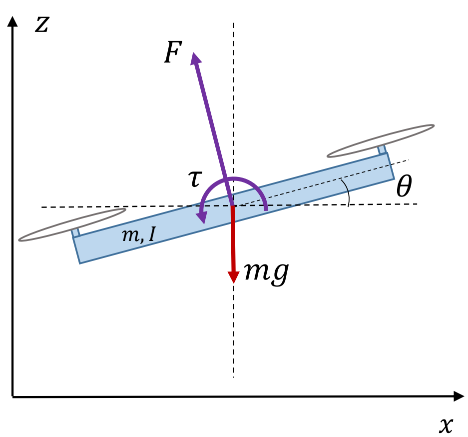

Figure 2: Simple model for quadrotor confined to the vertical plane in a gravitational field. Two rotors affixed symmetrically to the rigid chassis apply a net wrench proportional to applied voltages: the voltage sum is proportional to net force F; the voltage difference is proportional to net torque τ.

The system can also be described in normal form using the differentiation procedure in Section 3.

$$\begin{aligned} L_g h &= 0 \\ L_f h &= x_4 \end{aligned} \tag{23}$$

Since $L_g h \equiv 0$, we continue to differentiate:

$$\begin{aligned} L_g L_f h &= \cos x_3/m \\ L_f^2 h &= g(\cos x_3 - 1) \end{aligned} \tag{24}$$

$L_g L_f h \neq 0$, hence the system has relative degree $\gamma = 2$. In normal form, the system of equations is given by:

$$\begin{bmatrix} \dot{\xi}_1 \\ \dot{\xi}_2 \\ \dot{\zeta}_1 \\ \dot{\zeta}_2 \\ \dot{\zeta}_3 \\ \dot{\zeta}_4 \end{bmatrix} = \begin{bmatrix} \xi_2 \\ L_f^2 h(\zeta_2) + L_g L_f h(\zeta_2)u \\ \zeta_3 \\ \zeta_4 \\ g\sin\zeta_2 + k_{P,x}\zeta_1 + k_{D,x}\zeta_3 \\ k_{P,\theta}\zeta_2 + k_{D,\theta}\zeta_4 \end{bmatrix} \tag{25}$$

where we have used the the transformed coordinates $\xi_1 = x_1$, $\xi_2 = x_4$, $\zeta_1 = x_2$, $\zeta_2 = x_3$, $\zeta_3 = x_5$, and $\zeta_4 = x_6$. If the objective is to track a desired output y_d, the human operator can apply the feedback control law

$$u = \frac{1}{L_g L_f^{\gamma-1} h}\left(-L_f^{\gamma} h + v\right) = \frac{m}{\cos\zeta_2}\left(g(\cos\zeta_2 - 1) + v\right), \tag{26}$$

with v in a form similar to Eq. 15,

$$v = \ddot{y}_d - \alpha_1 \dot{e} - \alpha_0 e\,, \tag{27}$$

then we obtain the cascaded system:

$$\begin{bmatrix} \dot{\xi}_1 \\ \dot{\xi}_2 \\ \dot{\zeta}_1 \\ \dot{\zeta}_2 \\ \dot{\zeta}_3 \\ \dot{\zeta}_4 \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} 0 & 1 \\ -\alpha_0 & -\alpha_1 \end{bmatrix} \begin{bmatrix} \xi_1 - y_d \\ \xi_2 - \dot{y}_d \end{bmatrix} \\ \zeta_3 \\ \zeta_4 \\ g\sin\zeta_2 + k_{P,x}\zeta_1 + k_{D,x}\zeta_3 \\ k_{P,\theta}\zeta_2 + k_{D,\theta}\zeta_4 \end{bmatrix} \tag{28}$$

Note that the control law assumes knowledge of the state $\zeta_2 = x_3 = \theta$, which is unobservable. In physical terms, the quadrotor angle—which clearly affects the vertical dynamics—is rendered unobservable to the operator due to the inverting control input.

2.3.2 Horizontal Trajectory Tracking

Next, we consider a SISO version of the system designed to follow a *horizontal* trajectory $y_d(t) = x_d(t)$ while z and θ are stabilized about the origin. The sole input u is out-of-plane torque τ. In control-affine form, the system is given by:

$$\begin{aligned} \dot{x} = \begin{bmatrix} \dot{z} \\ \dot{x} \\ \dot{\theta} \\ \ddot{z} \\ \ddot{x} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} \dot{x_1} \\ \dot{x_2} \\ \dot{x_3} \\ \dot{x_4} \\ \dot{x_5} \\ \dot{x_6} \end{bmatrix} &= \begin{bmatrix} x_4 \\ x_5 \\ x_6 \\ g(\cos x_3 - 1) + k_{P,z}x_1 + k_{D,z}x_4 \\ g\sin x_3 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1/I \end{bmatrix} u \\ y = h(x) = x_2\,, & \end{aligned} \tag{29}$$

It is clear from the equations of motion that the dynamics are 4th order: the torque input directly affects angular acceleration, which must be integrated twice to affect horizontal acceleration; horizontal acceleration must then be integrated twice to affect horizontal position. As per the procedure in Section 3, we obtain:

$$\begin{aligned} L_g h &= 0 \\ L_f h &= x_5 \\ L_g L_f h &= 0 \\ L_f^2 h &= g\sin x_3 \\ L_g L_f^2 h &= 0 \\ L_f^3 h &= x_6 g\cos x_3 \\ L_g L_f^3 h &= \frac{g\cos x_3}{I} \\ L_f^4 h &= -x_6^2 g\sin x_3 \end{aligned} \tag{30}$$

