Toward experimental validation of a model for human sensorimotor learning and control in teleoperation

Eatai Roth, Darrin Howell, Cydney Beckwith, and Samuel A. Burden

Department of Electrical Engineering, University of Washington, Seattle, WA, USA

ABSTRACT

Humans, interacting with cyber–physical systems (CPS), formulate beliefs about the system's dynamics. It is natural to expect that human operators, tasked with teleoperation, use these beliefs to control the remote robot. For tracking tasks in the resulting human–cyber–physical system (HCPS), theory suggests that human operators can achieve exponential tracking (in stable systems) without state estimation provided they possess an accurate model of the system's dynamics. This internalized inverse model, however, renders a portion of the system state unobservable to the human operator—the zero dynamics. Prior work shows humans can track through observable linear dynamics, thus we focus on nonlinear dynamics rendered unobservable through tracking control. We propose experiments to assess the human operator's ability to learn and invert such models, and distinguish this behavior from that achieved by pure feedback control.

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

1. INTRODUCTION

Human interaction with the physical world is increasingly mediated by automation. Factory workers facilitate and monitor repetitive manufacturing tasks executed by robots to ensure process safety and product quality. Captains of air and watercraft oversee autopilot systems and intervene to compensate for unexpected equipment failures or environmental disturbances. Doctors perform minimally–invasive surgeries via compact but high–precision endoscopic robots. Drivers rely on lanekeeping assistance and automated braking during highway driving, and will soon relinquish control and allow the vehicles to drive themselves. The proliferation of these *human–cyber–physical* systems (HCPS) raises new questions regarding safety and reliability of the closed–loop behavior arising from humans embedded amid automation. How does a human learn to control a complex dynamical system? What is the appropriate mathematical representation of the learned behavior? Is the representation stable, or does it vary across time or task?

Several scientific disciplines have direct bearing on learning of complex motor actions, including motor control and neuromechanics. However, our goals as engineers differ in key ways from those of scientists. First, we require a system–level description of closed–loop interaction between humans and automation. Exceptions notwithstanding,[1–3] most scientific research in this area targets a reductionist description of subsystems isolated in function (sensory processing vs. motor action) or scale (individual neurons vs. entire limbs) without considering whole–body integration in service of a specified task. Second, we are ultimately concerned with accurately predicting the *phenomenology* of humans interacting with autonomy, rather than characterizing the behavior from *first principles*. If a model provides quantitative predictions that enable design of reliable semiautonomous systems, it is irrelevant to the engineer whether the model is faithful to the human's neurophysiology. For instance, in the context of flight control, data–driven quasi–linear continuous–time models with delay have been known for five decades to provide adequate description of closed–loop pilot behavior.[4] Thus, although such models do not

necessarily estimate true human sensorimotor delay,[5] and there are neurobiologically–plausible alternatives to the continuous feedback control hypothesis,[6] the quasi–linear models nevertheless prove useful in engineering applications.

To enable systematic engineering of human/automation sensorimotor loops, we require validated predictive dynamical models for motor control. Human behavioral repertoire is so rich that we must inevitably limit both the complexity of models and scope of their applicability. Since decades of work has focused on observable linear dynamics,[4,7–9] we address nonlinear dynamics rendered unobservable through tracking control. Building on our prior work,[10] we hypothesize that in applications of interest the human operator learns to *invert* automation dynamics by directly translating from desired task to required control input.

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

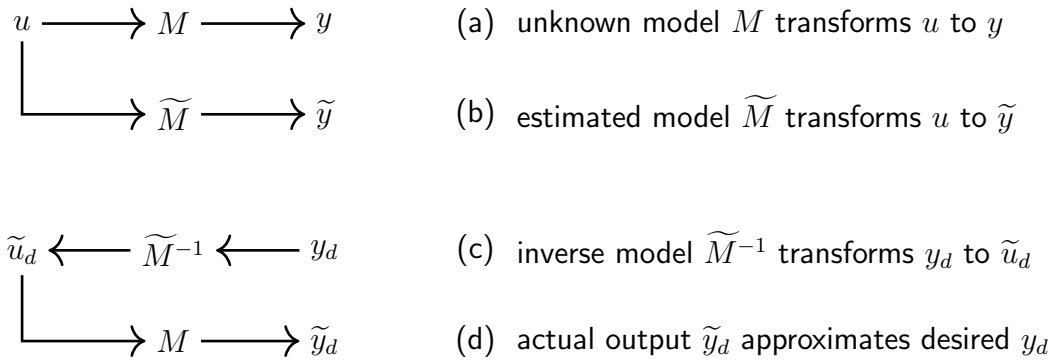This *inverse modeling* framework is illustrated in Figure 1.

$$u \longrightarrow M \longrightarrow y \qquad\qquad \text{(a)} \ \ \text{unknown model } M \text{ transforms } u \text{ to } y$$

$$ \longrightarrow \widetilde{M} \longrightarrow \widetilde{y} \qquad\qquad \text{(b)} \ \ \text{estimated model } \widetilde{M} \text{ transforms } u \text{ to } \widetilde{y}$$

$$\widetilde{u}_d \longleftarrow \widetilde{M}^{-1} \longleftarrow y_d \qquad\qquad \text{(c)} \ \ \text{inverse model } \widetilde{M}^{-1} \text{ transforms } y_d \text{ to } \widetilde{u}_d$$

$$\phantom{\widetilde{u}_d} \longrightarrow M \longrightarrow \widetilde{y}_d \qquad\qquad \text{(d)} \ \ \text{actual output } \widetilde{y}_d \text{ approximates desired } y_d$$

Figure 1. *Inverse modeling framework for human sensorimotor learning and control in teleoperation.*
(a) A forward model M, initially unknown to the human oeprator, transforms input u to output y.
(b) With experience, the human estimates a model $\widetilde{M}$ that transforms input u to predicted output $\widetilde{y}$.
(c) Inverting the estimated model yields $\widetilde{M}^{-1}$ that transforms desired output y_d to predicted desired input $\widetilde{u}_d$.
(d) Applying the prediction $\widetilde{u}_d$ to the original model M yields output $\widetilde{y}_d$ that approximates y_d.
As the fidelity of the model estimate improves $\left(\widetilde{M} \to M \right)$ the actual output converges to the desired $(\widetilde{y}_d \to y_d)$.

