

Characteristics That Make Linear Time-Invariant Dynamic Systems Difficult for Humans to Control

Seyyed Alireza Seyyed Mousavi, Xingye Zhang , T. Michael Seigler, and Jesse B. Hoagg 

Abstract—We present results from an experiment in which 55 human subjects interact with a dynamic system 40 times over a one-week period. The subjects are divided into five groups. For each interaction, a subject performs a command-following task, where the reference command is the same for all subjects and all trials; however, each group interacts with a different linear time-invariant dynamic system. We use a subsystem identification algorithm to estimate the control strategy that each subject uses on each trial. The experimental and identification results are used to examine the impact of the system characteristics (e.g., poles, zeros, relative degree, system order, phase lag) on the subjects' command-following performance and the control strategies that the subjects learn. Results demonstrate that phase lag (which arises from higher relative degree and nonminimum-phase zeros) tends to make dynamic systems more difficult for humans to control, whereas higher system order does not necessarily make a system more difficult to control. The identification results demonstrate that improvement in performance is attributed to: 1) using a comparatively accurate approximation of the inverse dynamics in feedforward; and 2) using a feedback controller with comparatively high gain. Results also demonstrate that system phase lag is an important impediment to a subject's ability to approximate the inverse dynamics in feedforward, and that a key aspect of approximating the inverse dynamics in feedforward is learning to use the correct amount of phase lead in feedforward.

Index Terms—Human behavior, human-in-the-loop (HITL) systems, nonminimum-phase zeros, phase lag, relative degree.

I. INTRODUCTION

HUMANS learn to interact with and control a wide range of dynamic systems. For example, humans learn to drive automobiles, fly aircraft, and ride bicycles. From an automatic-control-system perspective, the characteristics of a dynamic system (e.g., poles, zeros, relative degree, system order, stability, phase lag) often impact the choice of control strategy, and can limit achievable performance [1], [2]. In contrast, the impact of system characteristics on human-in-the-loop (HITL) behavior is less well understood [3], [4].

Manuscript received December 31, 2019; revised May 31, 2020 and October 2, 2020; accepted November 17, 2020. This work was supported in part by National Science Foundation under Grants CMMI-1405257 and OIA-1849213, and in part by the Kentucky Science and Engineering Foundation under Grant KSEF-148-502-15-364. This article was recommended by Associate Editor M. Mulder. (Corresponding author: Jesse B. Hoagg.)

The authors are with the Department of Mechanical Engineering, University of Kentucky, Lexington, KY 40506 USA (e-mail: alirezamousavi@uky.edu; xingyechang86@gmail.com; tmseigler@uky.edu; jesse.hoagg@uky.edu).

This article has supplementary material provided by the authors and color versions of one or more figures available at <https://doi.org/10.1109/THMS.2020.3046164>.

Digital Object Identifier 10.1109/THMS.2020.3046164

Human learning and motor control have been studied using a variety of experimental methods (e.g., [5]–[9]). Furthermore, a wide range of mathematical models for human learning and motor control have been proposed (e.g., [10]–[20]). For example, the internal model hypothesis proposes that the central nervous system constructs models of the body's interactions with the physical world and that those models are continually updated and used for control [21]–[23].

HITL behavior is also studied using system identification approaches to obtain models of the human control behavior using data from HITL experiments [24]–[39]. However, many of these studies incorporate assumptions regarding the human's control strategy. For example, the models in [31] and [36] incorporate an assumed feedforward control strategy, specifically, the feedforward controller approximates the inverse dynamics of the system with which the human interacts. It is also worth noting that many existing studies examine humans interacting with relatively simple dynamic systems (often referred to as *controlled elements* in the manual control literature). For example, Pool *et al.* [31]–[35] considered systems that are either a static gain, a single integrator, or a double integrator. Similarly, Pool *et al.* [36]–[38] considered dynamic systems that consist of a single integrator cascaded with a first-order low-pass filter. However, these studies do not consider more complex dynamic systems, for example, dynamic systems that have order greater than two, relative degree greater than two, complex-conjugate poles, and/or system zeros.

In [40], a subsystem identification (SSID) algorithm is presented, which can be used to identify the best-fit linear time-invariant (LTI) model of the control strategy that a human uses in an HITL experiment. This method does not require that a specific control strategy is assumed *a priori*. Instead, the SSID method identifies the LTI control strategy that is the best fit to the data from the HITL experiment. In [41], this method is used to model the feedforward and feedback control that subjects use in an HITL experiment, where subjects interact with an LTI system and perform a command-following (i.e., pursuit tracking) task. The results in [41] demonstrate that subjects learn to update the feedforward (i.e., anticipatory) control until it approximates the inverse dynamics of the system with which the subjects interact; this result supports the internal model hypothesis [21]–[23]. However, Zhang *et al.* [41] examined only one LTI system, which raises several questions. First, do humans learn to approximate the inverse dynamics in feedforward for other systems? Second, what system characteristics (e.g., poles, zeros, relative degree, system order, phase lag) can make a dynamic

system difficult for humans to control and why? Third, how do system characteristics impact a human's ability to approximate the inverse dynamics in feedforward?

Understanding the impact of system characteristics on HITL control behavior could have application to a variety of HITL technologies, including human-operated vehicles and devices. For example, higher relative degree can arise in structural systems where sensors and actuators are not collocated [42]. Furthermore, nonminimum-phase zeros can arise in human-operated vehicles (e.g., automobiles and aircraft). For example, driving a car backwards (e.g., parallel parking) is a dynamic system that often has real nonminimum-phase zeros [1]. Similarly, an aircraft's elevator-to-vertical-acceleration transfer function is often nonminimum phase [43].

This article examines the impact of system characteristics on HITL control strategies. We present results from an HITL experiment, where 55 subjects interact with an LTI system 40 times over a one-week period. The subjects are divided into five groups, where each group interacts with a different LTI system. For each interaction, each subject performs a command-following task, where the reference command (i.e., forcing function) is the same for all trials and all subjects. We use an SSID algorithm to identify the control strategy (feedforward, feedback, feedback time delay) that each subject uses on each trial. The results are used to examine the impact of the system characteristics on the subjects' performance and the control strategies that the subjects learn.

This article is organized as follows. Section II describes the experimental setup and approach. Section III examines the subjects' command-following performance in both the time domain and frequency domain. The remainder of this article is dedicated to explaining why the command-following performances differs between the groups (i.e., between the different dynamic systems). Our approach is to estimate a model of each subject's control strategy on each trial, and use these models to help explain the performance differences. Section IV presents the SSID algorithm that is used to obtain the best-fit models to each subject's experimental data on each trial. Then, Sections V and VI examine the identified feedforward and feedback controllers, respectively. Finally, Section VII summarizes the results and provides conclusions.

II. EXPERIMENTAL METHODS

In total, 55 people voluntarily participated in this study. At the time of this study, the subjects were 18 to 35 years of age, and they had no known motor control or neurological disorders. The University of Kentucky's Institutional Review Board approved this study under IRB protocol 44649.

In this study, subjects use a single-degree-of-freedom rotational joystick to affect the horizontal position of a controlled object that is displayed on the screen of a computer. The position of the joystick is denoted by u , which is the input to an LTI system. The horizontal position of the controlled object is denoted by y , which is the output of the LTI system. Another object also moves on the computer screen, and its horizontal position

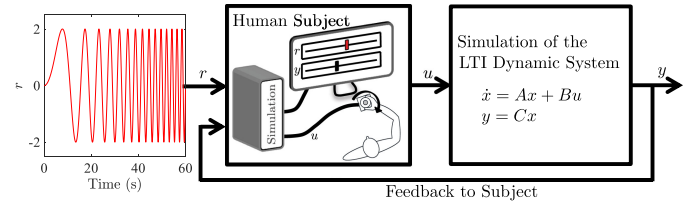


Fig. 1. Experimental setup. A subject uses a joystick to affect the horizontal position y of a controlled object displayed on a computer screen. The joystick position u is the input to an LTI system, and the controlled object's position y is the output of the LTI system. A reference object is also displayed on the computer screen, and its position r is a 60-s chirp signal.

is denoted by r , which is independent of u . This object is called the reference object. The signals u , y , and r are functions of time t . Fig. 1 is a diagram of the experimental setup.

Prior to interacting with the experimental setup, each subject is shown the computer screen and told that manipulating the joystick moves the controlled object. Subjects are told that their objective is to manipulate the joystick and attempt to make the controlled and reference objects have the same horizontal position at each instant of time. Thus, each subject's objective is to generate a control u that makes the magnitude of the error $e \triangleq r - y$ as small as possible. Prior to the experiment, the subjects have no knowledge of the LTI system relating u and y , or the reference object's trajectory r .

Each subject performs 40 trials of the experiment over 7 days. A trial is a 60-s time period during which a subject operates the joystick. Each subject's trials are divided into 4 sessions, and each session consists of 10 trials, which are completed within a period of 20 min. No subject participates in more than one session in a 12-h period.