$L_g L_f^3 h \not\equiv 0$, hence the system has relative degree $\gamma = 4$. In normal form, the system of equations is given by:

$$\begin{bmatrix} \dot{\xi}_1 \\ \dot{\xi}_2 \\ \dot{\xi}_3 \\ \dot{\xi}_4 \\ \dot{\zeta}_1 \\ \dot{\zeta}_2 \end{bmatrix} = \begin{bmatrix} \xi_2 \\ \xi_3 \\ \xi_4 \\ L_f^4 h(\xi_3, \xi_4) + L_g L_f^3 h(\xi_3) u \\ \zeta_2 \\ \sqrt{g^2 - \xi_3^2} - g + k_{P,z}\zeta_1 + k_{D,z}\zeta_2 \end{bmatrix} \tag{31}$$

where we have used the the transformed coordinates $\xi_1 = x_2$, $\xi_2 = x_5$, $\xi_3 = g \sin x_3$, $\xi_4 = x_6 g \cos x_3$, $\zeta_1 = x_1$, and $\zeta_2 = x_4$. If the human operator applies the control law

$$u = \frac{1}{L_g L_f^3 h}(-L_f^4 h + v) = \frac{I}{g \cos x_3}(x_6^2 g \sin x_3 + v) = \frac{I}{\sqrt{g^2 - \xi_3^2}} \left(\frac{\xi_4^2}{\sqrt{g^2 - \xi_3^2}} + v \right) \tag{32}$$

to track y_d, with

$$v = y_d^{(4)} - \alpha_3 \dddot{e} - \alpha_2 \ddot{e} - \alpha_1 \dot{e} - \alpha_0 e \,, \tag{33}$$

then the subsequent closed-loop cascaded system can be represented as:

$$\begin{bmatrix} \dot{\xi}_1 \\ \dot{\xi}_2 \\ \dot{\xi}_3 \\ \dot{\xi}_4 \\ \dot{\zeta}_1 \\ \dot{\zeta}_2 \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -\alpha_0 & -\alpha_1 & -\alpha_2 & -\alpha_3 \end{bmatrix} \begin{bmatrix} \xi_1 - y_d \\ \xi_2 - \dot{y}_d \\ \xi_3 - \ddot{y}_d \\ \xi_4 - \dddot{y}_d \end{bmatrix} \\ \zeta_2 \\ (\sqrt{g^2 - \xi_3^2} - g) + k_{P,z}\zeta_1 + k_{D,z}\zeta_2 \end{bmatrix} \tag{34}$$

In contrast with the control law for vertical trajectory tracking, the states ζ are not required for the application of Eqs. 32 and 33 (only ξ_3 and ξ_4, which can be determined from the output, are required). This is physically intuitive, as vertical position and velocity should have no effect on the horizontal dynamics of a quadrotor.

3. AMBIGUITIES IN STATIC AND DYNAMIC INVERSE MODELS

According to contemporary theories of forward and inverse modeling, humans invert the causal relationship between inputs and outputs acquired through a learning process[†]. At an abstract level, it is straightforward to describe how a sensorimotor system utilizes internal models: *forward models* predict sensory consequences of motor actions, while *inverse models* compute motor commands that generate desired sensory signals. However, the structure and properties of these models generally vary with task and environment. Even given a plausible forward model, there may not be an obvious candidate for the model's inverse. For instance, a *static* map $M : U \to Y$ (a map having no internal states) that is injective but not surjective has a well defined inverse $M^{-1} : M(U) \to U$ defined over the subset $M(U) \in Y$, but this "inverse" yields no prediction for any output $y \in (Y \setminus M(U))$. If the static map is surjective but not injective, it may be possible to define a (pseudo-)inverse that selects from the set of inputs that yield the given output, for instance based on optimizing a cost function.

Inverse models with *dynamic* internal states may contain additional ambiguities because the output is related to the input only indirectly, through an intermediate dynamical state. In the form proposed by Thoroughman and Shadmehr,[15] the inverse model of Lagrangian dynamics that govern limb movement may be interpreted as acausal, or whose output at time t depends partly on its input for some later time $t + n$. In contrast, a

[†]The question of *how* the forward and inverse models are learned, though a worthy subject of study in and of itself, does not affect our conclusions, hence we refer the interested reader to other studies on the subject.[26–28]

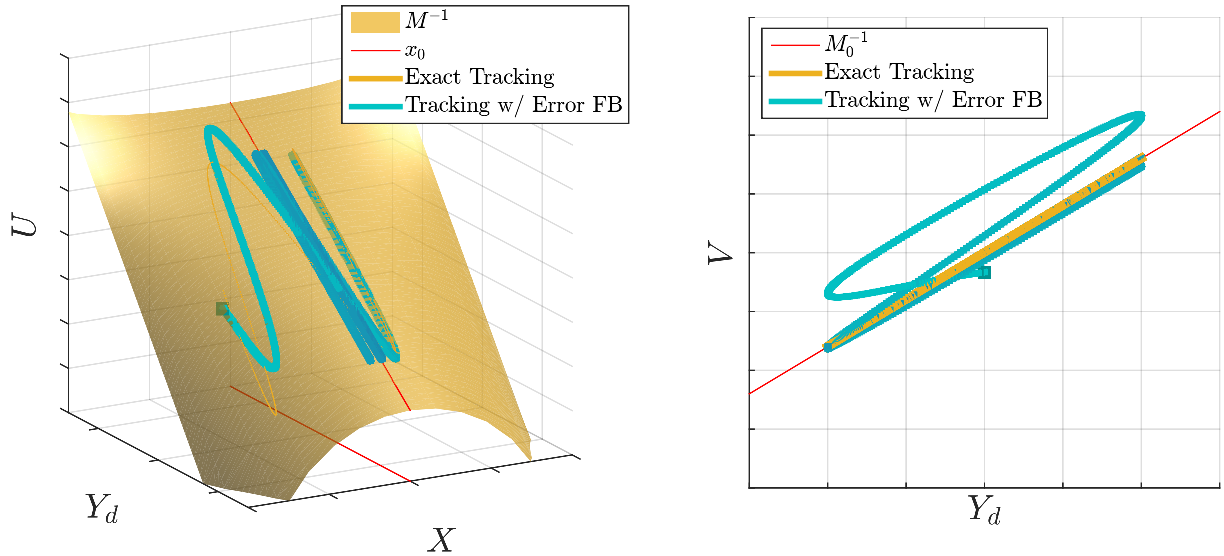

(a) (b)