This paper concerns experimental validation of the inverse modeling framework proposed in[10] and illustrated in Figure 1. After reviewing the mathematical background in Section 2, we describe a sequence of experimental assays in Section 3 designed to test predictions from the inverse modeling framework and its plausible alternatives. Results from these experiments are given in Section 4 We conclude in Section 5 with a discussion of the results and our plans for future work.

2. BACKGROUND: DYNAMIC INVERSE MODELS IN HCPS

For completeness, in this section we recapitulate some of the mathematical developments from [10, Sec. 2]. Specifically, we begin in Section 2.1 by providing the mathematical framework for inverse modeling, and continue in Section 2.2 by discussing properties of inverse models that have important implications for experiment design.

2.1 Dynamic forward and inverse models

2.1.1 Forward model

We closely follow the development in [10, Sec. 2.1] (which, in turn, follows in [11, Ch. 9]). Consider a single–input, single–output system in control–affine form,

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

where $x \in \mathbb{R}^n$, $f, g \in C^r(\mathbb{R}^n, \mathbb{R}^n)$, $h \in C^r(\mathbb{R}^n, \mathbb{R})$, and $r \in \mathbb{N}$.

Suppose that (1) has *strict relative degree* $\gamma \in \mathbb{N}$ *at* $x_0 \in \mathbb{R}^n$ [11, Def. 9.1]. Then (cf. [11, Sec. 9.2]) there exists a local diffeomorphism $\Phi \in C^r(X, \Phi(X))$ of the form

$$\Phi(x) = \begin{bmatrix} h(x) \\ L_f h(x) \\ \vdots \\ L_f^{\gamma-1} h(x) \\ \zeta(x) \end{bmatrix} \tag{2}$$

where $\zeta \in C^r(X, \mathbb{R}^{n-\gamma})$ and $L_f^\ell h$ denotes the ℓ–th Lie derivative of h along f [11, Sec. 9.1], and $X \subset \mathbb{R}^n$ is a neighborhood containing x_0. Pushing the dynamics in (1) forward via Φ, we obtain the *normal form*

$$\begin{bmatrix} \dot{\xi}_1 \\ \vdots \\ \dot{\xi}_{\gamma-1} \\ \dot{\xi}_\gamma \\ \dot{\zeta} \end{bmatrix} = \begin{bmatrix} \dot{\xi}_2 \\ \vdots \\ \xi_\gamma \\ b(\xi, \zeta) + a(\xi, \zeta)u \\ q(\xi, \zeta) \end{bmatrix}, \tag{3}$$

where $\xi \in \mathbb{R}^\gamma$, $b(\xi, \zeta) = \left(L_f^\gamma h \right) \left(\Phi^{-1}(\xi, \zeta) \right)$, $a(\xi, \zeta) = \left(L_g L_f^{\gamma-1} h \right) \left(\Phi^{-1}(\xi, \zeta) \right)$, and $q : \Phi(X) \to \mathbb{R}^{n-\gamma}$. Note that the input u only affects the states ζ indirectly, through their dynamic coupling with ξ via $\dot{\zeta} = q(\xi, \zeta)$.

Consider the *zero dynamics manifold*[*]

$$Z = \left\{ x \in X \mid \forall \ell \in \{0, \ldots, \gamma - 1\} : L_f^\ell(x) = 0 \right\} \tag{4}$$

with associated *zero dynamics*

$$\forall \zeta \in Z : \dot{\zeta} = q(0, \zeta). \tag{5}$$

If x_0 is an equilibrium point of the undriven dynamics in (1), i.e. $f(x_0) = 0$, then $(\xi_0, \zeta_0) = \Phi(x_0)$ is an equilibrium of the undriven dynamics in (1) since

$$\begin{bmatrix} b(\xi_0, \zeta_0) \\ q(\xi_0, \zeta_0) \end{bmatrix} = D\Phi(x_0)f(x_0) = 0; \tag{6}$$

note that necessarily

$$\xi_0 = \begin{bmatrix} h(x_0) \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \tag{7}$$

Following [11, Def. 9.10], we say that (1) is *locally exponentially minimum phase at* x_0 if $\text{spec}\, D_\zeta q(\xi_0, \zeta_0) \subset \mathbb{C}_-^o$, i.e. if ζ_0 is a locally exponentially stable equilibrium of the dynamics

$$\dot{\zeta} = q(\xi_0, \zeta) \tag{8}$$

obtained by fixing ξ at the equilibrium ξ_0.

---

[*]The neighborhood $X \subset \mathbb{R}^n$ of x_0 must be chosen sufficiently small to ensure Z is a smooth manifold.

2.1.2 Tracking

Suppose the goal is to design a control input $u : \mathbb{R} \to \mathbb{R}$ such that the output tracks a desired reference $y_d \in C^r(\mathbb{R}, \mathbb{R})$ and its derivatives up to order r from a given initial time $t_0 \in \mathbb{R}$. Restricting for simplicity to inputs that are continuously differentiable, $U = C^1(\mathbb{R}, \mathbb{R})$, (1) admits a *flow* $\phi : \mathbb{R} \times \mathbb{R} \times \mathbb{R}^n \times U \to \mathbb{R}^n$ that satisfies

$$\forall t, s \in \mathbb{R}, x \in \mathbb{R}^n, u \in U : \phi(t, s, x, u) = x + \int_s^t \left[f(\phi(\tau, s, x, u)) + g(\phi(\tau, s, x, u))u(\tau) \right] d\tau. \tag{9}$$

Thus for *exact tracking* we seek an initial state $x_d \in \mathbb{R}^n$ and input $u_d \in U$ such that

$$\forall t \geq t_0, \ \ell \in \{0, \dots, r\} : \frac{d^\ell}{dt^\ell} h \circ \phi(t, t_0, x_d, u_d) = \frac{d^\ell}{dt^\ell} y_d(t). \tag{10}$$