For each session, a subject is placed in an isolated area, which is free from distraction. The subject sits in a chair facing a computer screen, which is located approximately 60 cm from the subject's eyes and measures 47.6 cm high by 26.8 cm wide. The subject uses a hand of their choice to manipulate the single-degree-of-freedom rotational joystick, which is a Teledyne Gurley model number 8225-6000-DQSD.

The reference object's position r is a 60-s chirp with frequency content between 0 and 0.5 Hz. Specifically, for $t \in [0, 60]$, $r(t) \triangleq 2 \sin \frac{\pi t^2}{120}$. The upper frequency limit of 0.5 Hz is selected such that the reference r is within the frequency range of achievable human motion. Note that [44, Ch. 21] states that well-trained human subjects interacting with a static-gain system can accurately track sinusoidal commands up to approximately 2 Hz. However, Yu *et al.* [45] indicated that for subjects interacting with a single integrator, the upper frequency limit is significantly less (approximately 0.6 Hz). For the more complicated LTI dynamic systems considered in this study, experimental tuning was used to determine the upper frequency limit of 0.5 Hz. We note that this upper frequency limit is consistent with some of the sinusoidal commands used in [46]. The units of the reference r are hash marks (hm), which are vertical lines shown on the computer screen to denote position. The distance between hash




Fig. 2. Bode plots of $\mathcal{G}_m$, $\mathcal{G}_n$, $\mathcal{G}_{rd1}$, $\mathcal{G}_{rd2}$, and $\mathcal{G}_{rd3}$.

marks on the computer screen is 2.5 cm. The amplitude of the reference r is 2 hm, and the range of motion displayed on the computer screen is ± 8 hm.

For $t \in [0, 60]$, the controlled object's position y satisfies the LTI differential equation

$$\dot{x}(t) = Ax(t) + Bu(t) \quad (1)$$

$$y(t) = Cx(t) \quad (2)$$

where $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times 1}$, $C \in \mathbb{R}^{1 \times n_x}$, $x(t) \in \mathbb{R}^{n_x}$ is the state, and the initial condition is zero (i.e., $x(0) = 0$). It follows from (1) and (2) that the transfer function from u to y is $\mathcal{G}(s) \triangleq C(sI - A)^{-1}B$. The dynamic system (1) and (2) is simulated using a dSPACE DS1103 control board. The DS1103 board also measures u , and the ControlDesk software is used to display the controlled and reference objects on the computer screen. The joystick is connected to the dSPACE DS1103 connector panel using a 10-pin encoder cable.

The 55 subjects are divided into 5 groups, where each group has 11 subjects. The subjects interact with the LTI system (1) and (2), where the system matrices (i.e., A , B , C) are selected to explore the effects of different system characteristics. In particular, each group interacts with the LTI system (1) and (2), where the associated transfer function $\mathcal{G}$ is equal to one of the following five transfer functions:

$$\begin{aligned} \mathcal{G}_m(s) &\triangleq \frac{2(s+2.2)}{s^2+3.6s+4}, & \mathcal{G}_n(s) &\triangleq \frac{-2(s-2.2)}{s^2+3.6s+4} \\ \mathcal{G}_{rd1}(s) &\triangleq \frac{1.6(s+1)}{(s+1.6)}\mathcal{G}_m(s), & \mathcal{G}_{rd2}(s) &\triangleq \frac{1.6}{(s+1.6)}\mathcal{G}_m(s) \\ \mathcal{G}_{rd3}(s) &\triangleq \frac{27.1}{(s+2.2)(s^2+3.6s+4)}. \end{aligned}$$

These dynamic systems are selected to examine the impact of system characteristics, including system zeros, relative degree, system order, and phase shift. Fig. 2 shows the Bode plots for $\mathcal{G}_m$, $\mathcal{G}_n$, $\mathcal{G}_{rd1}$, $\mathcal{G}_{rd2}$, and $\mathcal{G}_{rd3}$.

All five transfer functions have a pair of complex-conjugate poles at $-1.8 \pm j0.872$. However, the transfer functions differ by the other poles and zeros. For example, $\mathcal{G}_m$ and $\mathcal{G}_n$ have the

same poles and the same dc gain. However, $\mathcal{G}_m$ is minimum phase with a zero at -2.2 , while $\mathcal{G}_n$ is nonminimum phase with a zero at $+2.2$. The magnitude of $\mathcal{G}_m$ and $\mathcal{G}_n$ coincide, whereas $\mathcal{G}_n$ has 180° more asymptotic phase lag than $\mathcal{G}_m$. We examine the impact of nonminimum-phase zeros by comparing subjects' behavior with $\mathcal{G}_m$ to that with $\mathcal{G}_n$.

We examine the impact of system order and relative degree by comparing subjects' control behavior with $\mathcal{G}_m$ (order 2, relative degree one) to that with $\mathcal{G}_{rd1}$ (order 3, relative degree one) and that with $\mathcal{G}_{rd2}$ (order 3, relative degree 2).

To examine the impact of phase lag, note that $\mathcal{G}_m$ and $\mathcal{G}_n$ have the same magnitude but different phases. In contrast, $\mathcal{G}_{rd3}$ and $\mathcal{G}_n$ have the same phase but different magnitudes. Thus, we examine the impact of phase shift by comparing the subjects' control behavior with $\mathcal{G}_m$, $\mathcal{G}_n$, and $\mathcal{G}_{rd3}$.

The transfer functions are selected such that all poles and zeros lie in the 0 -to- π rad/s frequency range of the chirp r , which is also within the frequency range of achievable human motion. In addition, the magnitudes of all poles and zeros are separated by at least 0.2 rad/s to ensure that the impact of poles and zeros do not approximately cancel (i.e., to prevent approximate pole-zero cancellation). First, the poles and zeros of $\mathcal{G}_m$ and $\mathcal{G}_n$ are selected through experimental tuning. These poles and zeros also determine the poles of $\mathcal{G}_{rd3}$, because $\mathcal{G}_{rd3}$ is selected to have the same phase as $\mathcal{G}_n$. Next, $\mathcal{G}_{rd2}$ is selected to have the poles and zero of $\mathcal{G}_m$ and an extra pole at -1.6 , which is selected to have magnitude less than the magnitudes of the other poles and zero of $\mathcal{G}_{rd2}$. This ensures that the extra pole at -1.6 has significant impact over the 0 -to- π rad/s. Finally, $\mathcal{G}_{rd1}$ is selected to have the poles and zero of $\mathcal{G}_{rd2}$ and an extra zero at -1 , which is selected to have magnitude less than the magnitudes of the poles of $\mathcal{G}_{rd1}$.

The zero-frequency gain for each transfer function is selected so that the joystick motion required to follow r is within the comfortable range of human motion (i.e., wrist twist). Note that $\mathcal{G}_m$, $\mathcal{G}_n$, $\mathcal{G}_{rd1}$, and $\mathcal{G}_{rd2}$ have the same zero-frequency gain (i.e., $\mathcal{G}_m(0) = \mathcal{G}_n(0) = \mathcal{G}_{rd1}(0) = \mathcal{G}_{rd2}(0) = 1.1$). In contrast, the zero-frequency gain $\mathcal{G}_{rd3}(0) = 3.1$ is higher than that of the other transfer functions because $\mathcal{G}_{rd3}$ has relative degree 3, which results in more magnitude change over the 0 -to- π rad/s range. Thus, the zero-frequency gain is selected as $\mathcal{G}_{rd3}(0) = 3.1$ so that the joystick motion required to follow r is within the comfortable range of human motion (i.e., less than 180° of wrist twist).

For each of the 2200 trials, we record r , u , and y with a sample time of $T_s = 0.02$ s. The sampled data are denoted by $\{r_k\}_{k=1}^n$, $\{u_k\}_{k=1}^n$, and $\{y_k\}_{k=1}^n$, where $n = 3000$ samples. For $k \in \{1, \dots, n\}$, we define $e_k \triangleq r_k - y_k$, which is the command-following error.

III. ANALYSIS OF COMMAND-FOLLOWING PERFORMANCE

A *divergent trial* is a trial in which for any $k \in \{1, \dots, n\}$, y_k exceeds ± 8 hm display limits. As shown in Table I, there are more divergent trials during the earlier trials than during the later trials. For each group, there are no divergent trials over the last

TABLE I
NUMBER OF DIVERGENT TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Total
$\mathcal{G}_m$	1	0	0	0	0	0	1
$\mathcal{G}_n$	7	2	1	0	0	0	10
$\mathcal{G}_{rd1}$	0	0	0	0	0	0	0
$\mathcal{G}_{rd2}$	0	0	1	0	0	0	1
$\mathcal{G}_{rd3}$	5	1	5	5	3	0	19




Fig. 3. Mean and standard deviation of $\|e\|$ on each trial for each group. For each group, the mean and standard deviation of $\|e\|$ tend to decrease over the trials. The $\circ$ is the mean, and the lines indicate standard deviation.