Figure 3: (a) Illustration of a system with stable dynamic inverse model $M^{-1} : Y \to X \times U$. Suppose we must determine the appropriate input $u_d(t)$ to achieve a desired trajectory $y_d(t)$. Assuming no initial tracking error, the trajectory begins and remains on the surface, bringing the system to equilibrium x_0 (note that X may be multi-dimensional, but is represented in this figure by a single dimension). For arbitrary initial conditions, the dynamic inverse model is applied in unison with error feedback (Eq. 15) and will simultaneously converge to the surface and the equilibrium. (b) Implementation of the dynamic inverse model is equivalent to feedback linearization. Since feedback linearization cancels the nonlinear dynamics, the inverse model can be represented as a *static* map $M_0^{-1} : Y \to V$, where v represents the artificial input from Eq. 14. For arbitrary initial conditions, error feedback terms are added to v, causing it to deviate from the pure inverse.

state feedback inverse model proposed in Diedrichsen et al.[29] is dependent only on the current estimate of the state. The model is given by $M : X \times U \to Y$, while its inverse is $M^{-1} : Y \to X \times U$. Though not necessarily incompatible, these two postulated inverse model structures generally yield different predictions for the sensorimotor system's behavior.

By requiring that the nonlinear systems have *strict relative degree*, a property described in the previous section, the ambiguity due to acausality and redundancy is eliminated. Systems with finite strict relative degree are feedback linearizable. Under stability conditions about an equilibrium x_0, the dynamic map of the feedback linearized system (originally dependent on the input and state) becomes solely dependent on the artificial input v so that $M_0 : V \to Y$ (see Figure 3), where $M_0(\cdot) = M(x_0, \cdot)$. As a consequence, the inverse model $M_0^{-1} : Y \to V$ is valid in a contraction region B regardless of the state. The implications of this phenomenon for human teleoperation are discussed in the next section.

4. DYNAMIC INVERSE MODELS IN HUMAN-CYBER-PHYSICAL SYSTEMS

We hypothesize that humans are able to learn forward and inverse models of external systems, and are able to implement inverse models in the form of a feedback control law. In this section, we describe how a human can achieve exponential tracking for the class of nonlinear systems from Section 2 without observing the states ζ. The details of this result have important implications for the design of HCPS. The dynamic inverse modeling concept also leads us to propose a methodology for autonomous interventions in HCPS trajectory-tracking tasks (Section 4.2).

4.1 Unobservability and the Virtual Internal Model

Whenever the relative degree of the system is strictly smaller than the dimension of the state, the inverse model *always* renders some states unobservable:

Property 2: *As noted in (Sastry, 1999, §9.2.2),*[24] *the feedback in Eq. 14 renders the states ζ unobservable.*

Since the unobservability of the internal states precludes direct implementation of the feedback law, we describe the use of a *virtual system state* to achieve exponential tracking with a stable model pair.

Theorem 1: (exponential tracking with a stable model pair). *If the forward model Eq. 1 has strict relative degree $\gamma \in \mathbb{N}$ and is exponentially stable and exponentially minimum phase at an equilibrium $f(x_0) = 0$ then there exists an open ball $B \subset \mathbb{R}^n$ containing x_0 and a rate $c > 0$ such that the time-varying feedforward input obtained by applying Eqs. 14 and 15 along the trajectory initalized from (ξ_0, ζ_0) achieves exponential tracking (Eq. 21) at rate c for any desired output $y_d \in Y$ and its first γ derivatives.*

The critical detail here is that the input applied by the human is not obtained from a feedback law utilizing the physical system state (since the states ζ are unobservable); it is instead obtained by applying the inverse model to *virtual* system states $\hat{\zeta}$ initialized at their equilibrium values ζ_0. The control law

$$u = \frac{1}{a(\xi, \hat{\zeta})}(-b(\xi, \hat{\zeta}) + v), \tag{35}$$

generates the dynamics

$$y^{(\gamma)} = b(\xi, \zeta) + \frac{a(\xi, \zeta)}{a(\xi, \hat{\zeta})}(-b(\xi, \hat{\zeta}) + v), \tag{36}$$

which clearly reduces to $y^{(\gamma)} = v$ when $\hat{\zeta} = \zeta$.

Stability (specifically, infinitesimal contractivity) of the forward and inverse model pair ensure that the *virtual* and *physical* system states converge to the same trajectory, and this feedforward input achieves exponential tracking. Figure 4 depicts an example system with a dynamic inverse model $M^{-1}(\zeta)$ and its corresponding virtual model $M^{-1}(\hat{\zeta})$ constructed with equilibrium values ζ_0. Applying a feedback signal prescribed by the virtual-model-based control law yields a different output than exact state feedback; however, because the forward-inverse model pair is (exponentially) stable and the unobservable states are (exponentially) stable, the trajectories will exponentially converge to each other (and to y_d).

Given that the output is equal to z in the example of a quadrotor following a vertical trajectory (Section 2.3.1), the human does not have access to the states $\theta = \zeta_2$, $\dot{\theta} = \zeta_4$, $x = \zeta_1$, and $\dot{x} = \zeta_3$. If we choose the equilibrium conditions $(\theta, \dot{\theta}, x, \dot{x}) = (0, 0, 0, 0)$ to initialize a virtual model of the dynamics, then the control law is given by

$$u = \frac{m}{\cos \hat{\zeta}_2}\Big(g(\cos \hat{\zeta}_2 - 1) + v\Big) \tag{37}$$

The desired trajectory is $y_d(t) = 2\sin(2t)$. The states and output of the system with virtual-model-based feedback will initially deviate from that of the system with exact feedback. Simulations of the quadrotor response to the virtual-model-based control law Eq. 35, both from the human perspective (no knowledge of θ or x deviation) and in (simulated) reality, are illustrated in Figure 5a. Like the virtual model, the actual system converges to the desired trajectory. Simulations of the quadrotor comparing exact feedback and feedback using the virtual model are shown in Figure 5b, illustrating the (minor) difference between output trajectories.