Equivalently, using the smooth function norm [12, Ch. 2.1] we may write (10) as

$$\left\| h \circ \phi^{(t_0, x_d, u_d)} - y_d \right\|_{C^r([t_0, \infty))} = 0, \tag{11}$$

where $\phi^{(t_0, x_d, u_d)} \in C^r(\mathbb{R}, \mathbb{R}^n)$ is obtained from ϕ by restricting its last three arguments to (t_0, x_d, u_d). It is possible to achieve exact tracking for the dynamical model only if the initial state x_d can be prepared precisely. In a more general setting, the goal may instead be to design a state feedback law $u_d \in C^1(X, \mathbb{R})$ that achieves *exponential tracking* wherein there exists a neighborhood $C \subset \mathbb{R}^n$ of x_0 and a rate $c > 0$ such that

$$\forall t \geq t_0, \ x \in C, \ y_d \in Y : \left\| h \circ \phi^{(t_0, x, u_d)} - y_d \right\|_{C^r(t)} \leq e^{-c(t-t_0)} \|x - x_0\| \tag{12}$$

for any desired output $y_d \in Y \subset C^r(\mathbb{R}, \mathbb{R})$. As we will see in the next section, some states may not contribute to the output or its derivatives up to order r. In that case, our goal may be to achieve *bounded exponential tracking*, i.e. ensure the state x remains bounded while exponential tracking (12) is achieved.

2.1.3 Inverse model

If (1) has strict relative degree, then the function a is bounded away from zero on $\Phi(X)$, hence if the full state were observable we could apply the feedback–linearizing state feedback

$$u(v, \xi, \zeta) = \frac{1}{a(\xi, \zeta)} (v - b(\xi, \zeta)) \tag{13}$$

to obtain $\frac{d^\gamma}{dt^\gamma} \xi_1 = \dot{\xi}_\gamma = v$, where $v \in \mathbb{R}$ is a new input. Note that $h(x) \equiv \xi_1$, hence this feedback renders the states ζ unobservable. It is straightforward to verify that with $\varepsilon(t, x) = h(x) - y_d(t)$ denoting the *tracking error*, and with z^γ denoting the γ–th derivative of the function z with respect to time, the input

$$v(t, x) = y_d^\gamma(t) - \alpha_{\gamma-1}\varepsilon^{\gamma-1}(t, x) - \alpha_{\gamma-2}\varepsilon^{\gamma-2}(t, x) - \cdots - \alpha_0\varepsilon(t, x) \tag{14}$$

results in exponential tracking for y_d and its first γ derivatives[†] so long as the polynomial $s^\gamma + \alpha_{\gamma-1}s^{\gamma-1} + \cdots + \alpha_1 s^1 + \alpha_0$ is Hurwitz and the unobserved states ζ remain close enough to ζ_0 to ensure $(\xi, \zeta) \in X$ for all time $t \geq t_0$ (cf. [11, Remark 1 to Theorem 9.14]). It will be helpful in what follows to note that v contains a *feedforward* term involving the γ–th derivative of the desired output (y_d^γ) and a *feedback* term involving γ derivatives of the tracking error $(\varepsilon, \varepsilon^1, \dots, \varepsilon^{\gamma-1})$.

We were lead in the preceding paragraph to consider conditions that will ensure ζ remains near ζ_0 under the feedback (14). Applying the feedback in (14), we obtain a *cascaded* system governing the closed–loop dynamics

$$\begin{bmatrix} \dot{\xi} \\ \dot{\zeta} \end{bmatrix} = \begin{bmatrix} A(\xi - \eta(t)) \\ q(\xi, \zeta) \end{bmatrix} =: F(t, \xi, \zeta) \tag{15}$$

[†]Note that if the system (1) has strict relative degree γ then it is (i) trivial to track more than γ derivatives of the output and (ii) impossible to track fewer than γ derivatives using a continuously differentiable input (unless the desired output happens to be γ times continuously differentiable). Therefore in what follows we will assume tracking is desired for $r = \gamma$ derivatives.
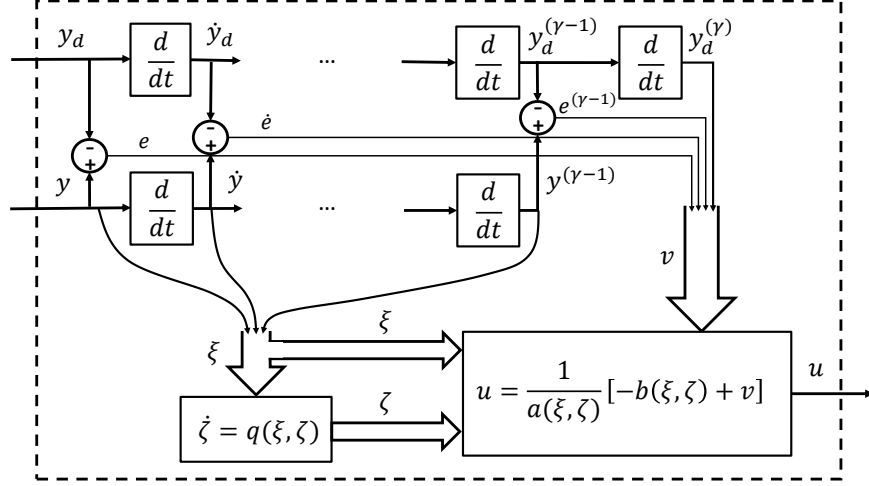
Figure 2. *Block diagram illustration of dynamic inverse model in* (13) *and* (14), *adapted from [11, Fig. 9.3].* The desired output y_d and its derivatives are transformed into coordinates ξ, which are in turn used to determine the states ζ via the zero dynamics $\dot{\zeta} = q(\xi, \zeta)$. The input signal is a combination of the (state–dependent) feedforward inverting input and error feedback.