TABLE II
MEAN $\|e\|$ AND CHANGE IN MEAN $\|e\|$ FROM FIRST FIVE TO
LAST FIVE TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
$\mathcal{G}_m$	0.52	0.35	0.30	0.25	0.22	0.22	−0.30
$\mathcal{G}_n$	1.35	1.00	0.82	0.70	0.64	0.53	−0.82
$\mathcal{G}_{rd1}$	0.38	0.30	0.26	0.24	0.22	0.22	−0.16
$\mathcal{G}_{rd2}$	0.72	0.57	0.45	0.41	0.36	0.34	−0.38
$\mathcal{G}_{rd3}$	0.95	0.75	0.68	0.57	0.56	0.54	−0.41

five trials. Divergent trials are omitted from the results reported in the rest of this article.

A. Time-Domain Analysis

For each trial, we define the time-averaged error

$$\|e\| \triangleq \frac{1}{n} \sum_{k=1}^n |e_k|.$$

Fig. 3 shows the mean and standard deviation of $\|e\|$ on each trial for each group, and Table II shows the mean $\|e\|$ on different sets of trials for each group. For each group, the mean and standard deviation of $\|e\|$ tend to decrease over the trials. In particular, for each group, the mean $\|e\|$ decreases by at least 42% from the first five trials to the last five trials. The group with $\mathcal{G}_n$ has the largest decrease (in absolute and percent) from the first five trials to the last five trials. For all groups, the open-loop (i.e., $u = 0$ and thus, $e = r$) time-averaged error is approximately 1.2. By trial 40, the mean $\|e\|$ for each group is at least 50% better than the open loop. We note that the mean $\|e\|$ for each group could continue to decrease with additional practice (e.g., trials). In particular,

the group with $\mathcal{G}_n$ appears to have a downward trajectory in mean $\|e\|$ over the last few trials in comparison to the other groups.

A one-way ANOVA comparing each group's mean $\|e\|$ over the last five trials yields $F_{4,11} = 23.7$ and $p < 0.001$, thus confirming a statistical difference between groups. A Tukey post-hoc pairwise comparison yields $p < 0.001$ for all pairs except $(\mathcal{G}_m, \mathcal{G}_{rd2})$, $(\mathcal{G}_{rd1}, \mathcal{G}_{rd2})$, $(\mathcal{G}_m, \mathcal{G}_{rd1})$, and $(\mathcal{G}_n, \mathcal{G}_{rd3})$. The p -values for pairs $(\mathcal{G}_m, \mathcal{G}_{rd2})$ and $(\mathcal{G}_{rd1}, \mathcal{G}_{rd2})$ are $p = 0.098$ and $p = 0.093$. In contrast, there is no statistical difference for pairs $(\mathcal{G}_m, \mathcal{G}_{rd1})$ and $(\mathcal{G}_n, \mathcal{G}_{rd3})$.

Therefore, over the last five trials, the mean $\|e\|$ with $\mathcal{G}_m$ is smaller than that with $\mathcal{G}_n$, which aligns with control theory results showing that nonminimum-phase zeros can limit performance [1], [2]. Table II also shows that over the last five trials, the mean $\|e\|$ with $\mathcal{G}_m$ is less than that with $\mathcal{G}_{rd2}$. Recall that the only difference between $\mathcal{G}_m$ and $\mathcal{G}_{rd2}$ is that $\mathcal{G}_{rd2}$ has an additional pole at -1.6 . Thus, $\mathcal{G}_{rd2}$ is higher order and higher relative degree than $\mathcal{G}_m$, which suggests that higher system order and/or higher relative degree could be impediments to a human's ability to control a dynamic system. However, the mean $\|e\|$ with $\mathcal{G}_{rd1}$ over the last five trials is comparable to that with $\mathcal{G}_m$; both of which are smaller than the mean $\|e\|$ with $\mathcal{G}_{rd2}$ over the last five trials. Furthermore, the only difference between $\mathcal{G}_{rd2}$ and $\mathcal{G}_{rd1}$ is that $\mathcal{G}_{rd1}$ has an additional zero at -1 . Together, these observations suggest that higher relative degree can make a system more difficult to control, whereas higher system order does not necessarily make a system more difficult to control.

Both nonminimum-phase zeros and relative degree contribute to phase lag. In particular, each nonminimum-phase zero and each integer increase in relative degree causes 90° of asymptotic phase lag. Thus, another interpretation of the abovementioned observations is that phase lag can make systems difficult for humans to control. Table II shows that over the last five trials, the mean $\|e\|$ with $\mathcal{G}_n$ and $\mathcal{G}_{rd3}$ are approximately equal but greater than the mean $\|e\|$ with $\mathcal{G}_{rd2}$, which is greater than the mean $\|e\|$ with $\mathcal{G}_{rd1}$ and $\mathcal{G}_m$. Furthermore, over the 0-to-0.5 Hz range, $\mathcal{G}_n$ and $\mathcal{G}_{rd3}$ have the same phase lag, which is greater than the phase lag of $\mathcal{G}_{rd2}$, which is greater than the phase lag of $\mathcal{G}_m$ and $\mathcal{G}_{rd1}$. These observations suggest that phase lag is a key characteristic that makes systems difficult for humans to control.

B. Frequency-Domain Analysis

For each trial, we calculate the discrete Fourier transform (DFT) of $\{y_k\}_{k=1}^n$ and $\{r_k\}_{k=1}^n$ at frequencies $\omega_i = \pi(i - 1)/30$ rad/s, where $i \in \{1, 2, \dots, N\}$, which are $N = 31$ evenly spaced frequencies over the 0-to-0.5 Hz range. Let $y_{\text{dft}}(\omega_i)$ and $r_{\text{dft}}(\omega_i)$ denote the DFT of $\{y_k\}_{k=1}^n$ and $\{r_k\}_{k=1}^n$ at ω_i , respectively.

For each trial, define the *frequency-averaged error in the magnitude of the output*

$$E_m \triangleq \frac{1}{N} \sum_{i=1}^N \left| |y_{\text{dft}}(\omega_i)| e^{j\angle r_{\text{dft}}(\omega_i)} - |r_{\text{dft}}(\omega_i)| e^{j\angle y_{\text{dft}}(\omega_i)} \right|$$

which is the frequency-averaged magnitude of the difference between the y_{dft} and r_{dft} assuming that the phase of y_{dft} is




Fig. 4. Mean and standard deviation of E_m and E_p on each trial for each group. For each group, the mean E_m and E_p tend to decrease over the trials. However, the mean E_p decreases more (in absolute and percent) than the mean E_m . The $\circ$ is the mean, and the lines indicate standard deviation.

TABLE III
MEAN $E_m \times 10^{-2}$ AND CHANGE FROM FIRST FIVE TO LAST FIVE TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
G_m	4.06	2.82	2.62	2.27	1.82	1.89	-2.17
G_n	8.80	7.70	6.85	5.76	5.36	4.45	-4.35
G_{rd1}	2.81	2.39	2.13	2.01	1.82	1.82	-0.99
G_{rd2}	6.12	5.35	4.18	3.89	3.46	3.21	-2.91
G_{rd3}	10.58	6.85	6.81	6.00	6.09	5.73	-4.85

TABLE IV
MEAN $E_p \times 10^{-2}$ AND CHANGE FROM FIRST FIVE TO LAST FIVE TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
G_m	7.13	4.08	3.37	2.65	2.32	2.38	-4.75
G_n	17.61	12.37	9.51	7.92	6.85	5.55	-12.06
G_{rd1}	4.50	3.15	2.71	2.44	2.18	2.15	-2.35
G_{rd2}	10.30	7.57	5.66	4.87	4.06	3.69	-6.61
G_{rd3}	14.56	10.76	9.75	8.22	8.07	7.55	-7.01

equal to the phase of r_{dft} . Similarly, for each trial, define the *frequency-averaged error in the phase of the output*

$$E_p \triangleq \frac{1}{N} \sum_{i=1}^N \left| |r_{dft}(\omega_i)| e^{j\angle y_{dft}(\omega_i)} - |r_{dft}(\omega_i)| e^{j\angle r_{dft}(\omega_i)} \right|$$

which is the frequency-averaged magnitude of the difference between the y_{dft} and r_{dft} assuming that the magnitude of y_{dft} is equal to the magnitude of r_{dft} . By comparing E_m and E_p , we can determine if difference between y_{dft} and r_{dft} is due more to error in magnitude or error in phase.

Fig. 4 shows the mean and standard deviation of E_m and E_p for each group on each trial. Tables III and IV show the mean E_m and E_p on different sets of trials for each group. These results are similar to the time-domain results shown in Fig. 3 and Table II. Specifically, the mean E_m and mean E_p both decrease over the trials. However, for each group, the mean E_p is generally greater




Fig. 5. Model of the control strategy. The control strategy is modeled using the feedforward transfer function G_{ff} , feedback transfer function G_{fb} , and feedback delay d .

than the mean E_m , which suggests that the subjects' command-following error is a result of error in phase more than error in magnitude.

A paired t -test of change in E_m and change in E_p from the first five trials to the last five trials yields $p < 0.05$ for each group except G_{rd3} . Therefore, the mean E_p decreases more (in absolute and percent) than the mean E_m . In particular, for G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} , the mean E_m decreases by 53%, 49%, 35%, 48%, and 46% from the first five trials to the last five trials. In contrast, the mean E_p decreases by 67%, 68%, 52%, 64%, and 48% in the same order. This observation suggests that the improvement in $\|e\|$ is attributed more to improvement in matching the phase of the reference than improvement in matching its magnitude.