Simulations were performed for various initial conditions with the virtual-model-based control law Eq. 35. As illustrated in Figure 6a, the trajectories can be seen to converge to one another (and the desired trajectory) exponentially. Comparison of the output error between initial conditions using exact or virtual model control in Figure 6b confirms that the trajectories converge to one another (exponential tracking) and there is zero

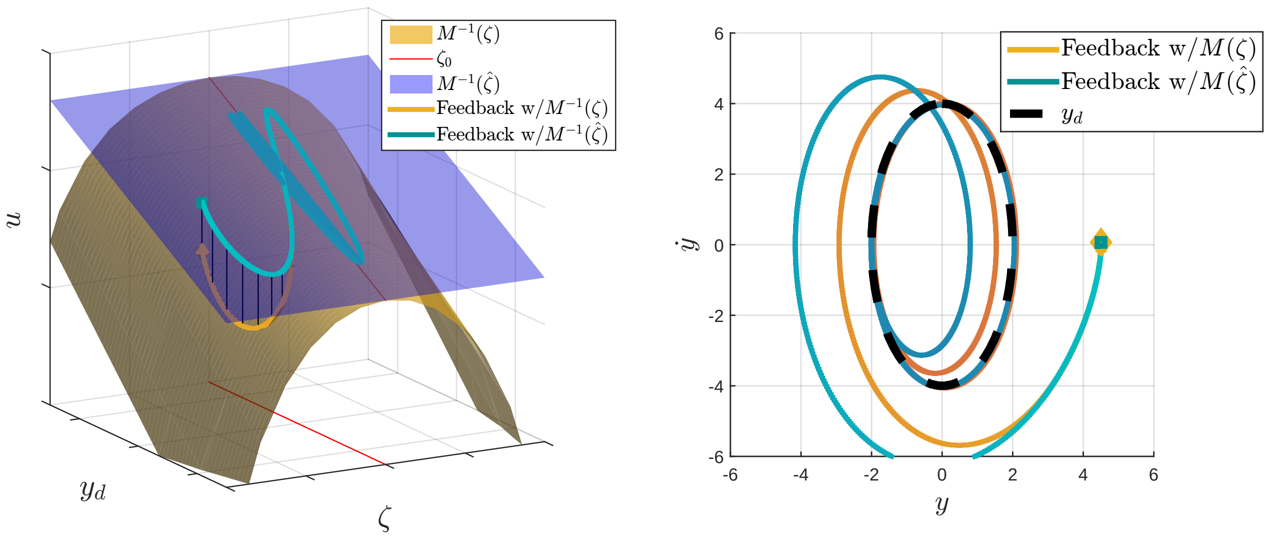

(a) (b)

Figure 4: (a) Illustration of a system with dynamic inverse model $M^{-1}(\zeta)$ and its corresponding virtual model $M^{-1}(\hat{\zeta})$ constructed with equilibrium values ζ_0. Suppose we must determine the appropriate input $u_d(t)$ to achieve a desired trajectory $y_d(t)$. If the state is known completely, the operator may execute perfect feedback linearization (Eq. 14, yellow-orange). However, if one or more states are unobservable, then the operator may instead implement the virtual-model-based control law (Eq. 35, blue) as if the unobservable states were at their equilibrium values. (b) The trajectory resulting from virtual-model-based control will generally differ from that produced by exact state feedback. Nevertheless, if the forward-inverse model pair is (exponentially) stable, and the unobservable states are (exponentially) stable, then the control law will yield (exponential) convergence to y_d (Theorem 1).

steady-state error. Depending on initial conditions, exact feedback produced less (red and green) or more (blue) overshoot.

We emphasize that, for the human operator to achieve exponential tracking in the presence of (unobservable) zero dynamics via inverse modeling, it is necessary for the physical plant to have an exponentially stable forward/inverse model pair. Though few physical systems are intrinsically stable about desirable operating points, many physical plants can be rendered stable via automated feedback control. Under the assumption that the human operator can learn to invert the plant's dynamics under stabilizing feedback, the preceding theory would apply to a system stabilized through low-level automation. Automated feedback, however, may be useful for goals beyond stabilization, including performance optimization about a desired trajectory, as described in the next section.

4.2 Autonomous Interventions

Due to the complexity and uncertain nature of many robotic tasks, high-level planning and/or shared control from a human (tele-)operator is often required for goal completion. While the human concentrates on planning and executing trajectories, an autonomous component can adapt the system dynamics in order to optimize other aspects of performance. However, in general, it is often difficult to predict the impact that alteration to plant dynamics may have on the system as a whole. For instance, the 2009 Air France disaster was caused by a poor response to autopilot shutdown—the pilots did not have an appropriate model of the high-altitude dynamics to recognize that their manual commands were stalling the plane.[30] Similarly, in the case of the 1992 YF-22 accident, the pilot transitioned from gear-down to gear-up at low altitude without realizing that this automatically increased pitch rate gains (intended to improve high-altitude performance), causing immediate pilot-induced oscillations (also known as aircraft-pilot coupling[31]) and a subsequent gear-up hard landing.[32]