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

$$\eta(t) = \begin{bmatrix} y_d(t) & \dot{y}_d(t) & \cdots & y_d^{\gamma-1}(t) \end{bmatrix}^\top, \tag{16}$$

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

$$A = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ -\alpha_0 & -\alpha_1 & -\alpha_2 & \cdots & -\alpha_{\gamma-1} \end{bmatrix}, \tag{17}$$

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

$$DF(\xi, \zeta) = \begin{bmatrix} A & 0 \\ D_\xi q(\xi, \zeta) & D_\zeta q(\xi, \zeta) \end{bmatrix}; \tag{18}$$

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

- $\dot{\xi} = A(\xi - z(t))$ is *infinitesimally contracting* [13, Definition 1] (since $\operatorname{spec} A \subset \mathbb{C}_-^o$);

- $\dot{\zeta} = q(\xi, \zeta)$ is *infinitesimally contracting when x is viewed as a parameter* [13, §2] (since $\operatorname{spec} D_\zeta q(\xi_0, \zeta_0) \subset \mathbb{C}_-^o$ and DF is continuous); and

- $D_\xi q(\xi, \zeta)$ is bounded (since DF is continuous).

This enables us to infer (by [13, Theorem 3] or [14, §3.8.3]) that the cascaded system in (15) is infinitesimally contracting on C with some contraction rate $c > 0$. Therefore we conclude (by [13, Theorem 1] or [14, Theorem 2]) that for any reference $y_d \in Y$ where

$$Y = \{ y \in C^\gamma(\mathbb{R}, \mathbb{R}) \mid \|y\|_{C^\gamma} < \infty, \ \forall t \in \mathbb{R} : (\eta(t), \zeta_0) \in C \}, \tag{19}$$

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

$$\forall t \geq t_0, \ (\xi, \zeta) \in C, \ y_d \in Y : \left\| h \circ \widetilde{\phi}^{(t_0, (\xi, \zeta), u_d)} - y_d \right\|_{C^r(t)} \leq e^{-c(t-t_0)} \left\| (\xi, \zeta) - (\eta(0), \zeta_0) \right\|, \qquad (20)$$

where $\widetilde{\phi}$ is the flow in the (ξ, ζ) coordinates, u_d combines the feedback–linearizing state feedback (13) with the exponential tracking feedback (14) determined by y_d, and $\eta(t)$ is given by (16). We summarize the above observations as follows.

Theorem 1 (exponential tracking with a stable model pair) [10, Thm. 1]. If the forward model (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 $C \subset \mathbb{R}^n$ containing x_0 and a rate $c > 0$ such that the time–varying feedforward input obtained by applying (13) and (14) along the trajectory initalized from (ξ_0, ζ_0) achieves bounded exponential tracking (12) at rate c for any desired output $y_d \in Y$ and its first γ derivatives.

2.2 Properties of the dynamic inverse model

In this section we summarize properties of the dynamic inverse model derived in the preceding section; these were previously reported in [10, Sec. 2.2, 4.1].

It is a well–known fact that system states can be rendered unobservable through feedback [15, §8.1].[16] Whenever the relative degree of (1) is strictly smaller than the dimension of the state, the derivations in Section 2.1 show that dynamic inverse models *always* render some states unobservable.

Property (inverting input renders system state unobservable). As noted in [11, §9.2.2], the inverting feedback in (13) renders the states ζ *unobservable*.

This seems to be an undesirable system property, particularly since the input u_d obtained by combining (13), (14) is apparently state–dependent. However, under the assumption that both the original forward model (1) and the closed–loop inverse model (15) are locally exponentially stable (at x_0 and $(\xi_0, \zeta_0) = \Phi(x_0)$, respectively), Theorem 1 [10, Thm. 1] shows that the inverse model does not need access to the actual system's state to achieve bounded exponential tracking. By exploiting the existence of a neighborhood C of the equilibrium (ξ_0, ζ_0) wherein the closed–loop dynamics (15) are infinitesimally contractive, we obtain an alternate route to bounded tracking than that provided by [11, Theorem 9.14], where global conditions were imposed on the closed–loop dynamics of the unobserved states $\dot{\zeta} = q(\xi, \zeta)$. This implies that, once a human operator has learned the full system dynamics through exploratory interactions, it is possible for them to implement the inverse model in (13) without estimating the full system state online—they can simply estimate the task–relevant portion of the state (ξ) and maintain a virtual (internal) copy of the unobservable portion of the state (ζ).

When (1) has strict relative degree, the derivations in Section 2.1 directly demonstrate that the input that achieves exact tracking for a given desired output is uniquely determined.

Property (feedforward input is unique). As noted in [11, §9.2.3], the feedforward input in (14) is unique.

Furthermore, under the hypotheses of Theorem 1, all inputs will converge to the feedforward input at an exponential rate.[‡] This is a valuable property from the experimentalist's perspective. In particular, uniqueness implies that our central hypothesis from Section 1—that humans learn to invert model dynamics—can be directly tested separately from questions concerning how forward or inverse models are learned, represented, or deployed to perform a particular task. Moreover, unlike the study of sensorimotor learning and control involving the person's own body, in the control of cyber–physical systems the experimentalist has access to high–fidelity (or, in the case of simulations like those employed in this paper, exact) models of the system's dynamics. Thus, if the desired trajectory is provided to the operator by the experimentalist, then the inverse modeling framework provides an exact prediction for the human's (asymptotic) input.

[‡]Since the dynamics are infinitesimally contractive, all trajectories converge to one another at an exponential rate.

3. EXPERIMENT DESIGN

3.1 Acquisition of forward and inverse models

3.1.1 Linear forward model with relative degree 1

This section contains the simplest dynamic forward model,

$$\dot{x} = f(x) + g(x)u = u,$$
$$y = h(x) = x, \tag{21}$$

which has strict relative degree 1 (one). This system is already in the normal form derived in Section 2.1, i.e. the change–of–coordinates in (2) is the identity function, $\Phi = \mathrm{id}_\mathbb{R}$, whence $\xi = x$. Therefore, given a desired output $y_d \in C^1(\mathbb{R}, \mathbb{R})$, the corresponding feedforward input is simply the derivative of the output: $u_d = y'_d$.