IV. MODELING CONTROL STRATEGIES USING SUBSYSTEM IDENTIFICATION

We discretize G using a zero-order hold on the input with sample time $T_s = 0.02$ s, which yields the discrete-time transfer function G . Thus, (1) and (2) imply that

$$\hat{y}(z) = G(z)\hat{u}(z) \quad (3)$$

where $\hat{u}$ and $\hat{y}$ are the z -transforms of u_k and y_k . Let G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} denote the discrete-time transfer functions obtained by discretizing G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} with a sample time $T_s = 0.02$ s and a zero-order hold on the input. Thus, G is equal to G_m , G_n , G_{rd1} , G_{rd2} , or G_{rd3} as appropriate.

Each subject's control strategy is modeled by the LTI control architecture shown in Fig. 5, which is given by

$$\hat{u}(z) = z^{-d}G_{fb}(z)\hat{e}(z) + G_{ff}(z)\hat{r}(z) \quad (4)$$

where $\hat{r}$ and $\hat{e}$ are the z -transforms of r_k and e_k ; G_{ff} is the feedforward transfer function; G_{fb} is the feedback transfer function; and the nonnegative integer d is the feedback delay. Feedforward is the anticipatory control determined solely from the reference r_k , whereas feedback is the reactive control determined from the observed error e_k . The feedback delay d can model sensory and/or motor time delay in the feedback path. Sensory time delay in the feedback path arises from optical perception and processing as a subject views the position r_k of the reference object and the position y_k of the controlled object on the computer screen, and compares those positions to determine the error e_k . Motor time delay in the feedback path indicates the time it takes

the subject to execute a desired feedback control signal (i.e., $G_{fb}\hat{e}$). We do not explicitly include feedforward time delay in (4) because the same reference r_k is used for each trial, and r_k is predictable (see [31], [47] for more discussion of predictable signals). Thus, subjects are able to learn to compensate for feedforward delay (which can arise from sensory and/or motor time delay). Note that subjects may not completely compensate for feedforward time delay on early trials; however, on these early trials, the feedforward transfer function G_{ff} can capture some time delay (i.e., poles at zero) even though such poles are not explicitly indicated in (4).

It follows from (3) and (4) that the closed-loop transfer function from r_k to y_k is

$$\tilde{G}_{yr}(z) \triangleq \frac{G(z) [G_{ff}(z) + z^{-d}G_{fb}(z)]}{1 + z^{-d}G_{fb}(z)G(z)}. \quad (5)$$

We use the SSID algorithm in [40] and [41] to determine the control strategy (i.e., G_{ff} , G_{fb} , d) of the form (4) that is the best-fit to the experimental data for each subject on each trial. Since y_k is bounded, we assume that $\tilde{G}_{yr}$ is asymptotically stable (i.e., its poles are in the open unit disk of the complex plane). Thus, (5) implies that G_{ff} is asymptotically stable. The SSID algorithm in [40] and [41] can be implemented in a computationally efficient manner if G_{ff} is finite impulse response (FIR). Thus, we let G_{ff} be FIR, which does not significantly restrict the class of feedforward behavior [40].

For each trial and for $i \in \{1, 2, \dots, N\}$, define the closed-loop frequency-response data $H(\omega_i) \triangleq y_{dft}(\omega_i)/r_{dft}(\omega_i)$. The objective of the SSID algorithm is to determine G_{ff} , G_{fb} , and d such that the modeled frequency response $\{\tilde{G}_{yr}(e^{j\omega_i T_s})\}_{i=1}^N$ approximates the frequency-response data $\{H(\omega_i)\}_{i=1}^N$. The SSID algorithm aims to find G_{ff} , G_{fb} , and d that minimize the cost function

$$J(G_{ff}, G_{fb}, d) \triangleq \sum_{i=1}^N \left| \tilde{G}_{yr}(e^{j\omega_i T_s}) - H(\omega_i) \right|^2 \quad (6)$$

subject to the constraint that $\tilde{G}_{yr}$ is asymptotically stable. The SSID algorithm is summarized as follows.

- 1) Select a feedback candidate pool that contains possible models for G_{fb} and d . Every element in the feedback candidate pool is such that $\tilde{G}_{yr}$ is asymptotically stable.
- 2) For each element in the feedback candidate pool, solve a convex optimization to find the FIR feedforward controller G_{ff} that is the best fit to the frequency-response data $\{H(\omega_i)\}_{i=1}^N$.
- 3) Determine the triple (G_{ff}, G_{fb}, d) that minimizes J over all models in the feedback candidate pool.

For details of the SSID method (see [41, Appendix A]). For an analysis of the algorithm's properties (with $d = 0$), see [40]. The SSID algorithm is implemented on a supercomputer using parallel processing such that the relatively large number of convex optimizations (required in Step 2) can be performed in a reasonable amount of time (see [40]).

For each trial, we identify the FIR feedforward controller G_{ff} , second-order strictly proper feedback transfer function G_{fb} , and feedback delay d that minimize J . The feedforward controller order is selected to allow G_{ff} to approximate G^{-1}

TABLE V
MEAN $\|G_{ff}G - 1\|_1$ AND CHANGE FROM FIRST FIVE TO LAST FIVE TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
G_m	0.44	0.25	0.30	0.20	0.15	0.17	−0.27
G_n	1.10	0.62	0.46	0.44	0.33	0.32	−0.78
G_{rd1}	0.27	0.18	0.18	0.15	0.13	0.14	−0.13
G_{rd2}	0.44	0.30	0.25	0.20	0.16	0.14	−0.30
G_{rd3}	0.68	0.50	0.40	0.32	0.28	0.33	−0.35

with approximately 0.1% error over the 0-to-0.5 Hz range of the chirp command r . Thus, the feedforward controller orders are 2, 6, 5, 2, and 3 for G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} , respectively. See Appendix A in the supplemental material for the details of the feedback candidate pool.

The SSID algorithm from [40] and [41] is used as opposed to traditional system identification techniques (e.g., [48]–[52]), because the unknown subsystem (4) is connected in feedback with G , and traditional system identification algorithms applied to closed-loop architectures can yield trivial solutions [53]–[55]. For example, $G_{ff} = G^{-1}$ and $z^{-d}G_{fb} = -G^{-1}$ is a solution with most classical system identification approaches. In contrast, $G_{ff} = G^{-1}$ and $z^{-d}G_{fb} = -G^{-1}$ does not minimize the SSID cost (6) and is not a solution with the SSID algorithm from [40] and [41]. Nevertheless, if the identified feedback controller satisfies $z^{-d}G_{fb} \approx -G^{-1}$, then the SSID results can be ill conditioned. See Appendix B in the supplemental material for more details and a conditioning analysis of the SSID results presented in this article. Appendix B shows that the qualitative results reported in this article are not impacted by ill-conditioned results. In addition, Appendix C in the supplemental material presents a validation analysis of the SSID results.

V. FEEDFORWARD IDENTIFICATION RESULTS

A. Learning to Approximate G^{-1} in Feedforward

The experimental study in [41] demonstrated that an important aspect of humans learning to control a dynamic system is updating the feedforward (i.e., anticipatory) control until it approximates the inverse system dynamics. However, the experimental study in [41] used only one dynamic system (order 4, relative degree 2, minimum phase). Thus, we are motivated to determine if humans learn to approximate the inverse dynamics in feedforward for other systems.

For each identified feedforward controller, we define

$$\|G_{ff}G - 1\|_1 \triangleq \frac{1}{\pi} \int_0^\pi |G_{ff}(e^{j\omega T_s})G(e^{j\omega T_s}) - 1| d\omega$$

where G is equal to G_m , G_n , G_{rd1} , G_{rd2} , or G_{rd3} as appropriate. Thus, $\|G_{ff}G - 1\|_1$ is the magnitude of the difference between the identified G_{ff} and the inverse dynamics G^{-1} normalized by the magnitude of G^{-1} and averaged over the 0-to-0.5 Hz range (i.e., the 0-to- π rad/s range).

Fig. 6 shows the mean and standard deviation of $\|G_{ff}G - 1\|_1$ on each trial for each group, and Table V shows the mean $\|G_{ff}G - 1\|_1$ on different sets of trials for each group. For each group, the mean $\|G_{ff}G - 1\|_1$ decreases over the trials. Figs. 7–11 are the Bode plots of the average identified feedforward controller G_{ff} over all 11 subjects on trials 1 and 40 for the




Fig. 6. Mean and standard deviation of $\|G_{ff}G - 1\|_1$ on each trial. The normalized difference between the identified G_{ff} and G^{-1} tends to decrease over the trials. The $\circ$ is the mean, and the lines indicate standard deviation.




Fig. 7. Average identified feedforward controller for the group with G_m on trials 1 and 40. The shaded region shows the standard deviation.




Fig. 8. Average identified feedforward controller for the group with G_n on trials 1 and 40. The shaded region shows the standard deviation.

group with G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} , respectively. For each group, the average identified G_{ff} approximates G^{-1} better on trial 40 than on trial 1. Thus, by the last trial, subjects learn an approximation of G^{-1} and use the approximate inverse dynamics in feedforward. This result agrees with the results in [41] for a different experiment, and supports the internal model hypothesis [21]–[23].