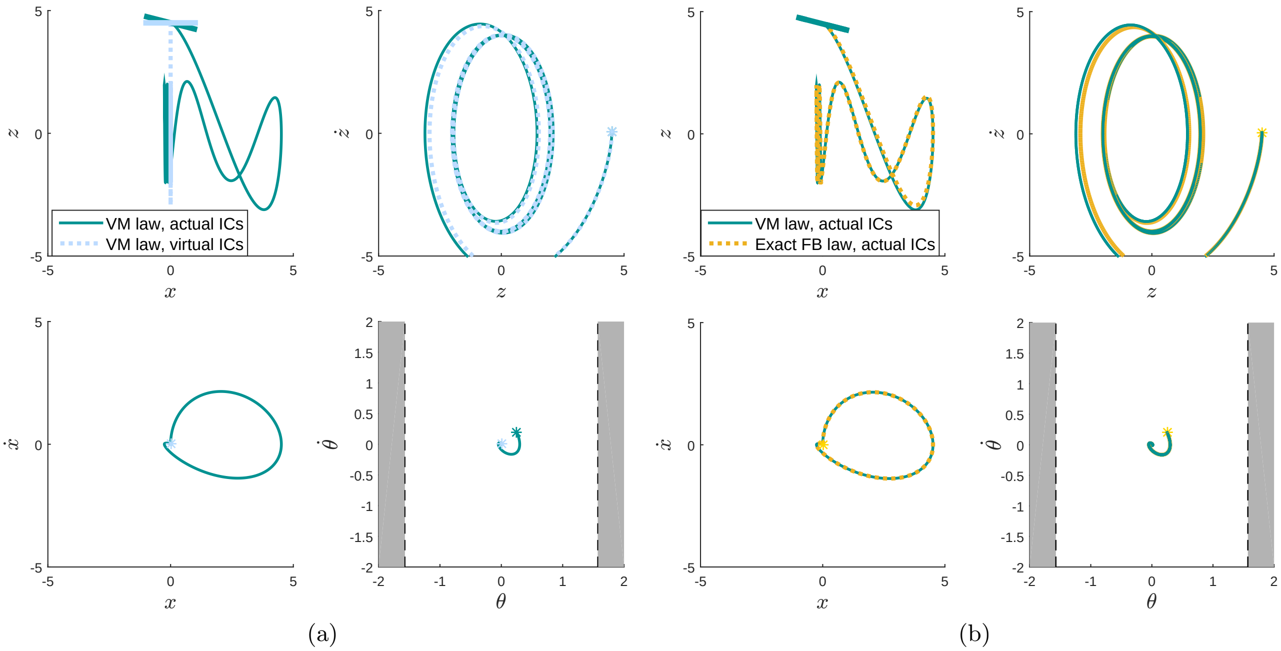

Figure 5: Simulations of a quadrotor performing vertical trajectory tracking. Initial conditions are $\{x_0, \dot{x}_0, z_0, \dot{z}_0, \theta_0, \dot{\theta}_0\} = \{0, 0, 4.5, 0.05, 0.25, 0.2\}$; gain values are $\{K_{P,x}, K_{D,x}, K_{P,\theta}, K_{D,\theta}\} = \{0.25, 0.7, 1, 1\}$. (a) The dotted curve represents virtual-model-based control if $\{x, \dot{x}, \theta, \dot{\theta}\} \equiv \{0, 0, 0, 0\}$ as the operator believes, while the solid curve is the actual result of the operator's virtual-model-based control law. Grayed-out areas represent $\theta < -\pi/2$ and $\theta > \pi/2$, from which convergence does not occur. (b) The dotted curve represents inverse-model-based control under complete state knowledge, while the solid curve represents control using the virtual model. Note that the $\{x, \dot{x}\}$ and $\{\theta, \dot{\theta}\}$ trajectories are the same for both control laws. This occurs because the unobservable dynamics in this example $\dot{\zeta} = q(\zeta)$ are independent of ξ and therefore completely unaffected by u; for general systems of the form Eq 13, the trajectories may differ.

We propose a specific methodology for autonomous adaptation based on the properties of the dynamic inverse model formulation. By applying interventions in a manner that does not change the basic underlying structure of the plant (e.g. relative degree, stability properties), we expect the human operator can learn to invert the augmented system's dynamics, possibly by adapting previously-acquired motor skills as was demonstrated by Shadmehr and Mussa-Ivaldi for reaching tasks in the presence of force fields.[7] In that study, learning to reach in a force field was mathematically represented as the sum of an original controller (modeling the inertial, Coriolis, centripetal, viscous, stiffness, and friction terms) used to reach in a null field and an environmental adaptation term to cancel the programmed force field. Through repeated trials, monotonic convergence to the original desired trajectory was observed. We envision analogous scenarios where the dynamics of a task are intentionally altered by the automated/cyber element (in a predictable manner), and the human learns the new dynamics to recover the desired trajectory.

We propose autonomous interventions that adapt the dynamics of the system through a set of parameters β to optimize secondary performance criteria (such as stabilizing a payload, minimizing fuel consumption, or extremizing other statistics) along the specified trajectory. It is expected that the human can learn and implement the new inverse model with the chosen β^* and regain the original tracking performance.

Given $y^* : [t_0, t_1] \to \mathbb{R}_1$ and $\beta \in B$, let

$$\beta^* = \operatorname{argmin} \left\{ J^M_{y^*}(\beta) \mid y^* \in Y \right\} \tag{38}$$

where $J^M_{y^*}$ represents the cost function for desired trajectory y^* using model M.