3.1.2 Linear forward model with relative degree 1, second–order zero dynamics

This section contains one of the simplest dynamic forward models with zero dynamics,

$$\dot{x} = f(x) + g(x)u = \begin{bmatrix} 0 \\ x_3 \\ -c_1(x_1 - x_2) - c_2 x_3 \end{bmatrix} + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} u,$$
$$y = h(x) = \frac{x_1 + x_2}{2}. \tag{22}$$

To transform into the normal form in (2), we compute Lie derivatives of the output with respect to f and g:

$$L_f h(x) = Dh(x)f(x) = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 0 \\ x_3 \\ \dot{x}_3 \end{bmatrix} = x_3,$$
$$L_g h(x) = Dh(x)f(x) = \begin{bmatrix} 1 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = 1; \tag{23}$$

since $Dhf \equiv 1 \neq 0$, the system has strict relative degree 1.

Letting $\xi = h(x) = \frac{1}{2}(x_1 + x_2)$ and choosing $\zeta = x_2$, $\dot{\zeta} = x_3$ to complete the change–of–coordinates, we find

$$\dot{\xi} = \frac{1}{2}(\dot{\zeta} + u), \ \ddot{\zeta} = -2c_1(\xi - \zeta) - c_2\dot{\zeta}; \tag{24}$$

this relative–degree–1 system has second–order zero dynamics.

Furthermore, the system is exponentially minimum phase for any $c_1, c_2 > 0$ since, for a fixed ξ, the zero dynamics are those of a linear spring–mass with setpoint ξ and stiffness $2c_1$, subject to viscous damping c_2.

3.2 Comparison with pure feedback control

3.2.1 Nonlinear forward model with relative degree 1

This section contains the simplest nonlinear dynamic forward model,

$$\dot{x} = f(x) + g(x)u = u,$$
$$y = h(x) = \mathrm{sig}(x), \tag{25}$$

where $\mathrm{sig} : [0, 1] \to [0, 1]$ is a smooth monotonically increasing nonlinear transformation (i.e. a *sigmoid* like that shown in Figure 3). So long as $\mathrm{sig} \in C^1$, the system has strict relative degree 1 (one); we chose the following specific algebraic form for sig and its inverse $\mathrm{gis} = \mathrm{sig}^{-1}$:

$$\forall y \in [-.5, +.5] : \mathrm{sig}(y) = \frac{1}{2\tan^{-1} 2} \tan^{-1} 4y; \tag{26}$$
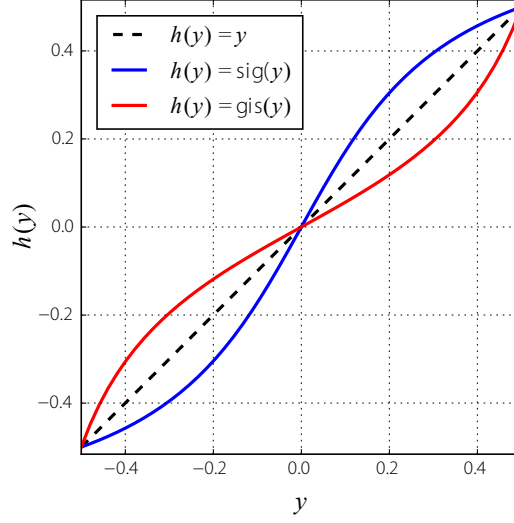
Figure 3. *Sigmoidal output nonlinearities employed in nonlinear forward model with relative degree 1 (Section 3.2.1):* dashed line indicates linear output; solid blue line (shallow slope at $y = 0$) indicates the function sig; solid red line (steep slope at $y = 0$) indicates the inverse function gis $= \text{sig}^{-1}$.

$$\forall y \in [-.5, +.5] : \text{gis}(y) = \frac{1}{4} \tan\left(\left[2 \tan^{-1} 2\right] y\right). \tag{27}$$

This system is not in the normal form derived in Section 2.1, i.e. the change–of–coordinates in (2) is not the identity function. However, if the desired output $y_d \in C^1(\mathbb{R}, [0, 1])$ presented to (21) is transformed into sig $\circ y_d$ before being presentated to (25), the feedforward inputs in both cases are simply the derivative of the original output: $u_d = y'_d$.

3.3 Experiment Instantiation

We propose an experimental assay to test the hypothesis that humans learn forward and inverse models in a human-cyber-physical interaction. We instantiate the above models (Section 3) in a simple path-following video game. Participants perform a SISO trajectory-tracking task (Figure 4), moving 1-DOF slider to guide a cursor (y, purple circle) along a prescribed trajectory (y_d, orange path). The sequence of trajectories comprises sigmoidal steps (cubic spline interpolation) of two different amplitudes (20% and 40% display height) presented in alternating upwards and downwards transitions. Subjects themselves initiate each trial (an individual sigmoidal step iteration), re-centering the input device between presentations. Trials conclude once the user has maintained the cursor to within one cursor radius of the final desired output value for a duration of 1 sec. Trials in which the input reaches an extrema of the input device or the output exceeds the bounds of the display area are immediately terminated (and excluded from analysis) and subsequently repeated.

4. EXPERIMENTAL RESULTS

4.1 Are human responses consistent with the model–inversion framework?

Towards testing whether human subjects learn the linear first-order forward model (Section 3.1), subjects (N=4) were presented a block of low amplitude step trajectories followed immediately by a block of high amplitude steps, each block consisting of 20 iterations. Subjects tended to undershoot the reference output trajectory in the sigmoidal transition. Nonetheless, the interquartile range (calculated across subjects) encompasses the reference trajectory for both amplitudes and for both upward and downward transitions (Figure 5, first and second columns). The symmetry of tracking responses in the step-up and step-down conditions motivates a pooled analysis of the data.