Fig. 9. Average identified feedforward controller for the group with G_{rd1} on trials 1 and 40. The shaded region shows the standard deviation.




Fig. 10. Average identified feedforward controller for the group with G_{rd2} on trials 1 and 40. The shaded region shows the standard deviation.




Fig. 11. Average identified feedforward controller for the group with G_{rd3} on trials 1 and 40. The shaded region shows the standard deviation.

B. Differences in the Groups' Approximations of G^{-1}

Although the SSID results show that the subjects learn to approximate G^{-1} in feedforward, Fig. 6 shows that on trial 40, the subjects interacting with G_m , G_{rd1} , and G_{rd2} approximate the inverse dynamics more accurately than those interacting with G_n and G_{rd3} . Specifically, mean $\|G_{ff}G - 1\|_1$ on trial 40 for G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} are 0.14, 0.22, 0.14, 0.13, and

0.35, respectively. Table V shows similar results over the last five trials.

A one-way ANOVA comparing each group's mean $\|G_{ff}G - 1\|_1$ over the last five trials yields $F_{4,11} = 7.33$ and $p < 0.001$, thus confirming a statistical difference between groups. A Tukey posthoc pairwise comparison yields $p < 0.05$ for each of G_m , G_{rd1} , or G_{rd2} paired with either G_n or G_{rd3} . Thus, the subjects interacting with G_m , G_{rd1} , and G_{rd2} learn to approximate G^{-1} in feedforward better than those interacting with G_n and G_{rd3} . This difference in the accuracy of the approximation of G^{-1} in feedforward helps to explain why the mean $\|e\|$ over the last five trials for G_m , G_{rd1} , and G_{rd2} are less than those for G_n and G_{rd3} .

One possible explanation for the difference in the accuracy of the approximations of G^{-1} is that the approximation tends to be worse at frequencies where G has significant phase lag, and that G_n and G_{rd3} have the most phase lag over the 0-to-0.5 Hz range. Note that this conjecture is supported qualitatively by the Bode plots in Figs. 7–11.

To examine this conjecture further, we consider a modified version of $\|G_{ff}G - 1\|_1$ that focuses on the frequency ranges over which each transfer function has the same amount of phase lag. First, note that G_{rd1} has the least phase lag of the five transfer functions over the 0-to-0.5 Hz range, and that phase lag is $\theta \triangleq \angle G_{rd1}(e^{j\pi T_s}) = -0.958$ rad (or -54.9°). For each dynamic system, let $\omega_\theta \in [0, \pi]$ be such that

$$\angle G(e^{j\omega_\theta T_s}) = \theta$$

where G is equal to G_m , G_n , G_{rd1} , G_{rd2} , or G_{rd3} as appropriate. Thus, ω_θ is the frequency at which the transfer function G has phase lag θ , which is the phase lag of G_{rd1} at 0.5 Hz. Therefore, ω_θ is equal to π rad/s (or 0.5 Hz) for G_{rd1} , and ω_θ is less than π rad/s for the other 4 transfer functions. Specifically, ω_θ is equal to 0.75π , 0.30π , 0.23π , and 0.23π rad/s for G_m , G_{rd2} , G_n , and G_{rd3} , respectively.

Next, for each identified feedforward controller, define

$$\|G_{ff}G - 1\|_{1,[0,\omega_\theta]} \triangleq \int_0^{\omega_\theta} \frac{|G_{ff}(e^{j\omega T_s})G(e^{j\omega T_s}) - 1|}{\omega_\theta} d\omega$$

where $\omega_\theta \in [0, \pi]$ satisfies $\angle G(e^{j\omega_\theta T_s}) = \theta$, and G is equal to G_m , G_n , G_{rd1} , G_{rd2} , or G_{rd3} as appropriate. Therefore, $\|G_{ff}G - 1\|_{1,[0,\omega_\theta]}$ is computed by averaging over the frequency range from 0 to ω_θ , which is the range where the phase of G is between 0 and θ . In other words, $\|G_{ff}G - 1\|_{1,[0,\omega_\theta]}$ is the magnitude of the difference between G_{ff} and G^{-1} normalized by the magnitude of G^{-1} and averaged over the frequency range where the phase of G is between 0 and θ . Similarly, for each identified feedforward controller, define

$$\|G_{ff}G - 1\|_{1,[\omega_\theta,\pi]} \triangleq \int_{\omega_\theta}^{\pi} \frac{|G_{ff}(e^{j\omega T_s})G(e^{j\omega T_s}) - 1|}{\pi - \omega_\theta} d\omega$$

which is the magnitude of the difference between G_{ff} and G^{-1} normalized by the magnitude of G^{-1} and averaged over the frequency range where the phase of G is between θ and the value it achieves at π rad/s.

Table VI shows the mean $\|G_{ff}G - 1\|_{1,[0,\omega_\theta]}$ on different sets of trials for each group, and Table VII shows the mean $\|G_{ff}G - 1\|_{1,[\omega_\theta,\pi]}$ on different sets of trials for each group.

TABLE VI
MEAN $\|G_{ff}G - 1\|_{1,[0,\omega_\theta]}$ AND CHANGE FROM FIRST FIVE TO LAST FIVE TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
G_m	0.36	0.20	0.22	0.16	0.12	0.14	−0.22
G_n	0.46	0.24	0.22	0.17	0.18	0.13	−0.33
G_{rd1}	0.27	0.18	0.18	0.15	0.13	0.14	−0.13
G_{rd2}	0.30	0.23	0.16	0.15	0.13	0.11	−0.19
G_{rd3}	0.31	0.18	0.19	0.14	0.10	0.13	−0.18

TABLE VII
MEAN $\|G_{ff}G - 1\|_{1,[\omega_\theta,\pi]}$ AND CHANGE FROM FIRST FIVE TO LAST FIVE TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
G_n	1.30	0.74	0.54	0.52	0.38	0.38	−0.92
G_{rd3}	0.79	0.59	0.47	0.38	0.33	0.39	−0.40

$\|G_{ff}G - 1\|_{1,[\omega_\theta,\pi]}$ on different sets of trials for the groups with G_n and G_{rd3} . Table VI shows that over the last five trials, the error in approximating G^{-1} in feedforward is approximately the same for all groups over the 0-to- ω_θ frequency range. In contrast, Table VII shows that over the last five trials, the groups with G_n and G_{rd3} have significantly greater error in approximating G^{-1} in feedforward over the higher frequency range. Thus, the groups with G_n and G_{rd3} learn to approximate G^{-1} in feedforward more poorly than the other groups at higher frequencies, where G_n and G_{rd3} have significantly more phase lag. This provides additional evidence that phase lag is an important impediment to a human's ability to learn to approximate the inverse dynamics in feedforward.

C. Learning the Phase Versus Magnitude of G^{-1}

Figs. 7–11 show that for each group, the average identified feedforward controller on trial 1 has significant phase lag relative to G^{-1} , and that this phase lag is significantly reduced (or eliminated) by trial 40. Thus, subjects learn to use phase lead in feedforward over the trials. This observation suggests that a key component of learning to approximate G^{-1} in feedforward is learning to use the correct amount of phase lead. To examine this observation in more detail, for each trial, define the *frequency-averaged error in the magnitude of $G_{ff}G$ relative to unity*

$$M_e(G_{ff}) \triangleq \frac{1}{\pi} \int_0^\pi \left| |G_{ff}(e^{j\omega T_s})G(e^{j\omega T_s})| - 1 \right| d\omega$$

which is the frequency-averaged magnitude of the difference between $G_{ff}G$ and 1 assuming that the phase of $G_{ff}G$ is equal to the phase of 1. Similarly, for each trial, define the *frequency-averaged error in the phase of $G_{ff}G$ relative to unity*

$$P_e(G_{ff}) \triangleq \frac{1}{\pi} \int_0^\pi \left| e^{j\angle[G_{ff}(e^{j\omega T_s})G(e^{j\omega T_s})]} - 1 \right| d\omega$$

which is the frequency-averaged magnitude of the difference between $G_{ff}G$ and 1 assuming that the magnitude of $G_{ff}G$ is equal to the magnitude of 1. Now, we compare M_e and P_e to determine if difference between $G_{ff}G$ and 1 is due more to




Fig. 12. Mean and standard deviation of M_e on each trial. The mean M_e tends to decrease over the trials. The $\circ$ is the mean, and the lines indicate standard deviation.




Fig. 13. Mean and standard deviation of P_e on each trial. The mean P_e tends to decrease over the trials. The $\circ$ is the mean, and the lines indicate standard deviation.