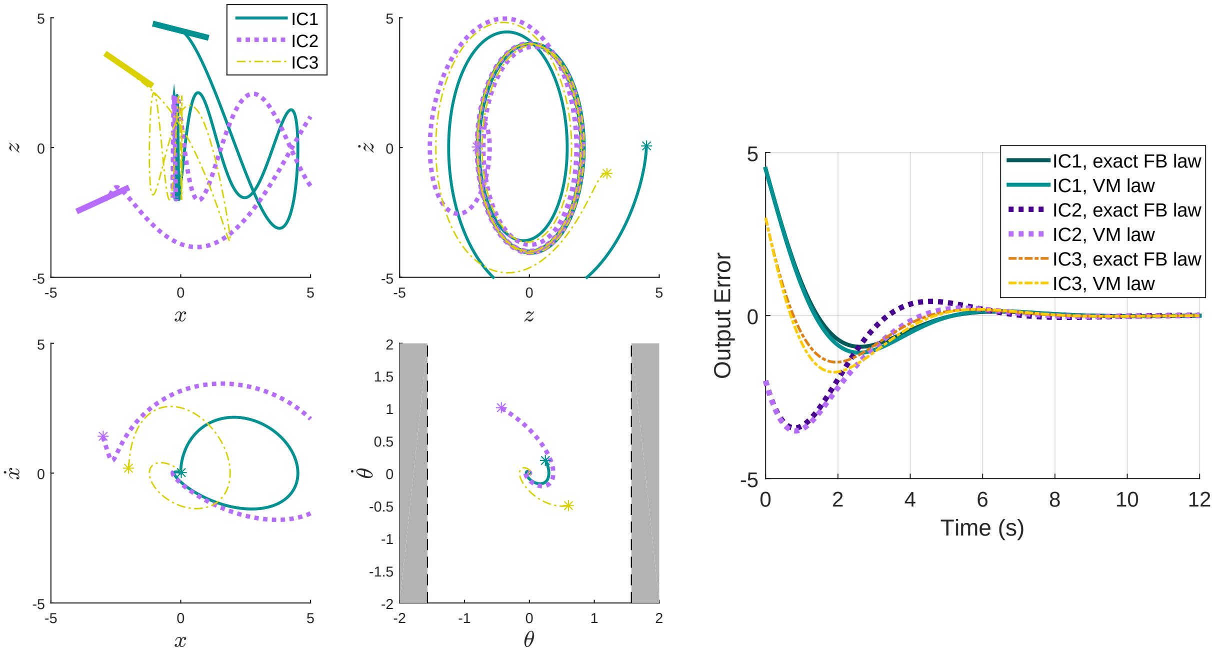

(a) (b)

Figure 6: Simulations of a quadrotor performing vertical trajectory tracking using a virtual-model-based control law, starting from three distinct initial conditions (IC1: $\{x_0, \dot{x}_0, z_0, \dot{z}_0, \theta_0, \dot{\theta}_0\} = \{0, 0, 4.5, 0.05, 0.25, 0.2\}$; IC2: $\{-3, 1.4, -2, 0, -0.43, 1\}$; IC3: $\{-2, 0.2, 3, -1, 0.6, -0.5\}$). Gain values are: $\{K_{P,x}, K_{D,x}, K_{P,\theta}, K_{D,\theta}\} = \{0.25, 0.7, 1, 1\}$. (a) Evolution of the states and output $y = z$; (b) Error in trajectory-following with an exact feedback law (darker) and with a virtual-model-based feedback law 35 (lighter).

The β parameters remain constant throughout the trajectory (until the next intervention), and do not affect the relative degree of the system, i.e.:

$$\begin{bmatrix} \dot{x} \\ \dot{\beta} \end{bmatrix} = \begin{bmatrix} f(x;\beta) + g(x;\beta)u \\ 0 \end{bmatrix} \iff \begin{bmatrix} \xi_1^{(\gamma)} \\ \dot{\zeta} \\ \dot{\beta} \end{bmatrix} = \begin{bmatrix} v \\ q(\xi, \zeta; \beta) \\ 0 \end{bmatrix} \tag{39}$$

We summarize the autonomous intervention procedure as follows:

1. Human learns to invert an original model of the system, completes a statistically significant number of orbits along a (operator-chosen) desired trajectory.

2. Autonomy adjusts parameters β to optimize a cost function along the trajectory; consequently, system dynamics are altered.

3. Human learns a new inverse model of the altered system. After learning, tracking behavior is equivalent to that of the previous system under our standing hypothesis, but performance as defined by the cost function improves.

4. If the environment and/or trajectory changes so that the parameters β are no longer near-optimal, repeat Steps 1-3.

Since not all orbits will be exactly alike (e.g., due to disturbances), it is not obvious how to select the trajectory y_β upon which to optimize. Based on the task, it may be sensible to optimize over the average trajectory:

Given observed trajectories $\{y_j^*\}_{j=1}^N$,

$$y_\beta = \text{argmin}\left\{\frac{1}{N}\sum_{j=1}^{N}||y_\beta - y_j^*||\right\}, \tag{40}$$

or, it may instead be prudent to minimize the maximum cost (worst case scenario) or the expected cost over all sampled trajectories:

$$y_\beta = \text{argmin}\left\{\max||y_\beta - y_j^*||\right\}, \tag{41}$$

$$y_\beta = \text{argmin}\left\{E||y_\beta - y_j^*||\right\}. \tag{42}$$

4.3 Example: Quadrotor Teleoperation

Consider a system composed of a semi-autonomous quadrotor with a slung payload connected by a spring, as well as a trained human teleoperator, tasked with building surveillance. We assume that a human operator is required to define the desired trajectory and to apply control inputs via a remote interface. We further assume that the human has learned an inverse model of the dynamic system with current parameters β. The task is performed repeatedly, generating a sample of periodic orbits. With adequate information on the desired trajectory y_d and coefficients α describing the error dynamics of these orbits, the autonomous/cyber component performs an optimization procedure to select parameters β which minimize a particular cost function. In this example, we choose the spring stiffness as a parameter, $\beta_1 = k$. The appropriate cost function may contain a payload jerk cost and a control effort cost, i.e.

$$J_{y^*}^M(k) = \int_{t_0}^{t_0+T}\left[c_1(\dddot{x}_P(t;k)^2 + \dddot{y}_P(t;k)^2) + c_2 u(t;k)^2\right] dt \tag{43}$$

When k is adjusted to the optimal setting k^*, the quadrotor dynamics have changed, and the human must adjust his/her own internal model through learning. Assuming the human can learn and implement the new dynamic inverse model, the trajectory tracking behavior will mirror the original behavior. Furthermore, performance as defined by the cost function is guaranteed to have improved. In future work, we intend to test predictions from this theory and protocol experimentally.

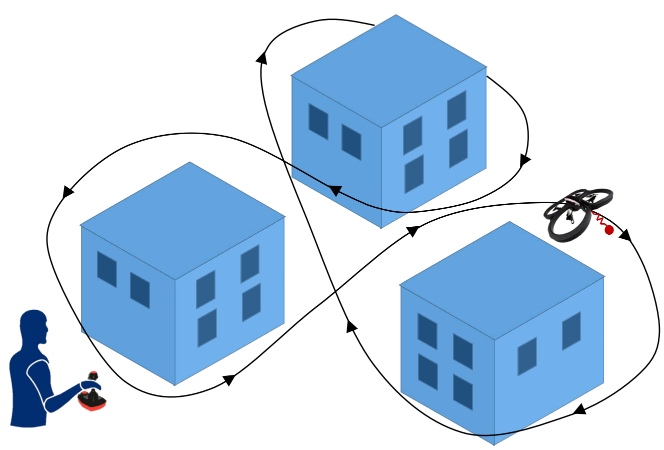

Figure 7: Teleoperated tracking of periodic orbit with a quadrotor and a slung payload.