For the first-order forward model, the model-inversion hypothesis prescribes the feedforward input, u_d, be equal to the derivative of the reference input, y'_d (Section 3.1.1). For the sigmoidal reference trajectory, this
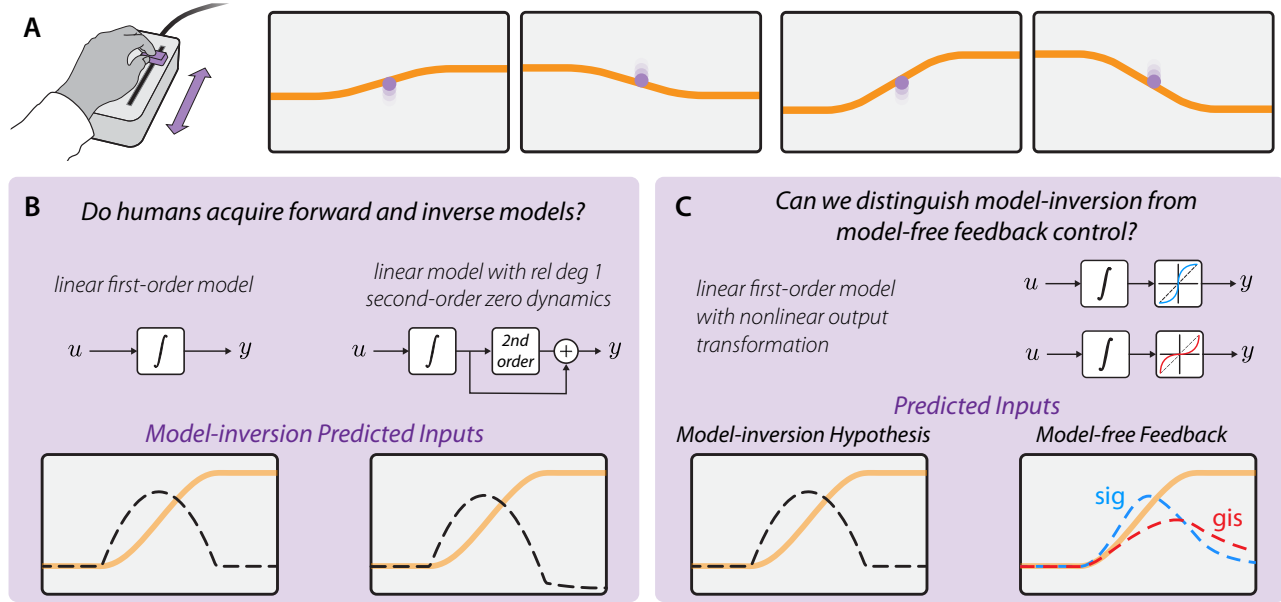
Figure 4. (A) *Experiment instantiation (Section 3.3):* Subjects move a 1-DOF slider to guide a cursor along a sequence of sigmoidal paths. (B) *Assessing model acquisition (Section 3.1):* Block diagrams for the forward models and predicted output (*black dashed*) for sigmoidal input (*gold*). (C) *Comparing model-inversion and model-free feedback control (Section 3.2):* Applying a static nonlinearity at the output does not affect the prediction from the model–inversion hypothesis but does affect the prediction from a pure feedback controller.

input is zero during the initial and final plateau and parabolic in the sigmoidal transition (Figure 5). Though human inputs are similar to the model–inversion predictions, there are notable differences. Both at the initiation and termination of the parabolic input regime, human operators exhibit a smooth transition whereas the model-predicted feedforward input requires an abrupt change in slider velocity. Despite the difference between empirical and model-predicted inputs, human subjects arrived at an input trajectory which sufficiently achieved the task, and moreover, this input trajectory was consistent across individuals. We fit P- and PD-control models (gains fit using linear regression to mean empirical errors and inputs; P-control $k_p = 3.248$; PD-control $k_p = 3.289$, $k_d = -0.113$) and simulate these pure feedback policies for the sigmoidal reference trajectories (Figure 5); P- and PD-feedback models furnish qualitatively similar controllers. These best-fit feedback controllers fail to capture the empirically observed tracking response, lagging noticeably behind measured outputs. Interestingly, empirical inputs fall between the best fit model-free feedback simulation and the model-inversion prediction.

In the subsequent assay, subjects perform the task under the forward model detailed in Section 3.1.2, linear dynamics with relative degree 1 and second-order zero dynamics. The model-predicted feedforward input is similar to that described for the linear first-order forward model above, prescribing zero input at the initial plateau followed by a parabolic input during the sigmoidal step. However, settling at the terminal plateau requires a negative input (Figure 6, bottom row); the model simulates the cursor towing a hidden mass subject to a spring force and viscous drag, and thus, an opposing input at the goal is required to arrest motion. Subjects performed this task for small amplitude transitions (20 replicates). Again, subjects achieved reasonably good output tracking performance; the reference output trajectory nearly lies within the interquartile range throughout. But as in the previous assay, humans produce smoother input trajectories than the model–inversion hypothesis dictates.