TABLE VIII
MEAN M_e AND CHANGE FROM THE FIRST FIVE TO LAST FIVE TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
G_m	0.24	0.14	0.16	0.12	0.08	0.10	−0.14
G_n	0.54	0.32	0.27	0.25	0.19	0.19	−0.35
G_{rd1}	0.15	0.13	0.12	0.09	0.09	0.11	−0.04
G_{rd2}	0.23	0.18	0.14	0.13	0.09	0.09	−0.14
G_{rd3}	0.35	0.26	0.22	0.18	0.15	0.19	−0.16

TABLE IX
MEAN P_e AND CHANGE FROM FIRST FIVE TO LAST FIVE TRIALS

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
G_m	0.35	0.19	0.23	0.15	0.12	0.13	−0.22
G_n	0.84	0.48	0.30	0.30	0.23	0.22	−0.62
G_{rd1}	0.19	0.11	0.09	0.09	0.08	0.07	−0.12
G_{rd2}	0.42	0.24	0.19	0.13	0.11	0.10	−0.32
G_{rd3}	0.72	0.47	0.35	0.26	0.23	0.27	−0.45

error in magnitude or error in phase. In other words, we use M_e and P_e to determine if the error in approximating the inverse dynamics in feedforward is attributed more to error in matching the magnitude or the phase of the inverse dynamics.

Figs. 12 and 13 show the mean and standard deviation of M_e and P_e on each trial for each group. For each group, the mean M_e and mean P_e tend to decrease over the trials. Tables VIII and IX show the mean M_e and P_e on different sets of trials for each group. A paired t -test of change in M_e and change in P_e




Fig. 14. Mean and standard deviation of $\|z^{-d}G_{fb}\|_1$ on each trial. The mean $\|z^{-d}G_{fb}\|_1$ tends to increase over the first 20 trials. The $\circ$ is the mean, and the lines indicate standard deviation.

TABLE X
MEAN $\|z^{-d}G_{fb}\|_1$

	Trials 1–5	Trials 6–10	Trials 11–15	Trials 16–20	Trials 21–25	Trials 26–30	Trials 31–40
G_m	1.71	1.73	2.76	2.74	2.42	2.03	2.17
G_n	1.23	1.25	0.96	1.27	1.31	1.33	1.29
G_{rd1}	1.28	1.36	1.42	1.43	1.48	1.57	1.55
G_{rd2}	1.35	1.75	1.80	2.79	1.88	2.63	2.20
G_{rd3}	0.68	1.01	0.91	1.14	0.97	1.29	1.15

from the first five trials to the last five trials yields $p < 0.05$ for each group except G_m , which has $p = 0.13$. Thus, the mean P_e decreases more (in absolute and percent) than the mean M_e from the first five trials to the last five trials. Specifically, for G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} , the mean M_e decreases by 58%, 65%, 27%, 61%, and 46% from the first five trials to the last five trials. In contrast, the mean P_e decreases by 63%, 74%, 63%, 76%, and 63% in the same order. This observation suggests that improvement in the approximation of G^{-1} in feedforward is attributed more to improvement in matching the phase of G^{-1} than improvement in matching the magnitude of G^{-1} . Thus, learning the phase lead of G^{-1} is a critical aspect of learning to approximate G^{-1} in feedforward.

VI. FEEDBACK IDENTIFICATION RESULTS

For each identified feedback controller, we define

$$\|z^{-d}G_{fb}\|_1 \triangleq \frac{1}{\pi} \int_0^\pi |e^{-j\omega T_s d} G_{fb}(e^{j\omega T_s})| d\omega$$

which is the frequency-averaged magnitude of $z^{-d}G_{fb}$ over the 0-to-0.5 Hz range. Note that $\|z^{-d}G_{fb}\|_1$ does not depend on the feedback delay d . Fig. 14 shows the mean and standard deviation of $\|z^{-d}G_{fb}\|_1$ on each trial for each group, and Table X shows the mean $\|z^{-d}G_{fb}\|_1$ on different sets of trials for each group. For each group, the mean $\|z^{-d}G_{fb}\|_1$ tends to increase from trials 1–5 to trials 16–20 (although the increase for the group with G_n is comparatively small). These increases suggest that the subjects learn to use more frequency-averaged feedback gain over the first 20 trials. Thus, subjects learn to increase the feedback gain in a manner that maintains closed-loop stability. To examine this further, we compute the stability margins (i.e., upward gain margin and phase margin) associated with each identified feedback controller.

TABLE XI
MEAN UPWARD GAIN MARGIN (ABSOLUTE)

	Trials 1–5	Trials 6–10	Trials 11–15	Trials 16–20	Trials 21–25	Trials 26–30	Trials 31–40
G_m	3.26	1.98	1.92	1.87	2.19	2.04	1.90
G_n	1.59	1.66	1.39	1.46	1.42	1.32	1.62
G_{rd1}	1.88	1.83	1.48	1.71	1.49	1.55	1.69
G_{rd2}	1.73	1.87	1.53	1.25	1.32	1.28	1.25
G_{rd3}	1.44	1.18	1.24	1.22	1.22	1.19	1.25

TABLE XII
MEAN PHASE MARGIN (DEGREES)

	Trials 1–5	Trials 6–10	Trials 11–15	Trials 16–20	Trials 21–25	Trials 26–30	Trials 31–40
G_m	47	32	14	19	26	21	17
G_n	21	17	25	10	21	13	14
G_{rd1}	19	13	12	14	10	12	15
G_{rd2}	36	21	17	18	17	11	13
G_{rd3}	28	11	11	10	11	11	12

Tables XI and XII show the mean upward gain margins and mean phase margins on different sets of trials for each group. For each group, the mean upward gain margin and mean phase margin decrease from trials 1–5 to trials 16–20. Thus, over the first 20 trials, subjects learn to increase feedback gain by using controllers with smaller stability margins. This observation suggests that through repeated interaction, the subjects learn the system dynamics G well enough to use less conservative feedback controllers without causing closed-loop instability.

Note that the mean $\|z^{-d}G_{fb}\|_1$ and associated stability margins do not have a consistent trend over the last 20 trials. However, during the last 20 trials, the feedforward control is a comparatively accurate approximation of G^{-1} , and (5) implies that the closed-loop transfer function $\tilde{G}_{yr}$ is insensitive to $z^{-d}G_{fb}$. Thus, for these trials, the identified G_{fb} may not be an accurate representation of the feedback used by the subjects.

The average identified feedback time delay over all 40 trials for groups with G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} are 233, 299, 310, 241, and 310 ms, respectively. These results for human time delay with visual feedback are consistent with [41]. There is no apparent trend in the mean or standard deviation of the identified feedback time delay.

VII. CONCLUSION

This article contains several new contributions. First, this article determines characteristics that can make an LTI dynamic system difficult for humans to control. Second and more importantly, this paper provides insights into why those characteristics can make a system difficult for humans to control. Note that many existing studies (e.g., [24]–[38]) focus on the human's steady-state control behavior (i.e., the behavior after practice). In contrast, this work examines not only the steady-state behavior but also how the subjects' control behavior changes with practice (i.e., from the early to later trials).

This article shows that higher relative degree and nonminimum-phase zeros tend to make dynamic systems more difficult for humans to control, whereas higher system order (without change in relative degree) does not necessarily make

a system more difficult to control (see Fig. 3 and Table II). Together, these results demonstrate that phase lag is a key characteristic that makes systems difficult for humans to control. In addition, the frequency-domain analysis in Section III-B shows that the subjects' improvement in time-averaged error $\|e\|$ over the trials is attributed more to improvement in matching the phase of the reference than improvement in matching the magnitude of the reference.

The feedforward identification results in Section V show that approximating the inverse dynamics in feedforward is an important component of learning to control each of the five dynamic systems. This result agrees with previously reported results for different experiments (e.g., [41]). The results also show that subjects interacting with systems that have less phase lag (i.e., G_m , G_{rd1} , G_{rd2}) learn to approximate the inverse dynamics in feedforward more accurately than those interacting with systems that have more phase lag (G_n , G_{rd3}). This observation helps to explain why the subjects' interacting with G_m , G_{rd1} , and G_{rd2} have smaller command-following error than those interacting with G_n and G_{rd3} .

The feedforward identification results provide new insights into why system phase lag can make an LTI dynamic system difficult for humans to control. For each group, the average identified feedforward controller on the first trial has phase lag relative to the inverse dynamics, and this phase lag is significantly reduced or eliminated by the last trial (see Figs. 7–11). Furthermore, for each group, the subjects' improvement in approximating G^{-1} in feedforward over the trials is attributed more to improvement in matching the phase of G^{-1} than improvement in matching the magnitude of G^{-1} (see Tables VIII and IX). Thus, a key aspect of learning to approximate G^{-1} in feedforward is learning the phase of G^{-1} and learning to use the correct amount of phase lead. However, results also demonstrate that larger system phase lag is an important impediment to a subject's ability to approximate G^{-1} in feedforward (see Tables VI and VII).

The feedback identification results show that the frequency-averaged feedback gains tend to increase over the first 20 trials, and these higher feedback gains correspond to smaller stability margins. Thus, learning to use less-conservative feedback controllers is another factor that contributes to the improvement in $\|e\|$ over trials.

The results of this article provide new insights into the impact of system characteristics (e.g., relative degree, system order, zeros, phase lag) on HITL control behavior. These results could have application to design and analysis for a variety of HITL technologies, including active prostheses and exoskeletons, and human-operated devices and vehicles.