5. DISCUSSION: LIMITATIONS AND FUTURE WORK

There are several practical reasons why a human operator may not *exactly* learn and invert the dynamics of a nonlinear plant, even given infinite learning time:

1. Internal/external noise during learning may prevent the exact scaling of combinations of motor primitives, only achieving an approximate model. For approximate models, it is important to examine the robustness of the linearization to plant uncertainties to determine the sensitivity of the system to the inexact cancellation of nonlinearities.[33]

2. The dynamics of a cyber-physical system often depend on a changing, unpredictable environment; for instance, an aerial vehicle's operation can depend on wind disturbances. Internal disturbances affecting operator input (e.g. neuromotor noise, muscle spasms) may also prevent exact model inversion. The effects of these disturbances usually cannot be learned and must be compensated online. Although the coefficients α describing the human tracking behavior are not required for the optimization problem in Section 4.2, it may be helpful to estimate these parameters to predict sensitivity to disturbances and transient behavior. These α can be estimated from experimental data prior to the intervention (using system identification techniques, for example). It has been experimentally demonstrated in a wide range of tasks that in response to disturbances, humans weigh the costs of motor variability against the consequences to task completion (the so-called minimum intervention principle[34–36]), and the resulting behavior resembles the function of an optimal feedback controller.[34,37] Optimal feedback control can be incorporated into our current inverse modeling paradigm.

3. Time delay is an inherent part of the human sensorimotor loop that is not currently accounted for in our formulation, but may significantly degrade performance if the system dynamics and/or task require a high bandwidth response.[5] Automation may be employed to stabilize high-frequency dynamics as described in Section 4.1.

4. Task constraints, interface constraints, and human musculoskeletal kinematic/dynamic constraints[38] may saturate the human input, prohibiting exact inversion.

5. Human performance depends heavily on mental and physiological state. Past studies[39–43]suggest that factors such as attentiveness, stress, and fatigue will influence the recall and/or implementation of previously-learned internal models.

Lastly, autonomous interventions are only practical if learning takes place over a reasonably short duration. In many cases, minor changes in system parameters may require only minor adjustments in the inverse model which can be compensated quickly. Future work will investigate methods for estimating adaptation time, which could be applied as an added cost in the intervention.

6. CONCLUSION

This paper presented a control-theoretic framework for human teleoperation of semi-autonomous dynamic nonlinear systems based on the concept of dynamic inverse models in neuromechanical motor control. For a class of nonlinear systems (including Lagrangian systems), exponential tracking may be achieved by the human operator via inversion of an internal model of the system, combined with output error feedback. While implementation of the dynamic inverse model generally renders some system states unobservable, we show that the operator can implement a virtual model assuming equilibrium values for the unobservable states of the system dynamics, since the virtual and physical states exponentially converge if the forward model is both exponentially stable and exponentially minimum phase.

The proposed inverse modeling framework also led to the formulation of a concept for autonomous interventions, in which the human defines a desired trajectory and a set of parameters is adjusted to optimize desired performance criteria along that trajectory. If, as we hypothesize, the human learns an inverse model of the augmented system, then we guarantee that the trajectory learned before autonomous intervention will again be tracked, hence performance as defined by the cost function will improve.

In future work, we will focus on experimental validation of the dynamic inverse modeling framework proposed herein, as well as extensions for robustness with respect to approximate modeling, noise, disturbances, time delay, and human variability.

ACKNOWLEDGMENTS

This effort was partially funded through the Office of the Secretary of Defense Autonomy Research Pilot Initiative (ARPI), MIPR DWAM31168. This work is also supported by the Office of Naval Research under the Embedded Humans MURI (N00014-13-1-0341).

REFERENCES

[1] Moore, M. and Zuby, D., "Collision avoidance features: initial results," in [*23rd International Conference on the Enhanced Safety of Vehicles*], (2013).

[2] Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., Van Riemsdijk, M. B., and Sierhuis, M., "Coactive design: Designing support for interdependence in joint activity," *Journal of Human-Robot Interaction, 3 (1), 2014* (2014).

[3] Parasuraman, R., Sheridan, T. B., and Wickens, C. D., "A model for types and levels of human interaction with automation," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on* **30**(3), 286–297 (2000).

[4] Cummings, M. L. and Clare, A., "Holistic modelling for human-autonomous system interaction," *Theoretical Issues in Ergonomics Science* **16**(3), 214–231 (2015).

[5] McRuer, D. T. and Krendel, E. S., "Mathematical models of human pilot behavior," tech. rep., DTIC Document (1974).

[6] Bhushan, N. and Shadmehr, R., "Computational nature of human adaptive control during learning of reaching movements in force fields," *Biological Cybernetics* **81**(1), 39–60 (1999).

[7] Shadmehr, R. and Mussa-Ivaldi, F. A., "Adaptive representation of dynamics during learning of a motor task," *The Journal of Neuroscience* **14**(5), 3208–3224 (1994).

[8] Wolpert, D. M. and Kawato, M., "Multiple paired forward and inverse models for motor control," *Neural Networks* **11**(7), 1317–1329 (1998).

[9] Gandolfo, F., Mussa-Ivaldi, F., and Bizzi, E., "Motor learning by field approximation," *Proceedings of the National Academy of Sciences* **93**(9), 3843–3846 (1996).

[10] Goodbody, S. J. and Wolpert, D. M., "Temporal and amplitude generalization in motor learning," *Journal of Neurophysiology* **79**(4), 1825–1838 (1998).