4.2 Are human responses distinguishable from model–free responses?

Alternative to the model-inversion hypothesis, humans might not learn the inverse model, relying instead on a model-free feedback policy. And for some high-gain feedback policies, the model-free inputs may be indistinguishable from those predicted under a model–inversion assumption. We propose a follow-up to the first assay aimed at discriminating these two strategies. We select a forward–model condition under which participants
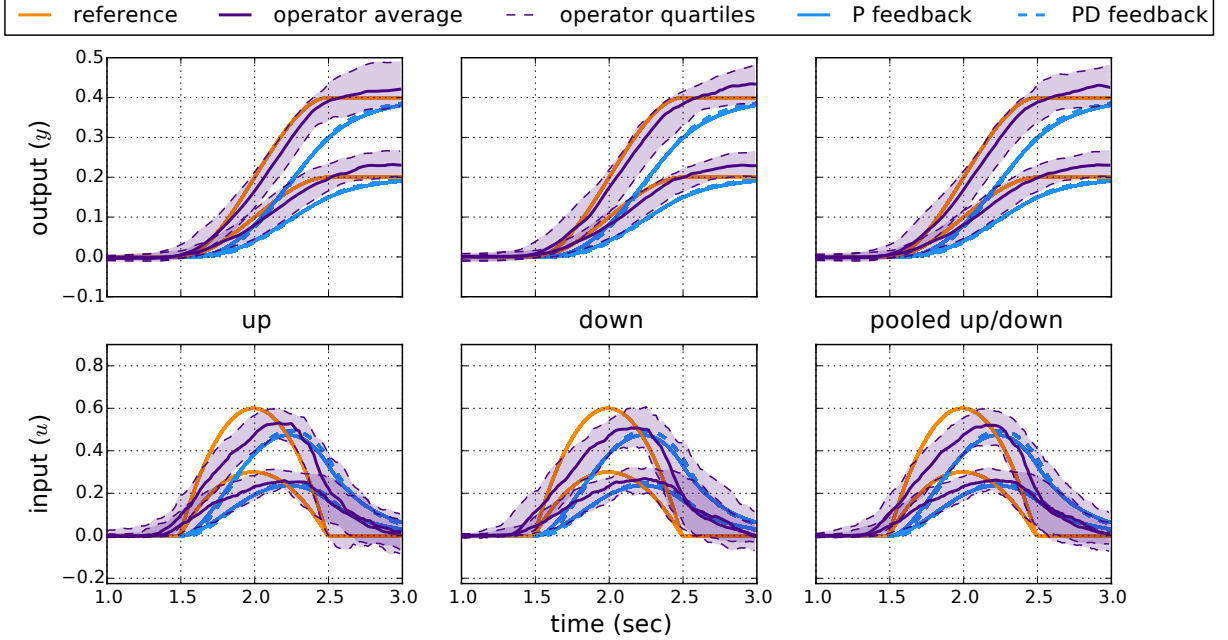
Figure 5. *Results from linear forward model with relative degree 1 (Section 3.1.1).* The output reference and corresponding feedforward input are drawn in gold; purple lines and fill indicate the operator average and interquartile range (4 subjects). The same reference was presented monotone increasing ("up", first column) and decreasing ("down", second column, direction flipped to facilitate comparison with "up") at two different magnitudes (20% and 40% display height); since no significant difference was detected between the presentations, we consider the pooled data in our final analyses (third column). In blue, we simulate responses for the best fit P- and PD-feedback controllers (parameters fit from the mean empirical input). Traces in the top row are baseline adjusted (by subtracting mean from 0 sec to 1 sec) to facilitate comparison between reference magnitudes.

successfully tracked the task trajectories, presumably via model inversion. In a subsequent assay, we repeat these trials with the same forward model, however, we apply a static nonlinear transformation to both the cursor position (the output, y) and the task trajectory (the desired output, y_d). Being that the underlying dynamics of the forward model are identical, and presuming that humans can invert the output nonlinearity, the model–inversion hypothesis predicts the same inputs for both the original assay (Section 3.1.1) and the transformed presentation (Section 3.2.1). In contrast, since the output tracking error ε is similarly nonlinearly transformed, a model-free feedback policy predicts different inputs for the output-transformed condition.

As in Section 4.1 and Figure 5, subjects tracked an assay of output reference trajectories comprising two blocks (low- and high-amplitude steps with 20 iterations per block) for models with the sig and gis output nonlinearities ((26) and (27)). Remarkably, subjects perform better in these tasks and empirical outputs agree more closely with model–prediced inputs as compared to the strictly linear case (Figure 5), regardless of which nonlinearity is appended to the forward model. The model-free feedback policies fit in Section 4.1 do a poor job predicting the output tracking performance. Again, empirical inputs are intermediate to the model–inversion and model–free predictions. It is unsurprising that the sig nonlinearity yields improved output tracking performance; the nonlinearity mimics the shape of the reference output, hastening the cursor through the sigmoid and minimizing output deviations at the plateaus. That subjects performed similarly for the gis–transformed input (surprisingly) supports the model–inversion hypothesis.

5. DISCUSSION: OUTLOOK FOR PREDICTIONS IN HCPS

The experimental results reported in Figures 5 and 6 provide mixed support for our central hypothesis, namely, that humans learn and invert dynamic models to control cyber–physical systems. With the simplest linear
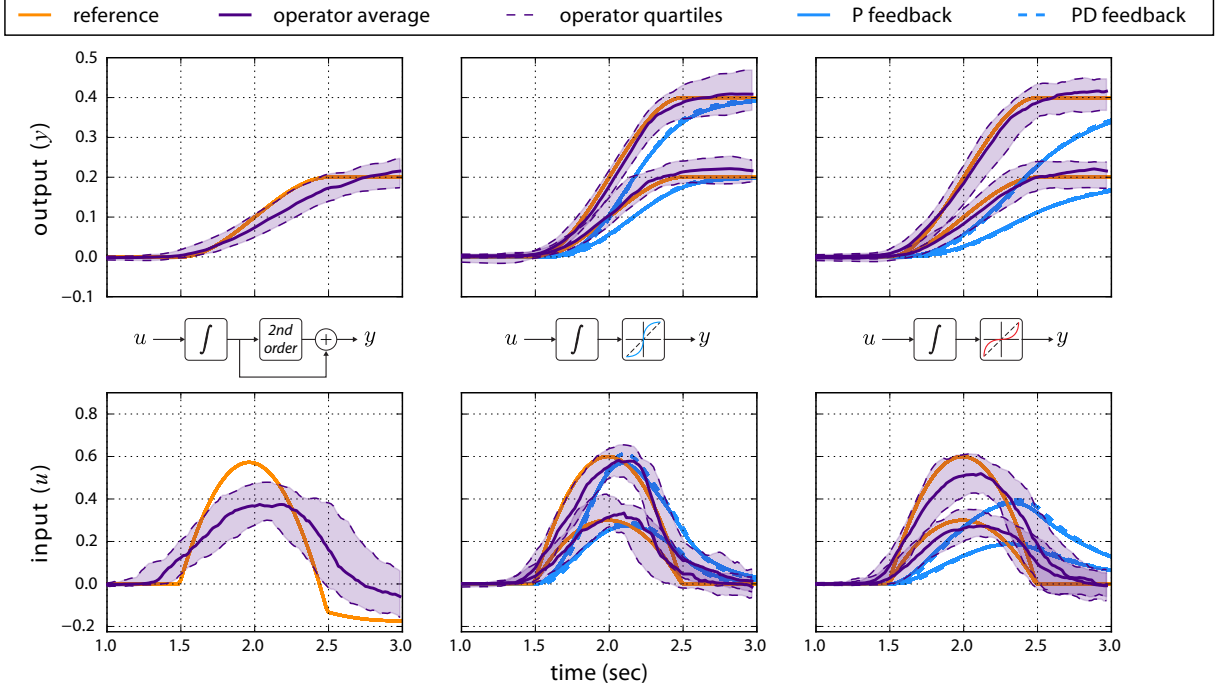
Figure 6. (left) *Results from linear forward model with relative degree 1, second–order zero dynamics (Section 3.1.2).* (middle and right) *Results from nonlinear forward model with relative degree 1 (Section 3.2.1).* Note that the nonlinear output transformation has been inverted to facilitate comparison. The *middle* plot shows the result with the nonlinearity sig (26), while the *right* depicts the result with the inverse nonlinearity gis = sig^{-1} (27). Both nonlinearities are illustrated in Figure 3. See Figure 5 caption for details about the plotting conventions; the same axis limits are employed to facilitate comparison with Figure 5.