REFERENCES

- [1] J. B. Hoagg and D. S. Bernstein, "Nonminimum-phase zeros: Much to do about nothing," *IEEE Control Syst. Mag.*, vol. 27, no. 3, pp. 45–57, Jun. 2007.
- [2] M. M. Seron, J. H. Braslavsky, and G. C. Goodwin, *Fundamental Limitations in Filtering and Control*. Berlin, Germany: Springer-Verlag, 1997.
- [3] A. J. Bastian, "Learning to predict the future: The cerebellum adapts feed-forward movement control," *Curr. Opin. Neurobiol.*, vol. 16, pp. 645–649, 2006.

- [4] M. Mulder *et al.*, "Manual control cybernetics: State-of-the-art and current trends," *IEEE Trans. Human-Mach. Syst.*, vol. 48, no. 5, pp. 468–485, Oct. 2018.
- [5] R. Happee, "Goal directed arm movements. III. feedback and adaptation in response to inertia perturbations," *J. Electromyography Kinesiology*, vol. 3, no. 2, pp. 112–122, 1993.
- [6] R. Shadmehr and F. A. Mussa-Ivaldi, "Adaptive representation of dynamics during learning of a motor task," *J. Neurosci.*, vol. 14, no. 5, pp. 3208–3224, 1994.
- [7] F. Gandolfo, F. A. Mussa-Ivaldi, and E. Bizzi, "Motor learning by field approximation," *Proc. Nat. Acad. Sci.*, vol. 93, no. 9, pp. 3843–3846, 1996.
- [8] R. L. Sainburg, C. Ghez, and D. Kalakanis, "Intersegmental dynamics are controlled by sequential anticipatory, error correction, and postural mechanisms," *J. Neurophysiol.*, vol. 81, no. 3, pp. 1045–1056, 1999.
- [9] J. R. Flanagan and A. M. Wing, "The role of internal models in motion planning and control: Evidence from grip force adjustments during movements of hand-held loads," *J. Neurosci.*, vol. 17, no. 4, pp. 1519–1528, 1997.
- [10] T. Flash and N. Hogan, "The coordination of arm movements: An experimentally confirmed mathematical model," *J. Neurosci.*, vol. 5, no. 7, pp. 1688–1703, 1985.
- [11] M. Kawato, K. Furukawa, and R. Suzuki, "A hierarchical neural-network model for control and learning of voluntary movement," *Biol. Cybern.*, vol. 57, no. 3, pp. 169–185, 1987.
- [12] R. C. Miall, D. J. Weir, D. M. Wolpert, and J. F. Stein, "Is the cerebellum a smith predictor?," *J. Motor Behav.*, vol. 25, no. 3, pp. 203–216, 1993.
- [13] M. Katayama and M. Kawato, "Virtual trajectory and stiffness ellipse during multi-joint arm movement predicted by neural inverse models," *Biol. Cybern.*, vol. 69, nos. 5/6, pp. 353–362, 1993.
- [14] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan, "An internal model for sensorimotor integration," *Sci.*, vol. 269, pp. 1880–1882, 1995.
- [15] E. Guigon, P. Baraduc, and M. Desmurget, "Computational motor control: Redundancy and invariance," *J. Neurophysiol.*, vol. 97, no. 1, pp. 331–347, 2007.
- [16] N. Bhushan and R. Shadmehr, "Computational nature of human adaptive control during learning of reaching movements in force fields," *Biol. Cybern.*, vol. 81, no. 1, pp. 39–60, 1999.
- [17] J.-L. Vercher, F. Sares, J. Blouin, C. Bourdin, and G. M. Gauthier, "Role of sensory information in updating internal models of the effector during arm tracking," *Prog. Brain Res.*, vol. 142, pp. 203–222, 2003.
- [18] J. B. Dingwell, C. D. Mah, and F. A. Mussa-Ivaldi, "Experimentally confirmed mathematical model for human control of a non-rigid object," *J. Neurophysiol.*, vol. 91, no. 3, pp. 1158–1170, 2004.
- [19] D. Liu and E. Todorov, "Evidence for the flexible sensorimotor strategies predicted by optimal feedback control," *J. Neurosci.*, vol. 27, no. 35, pp. 9354–9368, 2007.
- [20] A. J. Nagengast, D. A. Braun, and D. M. Wolpert, "Optimal control predicts human performance on objects with internal degrees of freedom," *PLOS Comput. Biol.*, vol. 5, 2009, Paper e1000419.
- [21] R. Shadmehr, M. A. Smith, and J. W. Krakauer, "Error correction, sensory prediction, and adaptation in motor control," *Annu. Rev. Neurosci.*, vol. 33, pp. 89–108, 2010.
- [22] M. Kawato, "Internal models for motor control and trajectory planning," *Curr. Opin. Neurobiol.*, vol. 9, no. 6, pp. 718–727, 1999.
- [23] D. M. Wolpert, R. C. Miall, and M. Kawato, "Internal models in the cerebellum," *Trends Cogn. Sci.*, vol. 2, no. 9, pp. 338–347, 1998.
- [24] D. T. McRuer and E. S. Krendel, "The human control operator as a servo system element: Part I," *J. Franklin Inst.*, vol. 5, no. 267, pp. 381–403, 1959.
- [25] D. T. McRuer and E. S. Krendel, "The human control operator as a servo system element: Part II," *J. Franklin Inst.*, vol. 6, no. 267, pp. 511–536, 1959.
- [26] D. T. McRuer, D. Graham, and E. S. Krendel, "Manual control of a single-loop system: Part I," *J. Franklin Inst.*, vol. 283, no. 1, pp. 1–29, 1967.
- [27] D. T. McRuer, D. Graham, and E. S. Krendel, "Manual control of a single-loop system: Part II," *J. Franklin Inst.*, vol. 283, no. 2, pp. 145–168, 1967.
- [28] R. L. Stapleford, D. T. McRuer, and R. E. Magdaleno, "Pilot describing function measurements in a multiloop task," *IEEE Trans. Hum. Factors Electron.*, vol. 8, no. 2, pp. 113–125, 1967.
- [29] D. T. McRuer and D. H. Weir, "Theory of manual vehicle control," *IEEE Trans. Human-Mach. Syst.*, vol. 10, no. 4, pp. 257–291, Dec. 1969.
- [30] R. J. Wasicko, D. T. McRuer, and R. E. Magdaleno, "Human pilot dynamics in single-loop systems with compensatory and pursuit displays," Air Force Flight Dynamics Laboratory, Research and Technology Division, Air Force Systems Command, Wright-Patterson Air Force Base, OH, USA, Tech. Rep. AFFDL-TR-66-137, 1966.
- [31] F. M. Drop, D. M. Pool, H. J. Damveld, M. M. van Paassen, and M. Mulder, "Identification of the feedforward component in manual control with predictable target signals," *IEEE Trans. Cybern.*, vol. 43, no. 6, pp. 1936–1949, Dec. 2013.
- [32] K. van der El, D. M. Pool, M. R. M. van Paassen, and M. Mulder, "Effects of preview on human control behavior in tracking tasks with various controlled elements," *IEEE Trans. Cybern.*, vol. 48, no. 4, pp. 1242–1252, Apr. 2018.
- [33] K. van der El, D. M. Pool, H. J. Damveld, M. R. M. van Paassen, and M. Mulder, "An empirical human controller model for preview tracking tasks," *IEEE Trans. Cybern.*, vol. 46, no. 11, pp. 2609–2621, Nov. 2016.
- [34] K. van der El, D. M. Pool, M. M. van Paassen, and M. Mulder, "Effects of linear perspective on human use of preview in manual control," *IEEE Trans. Human-Mach. Syst.*, vol. 48, no. 5, pp. 496–508, Oct. 2018.
- [35] K. van der El, S. Padmos, D. M. Pool, M. M. van Paassen, and M. Mulder, "Effects of preview time in manual tracking tasks," *IEEE Trans. Human-Mach. Syst.*, vol. 48, no. 5, pp. 486–495, Oct. 2018.
- [36] V. A. Laurence, D. M. Pool, H. J. Damveld, M. M. van Paassen, and M. Mulder, "Effect of controlled element dynamics on human feedforward behavior in ramp-tracking tasks," *IEEE Trans. Cybern.*, vol. 45, no. 2, pp. 253–265, Feb. 2015.
- [37] F. M. Drop, D. M. Pool, M. M. van Paassen, M. Mulder, and H. H. Bulthoff, "Effects of target signal shape and system dynamics on feedforward in manual control," *IEEE Trans. Cybern.*, vol. 49, no. 3, pp. 768–780, Mar. 2019.
- [38] S. Barendsward, D. M. Pool, M. M. V. Paassen, and M. Mulder, "Dual-axis manual control: Performance degradation, axis asymmetry, crossfeed, and intermittency," *IEEE Trans. Human-Mach. Syst.*, vol. 49, no. 2, pp. 113–125, Apr. 2019.
- [39] D. M. Pool, G. A. Harder, and M. M. van Paassen, "Effects of simulator motion feedback on training of skill-based control behavior," *J. Guid. Control Dyn.*, vol. 39, no. 4, pp. 889–901, 2016.
- [40] X. Zhang and J. B. Hoagg, "Subsystem identification of multivariable feedback and feedforward systems," *Automatica*, vol. 72, pp. 131–137, 2016.
- [41] X. Zhang, S. Wang, J. B. Hoagg, and T. M. Seigler, "The roles of feedback and feedforward as humans learn to control unknown dynamic systems," *IEEE Trans. Cybern.*, vol. 48, no. 2, pp. 543–555, Feb. 2018.
- [42] J. B. Hoagg, J. Chandrasekar, and D. S. Bernstein, "On the zeros, initial undershoot, and relative degree of lumped-mass structures," *J. Dyn. Syst. Meas. Control*, vol. 129 no. 4, pp. 493–502, 2007.
- [43] J. Yan, J. B. Hoagg, R. E. Hindman, and D. S. Bernstein, "Longitudinal aircraft dynamics and the instantaneous acceleration center of rotation," *IEEE Control Syst. Mag.*, vol. 31, no. 4, pp. 68–92, Aug. 2011.
- [44] R. J. Jagacinski and J. Flach, *Control Theory for Humans: Quantitative Approaches to Modeling Performance*. Mahwah, NJ: Erlbaum Associates, 2003.
- [45] B. Yu, R. B. Gillespie, J. S. Freudenberger, and J. A. Cook, "Human control strategies in pursuit tracking with a disturbance input," in *Proc. Conf. Dec. Control*, Dec. 2014, pp. 3795–3800.
- [46] D. T. McRuer and H. R. Jex, "A review of quasi-linear pilot models," *IEEE Trans. Hum. Factors Electron.*, vol. HFE-8, no. 3, pp. 231–249, Sep. 1967.
- [47] F. M. Drop, R. de Vries, M. Mulder, and H. H. Bulthoff, "The predictability of a target signal affects manual feedforward control," in *Proc. 13th IFAC Symp. Anal., Des., Eval. Human-Mach. Syst.*, Kyoto, Japan, Aug. 2016, pp. 177–182.
- [48] T. Soderstrom and P. Stoica, *System Identification*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1989.
- [49] J. N. Juang, *Applied System Identification*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1993.
- [50] P. V. Overschee and B. D. Moor, "A unifying theorem for three subspace system identification algorithms," *Automatica*, vol. 31, no. 12, pp. 1853–1864, 1995.
- [51] L. Ljung, *System Identification: Theory for the User*, 2nd ed. Englewood Cliffs, NJ, USA: Prentice-Hall, 1999.
- [52] R. Pintelon and J. Schoukens, *System Identification: A Frequency Domain Approach*. Piscataway, NJ, USA: IEEE, 2001.
- [53] L. L. I. Gustavsson and T. Soderstrom, "Identification of processes in closed loop-identifiability and accuracy aspects," *Automatica*, vol. 13, no. 1, pp. 59–75, 1977.
- [54] U. Forssell and L. Ljung, "Closed-loop identification revisited," *Automatica*, vol. 35, pp. 1215–1241, 1999.
- [55] R. Isermann and M. Münchhof, *Identification of Dynamic Systems: An Introduction With Applications*. Berlin, Germany: Springer, 2011.