[11] Conditt, M. A., Gandolfo, F., and Mussa-Ivaldi, F. A., "The motor system does not learn the dynamics of the arm by rote memorization of past experience," *Journal of Neurophysiology* **78**(1), 554–560 (1997).

[12] Ting, L. H. and Macpherson, J. M., "A limited set of muscle synergies for force control during a postural task," *Journal of Neurophysiology* **93**(1), 609–613 (2005).

[13] d'Avella, A. and Bizzi, E., "Shared and specific muscle synergies in natural motor behaviors," *Proceedings of the National Academy of Sciences of the United States of America* **102**(8), 3076–3081 (2005).

[14] Ting, L. H. and McKay, J. L., "Neuromechanics of muscle synergies for posture and movement," *Current Opinion in Neurobiology* **17**(6), 622–628 (2007).

[15] Thoroughman, K. A. and Shadmehr, R., "Learning of action through adaptive combination of motor primitives," *Nature* **407**(6805), 742–747 (2000).

[16] Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., and Kawato, M., "The central nervous system stabilizes unstable dynamics by learning optimal impedance," *Nature* **414**(6862), 446–449 (2001).

[17] Franklin, D. W., Osu, R., Burdet, E., Kawato, M., and Milner, T. E., "Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model," *Journal of Neurophysiology* **90**(5), 3270–3282 (2003).

[18] Ito, M., "Control of mental activities by internal models in the cerebellum," *Nature Reviews Neuroscience* **9**(4), 304–313 (2008).

[19] Bastian, A. J., "Learning to predict the future: the cerebellum adapts feedforward movement control," *Current Opinion in Neurobiology* **16**(6), 645–649 (2006).

[20] Wolpert, D. M., Miall, R. C., and Kawato, M., "Internal models in the cerebellum," *Trends in Cognitive Sciences* **2**(9), 338–347 (1998).

[21] Gribble, P. L. and Scott, S. H., "Overlap of internal models in motor cortex for mechanical loads during reaching," *Nature* **417**(6892), 938–941 (2002).

[22] Imamizu, H., Miyauchi, S., Tamada, T., Sasaki, Y., Takino, R., PuÈtz, B., Yoshioka, T., and Kawato, M., "Human cerebellar activity reflecting an acquired internal model of a new tool," *Nature* **403**(6766), 192–195 (2000).

[23] Johnson-Frey, S. H., "The neural bases of complex tool use in humans," *Trends in Cognitive Sciences* **8**(2), 71–78 (2004).

[24] Sastry, S., [*Nonlinear systems: analysis, stability, and control*], vol. 10, Springer Science & Business Media (1999).

[25] Sontag, E. D., "Contractive systems with inputs," in [*Perspectives in Mathematical System Theory, Control, and Signal Processing*], 217–228, Springer (2010).

[26] Wolpert, D. M., Diedrichsen, J., and Flanagan, J. R., "Principles of sensorimotor learning," *Nature Reviews Neuroscience* **12**(12), 739–751 (2011).

[27] Smith, M. A., Ghazizadeh, A., and Shadmehr, R., "Interacting adaptive processes with different timescales underlie short-term motor learning," *PLoS Biol* **4**(6), e179 (2006).

[28] Donchin, O., Francis, J. T., and Shadmehr, R., "Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control," *The Journal of Neuroscience* **23**(27), 9032–9045 (2003).

[29] Diedrichsen, J., Shadmehr, R., and Ivry, R. B., "The coordination of movement: optimal feedback control and beyond," *Trends in Cognitive Sciences* **14**(1), 31–39 (2010).

[30] Wise, J.

[31] Pavel, M. D., Jump, M., Dang-Vu, B., Masarati, P., Gennaretti, M., Ionita, A., Zaichik, L., Smaili, H., Quaranta, G., Yilmaz, D., et al., "Adverse rotorcraft pilot couplingspast, present and future challenges," *Progress in Aerospace Sciences* **62**, 1–51 (2013).

[32] Harris, J. J. and Black, G. T., "F-22 control law development and flying qualities," *AIAA Paper* , 96–3379 (1996).

[33] Hauser, J., Sastry, S., and Kokotovic, P., "Nonlinear control via approximate input-output linearization: The ball and beam example," *IEEE Transactions on Automatic Control* **37**(3), 392–398 (1992).

[34] Todorov, E. and Jordan, M. I., "Optimal feedback control as a theory of motor coordination," *Nature Neuroscience* **5**(11), 1226–1235 (2002).

[35] Krishnamoorthy, V., Goodman, S., Zatsiorsky, V., and Latash, M. L., "Muscle synergies during shifts of the center of pressure by standing persons: identification of muscle modes," *Biological Cybernetics* **89**(2), 152–161 (2003).

[36] Valero-Cuevas, F. J., Venkadesan, M., and Todorov, E., "Structured variability of muscle activations supports the minimal intervention principle of motor control," *Journal of Neurophysiology* **102**(1), 59–68 (2009).

[37] Todorov, E., "Optimality principles in sensorimotor control," *Nature Neuroscience* **7**(9), 907–915 (2004).

[38] Valero-Cuevas, F. J., [*Fundamentals of Neuromechanics*], Springer (2015).

[39] Parasuraman, R., Molloy, R., and Singh, I. L., "Performance consequences of automation-induced complacency," *The International Journal of Aviation Psychology* **3**(1), 1–23 (1993).

[40] Parasuraman, R. E. and Mouloua, M. E., [*Automation and human performance: Theory and applications.*], Lawrence Erlbaum Associates, Inc (1996).

[41] Donmez, B., Nehme, C., and Cummings, M. L., "Modeling workload impact in multiple unmanned vehicle supervisory control," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions On* **40**(6), 1180–1190 (2010).

[42] Lynas, D. and Horberry, T., "Human factor issues with automated mining equipment," *The Ergonomics Open Journal* **4**(1) (2011).

[43] Harrison, Y. and Horne, J. A., "The impact of sleep deprivation on decision making: a review.," *Journal of Experimental Psychology: Applied* **6**(3), 236 (2000).