model with relative degree 1 (Section 3.1.1, Figure 5), the interquartile distribution of operator outputs contains the prescribed references, but the predicted input does not lie within the interquartile distribution of operator inputs; when second–order zero dynamics are added (Section 3.1.2, Figure 6 *left*), the errors are exacerbated. However, adding a static nonlinearity to the simplest model (Section 3.2.1) results in much tighter tracking, and the predicted input is subsumed within the interquartile distribution of operator inputs (Figure 6, *middle* and *right*). Most provocatively, this result holds regardless of whether the increasing (sig (26)) or decreasing (gis (27)) nonlinearity from Figure 3 is employed. Furthermore, pure feedback appears to be inadequate to explain the results we observe in Figure 6.

Across forward model conditions and reference trajectories, subjects generated smooth inputs whereas the model–inversion prediction demands an abrupt acceleration at the beginning of the sigmoidal region and an abrupt cessation at the terminal plateau. The observed smoothness may reflect a biomechanical limitation, that human users have difficulty actuating the slider with discontinuous velocities. It has been demonstrated for unconstrained, volitional reaching tasks that human subjects minimize the magnitude-squared of jerk of the hand trajectory, minimizing deviations in acceleration and yielding smooth velocities.[17] This constraint may suggest that operators cannot implement the feedforward input for these trajectories, but this does not preclude that humans have learned the forward model.

Though promising, these results indicate that further experimental work is needed to determine when (or whether) the model inversion framework proposed in[10] (and summarized in Section 2) can be relied on to predict human behavior in robot teleoperation. In future work, we will conduct assays involving a wider variety of references, system models, and nonlinearities.

ACKNOWLEDGMENTS

REFERENCES

[1] Nishikawa, K., Biewener, A. A., Aerts, P., Ahn, A. N., Chiel, H. J., Daley, M. A., Daniel, T. L., Full, R. J., Hale, M. E., Hedrick, T. L., et al., "Neuromechanics: an integrative approach for understanding motor control," *Integrative and Comparative Biology* **47**(1), 16–54 (2007).

[2] Schwenk, K., Padilla, D. K., Bakken, G. S., and Full, R. J., "Grand challenges in organismal biology," *Integrative and Comparative Biology* **49**(1), 7–14 (2009).

[3] Roth, E., Sponberg, S., and Cowan, N., "A comparative approach to closed-loop computation," *Current opinion in neurobiology* **25**, 54–62 (2014).

[4] McRuer, D. T. and Jex, H. R., "A review of Quasi–Linear pilot models," *IEEE Transactions on Human Factors in Electronics* **HFE–8**(3), 231–249 (1967).

[5] Gawthrop, P., Loram, I., and Lakie, M., "Predictive feedback in human simulated pendulum balancing," *Biological cybernetics* **101**, 131–146 (2009).

[6] Gawthrop, P., Loram, I., Lakie, M., and Gollee, H., "Intermittent control: a computational theory of human control," *Biological cybernetics* **104**, 31–51 (2011).

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

[8] Allen, R. W. and McRuer, D., "The man/machine control interface—pursuit control," *Automatica: the journal of IFAC, the International Federation of Automatic Control* **15**(6), 683–686 (1979).

[9] McRuer, D. and Weir, D. H., "Theory of manual vehicular control," *Ergonomics* **12**(4), 599–633 (1969).

[10] Robinson, R. M., Scobee, D. R. R., Burden, S. A., and Shankar Sastry, S., "Dynamic inverse models in human-cyber-physical systems," in [*SPIE Defense + Security*], 98361X–98361X–20, International Society for Optics and Photonics (25 May 2016).

[11] Sastry, S. S., [*Nonlinear Systems: Analysis, Stability, and Control*], Springer (1999).

[12] Hirsch, M. W., [*Differential topology*], Springer–Verlag (1976).

[13] Sontag, E. D., "Contractive systems with inputs," in [*Perspectives in Mathematical System Theory, Control, and Signal Processing*], *Lecture Notes in Control and Information Sciences* **398**, 217–228, Springer Berlin Heidelberg (2010).

[14] Lohmiller, W. and Slotine, J.-J., "On contraction analysis for non–linear systems," *Automatica: the journal of IFAC, the International Federation of Automatic Control* **34**(6), 683–696 (1998).

[15] Callier, F. and Desoer, C., [*Linear system theory*], Springer New York (1991).

[16] Carver, S. G., Kiemel, T., Cowan, N. J., and Jeka, J. J., "Optimal motor control may mask sensory dynamics," *Biological cybernetics* **101**(1), 35–42 (2009).

[17] Flash, T. and Hogan, N., "The coordination of arm movements: an experimentally confirmed mathematical model," *Journal of neuroscience* **5**(7), 1688–1703 (1985).