APPENDIX A DESCRIPTION OF CANDIDATE POOL FOR SSID

For each trial, we use the SSID algorithm in [41, Algorithm 1] to identify the best-fit model of the subject's control (4). The controller orders are chosen sufficiently large to capture different control approaches that lead to good command-following performance. We select the controller orders to allow for high gain in feedback as well as approximate dynamic inversion in feedforward. Specifically, for all groups, G_{fb} is modeled as a second-order strictly proper transfer function. We select the feedforward controller order large enough to allow G_{ff} to approximate G^{-1} with approximately 0.1% error over the 0-to-0.5 Hz range. For G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} , the FIR feedforward transfer function is second order, sixth order, fifth order, second order, and third order, respectively.

The candidate pool (denoted by Φ in [41, Algorithm 1]) is designed to capture a wide range of behavior over the 0-to-0.5 Hz range. The candidate pools for G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} contain approximately 0.9, 0.7, 0.8, 1.0, and 0.4 billion elements, respectively. The candidate pool satisfies the following conditions:

- C1) If $\lambda \in \mathbb{C}$ is a pole of G_{fb} , then $|(\ln \lambda)/T_s| \leq 31.5$.
- C2) If $\lambda \in \mathbb{C}$ is a zero of G_{fb} , then $|(\ln \lambda)/T_s| \leq 31.5$.
- C3) $\max_{\omega \in [0, \pi]} |G_{fb}(e^{j\omega T_s})| \leq 30.5$.
- C4) If $\lambda \in \mathbb{C}$ is a pole of $\tilde{G}_{yr}$, then $|\lambda| < 0.998$.
- C5) $d \in \{4, 5, 6, \dots, 25\}$.

Conditions C1) and C2) constrain Φ to include only elements that have a significant impact on controller dynamics over the 0-to-0.5 Hz range. Specifically, C1) and C2) state that G_{fb} has continuous-time equivalent poles and zeros (that is, poles and zeros obtained from the matched z -transform mapping $s = (\ln z)/T_s$) that have magnitudes between 0 and 31.5 rad/s. This condition arises because $\{H(\omega_k)\}_{k=1}^N$ is at frequencies $\omega_1, \dots, \omega_N \in [0, \pi]$ rad/s, which corresponds to the frequency range of r . Thus, we seek to identify G_{fb} on the interval $[0, \pi]$ rad/s. The upper limit 31.5 rad/s on the magnitude of the continuous-time equivalent poles and zeros is one decade above the π rad/s limit on the chirp frequency. A continuous-time pole or zero with magnitude greater than 31.5 rad/s has negligible effect on the Bode plot over the frequency range $[0, \pi]$ rad/s. Thus, we restrict the candidate pool to elements that correspond to continuous-time equivalent poles and zeros with magnitude between 0 and 31.5 rad/s.

Condition C3) states that the peak magnitude of G_{fb} over $[0, \pi]$ rad/s is no more than 30.5. We impose an upper limit on the magnitude of G_{fb} because a human cannot use arbitrarily high gain in feedback. See [41] for a description of the experiment used to determine the 30.5 upper limit.

Condition C4) states that each closed-loop pole has magnitude less than 0.998. A discrete-time pole with magnitude 0.998 and sample time $T_s = 0.02$ s has a settling time of approximately 40 s. Thus, C4) restricts Φ to include only elements that result in closed-loop transfer functions with settling times less than 40 s. The behavior observed in this experiment exhibits settling times significantly less than 40 s.

Condition C5) restricts the feedback time delay to the range of [80, 500] ms.

The SSID algorithm is coded in C++ for parallel computation and implemented on the High Performance Computing Cluster at the University of Kentucky. For each trial, it takes approximately 0.3 h to run [41, Algorithm 1] on one compute node; each node has 16 Intel E5-2670 @ 2.6 GHz cores.

APPENDIX B CONDITIONING OF SSID RESULTS

We use the SSID algorithm from [40], [41] as opposed to classical system identification techniques, because the unknown subsystem (4) is connected in feedback with G , and classical system identification algorithms applied to closed-loop architectures can yield trivial solutions [55]. For example, $G_{ff} = G^{-1}$ and $z^{-d}G_{fb} = -G^{-1}$ is a solution with most classical system identification approaches. In contrast, $G_{ff} = G^{-1}$ and $z^{-d}G_{fb} = -G^{-1}$ does not minimize the SSID cost (6) and is not a solution with the SSID algorithm [41, Appendix A]. However, if the identified feedback controller satisfies $z^{-d}G_{fb} \approx -G^{-1}$, then the denominator of (5) is approximately zero. In this case, the SSID results are ill conditioned. The feedback candidate pool is selected in Step 1 to prevent $z^{-d}G_{fb} = -G^{-1}$. Appendix A in the Supplemental Material provides details on the SSID candidate pools used in this paper. Although the feedback candidate pool prevents $z^{-d}G_{fb} = -G^{-1}$, the SSID algorithm could yield ill-conditioned results if $z^{-d}G_{fb} \approx -G^{-1}$ at some frequencies. This appendix examines the conditioning of the SSID results presented in this paper and shows that the qualitative results reported in this paper are not impacted by ill-conditioned results.

Define the condition number

$$\mathcal{C}(G_{fb}, d) \triangleq \frac{1}{\pi} \int_0^\pi \left| \frac{1}{1 + e^{-j\omega T_s d} G_{fb}(e^{j\omega T_s}) G(e^{j\omega T_s})} \right| d\omega,$$

where G is G_m , G_n , G_{rd1} , G_{rd2} , or G_{rd3} as appropriate. Note that $\mathcal{C}(G_{fb}, d)$ is a measure of the conditioning of the identified feedback controller $z^{-d}G_{fb}$. A larger value of $\mathcal{C}(G_{fb}, d)$ indicates a more poorly conditioned identified $z^{-d}G_{fb}$.

Figures 15–19 show $\mathcal{C}$ for each SSID result for the groups with G_m , G_n , G_{rd1} , G_{rd2} , and G_{rd3} , respectively. The SSID results are organized from the largest to smallest $\mathcal{C}$. For group with G_m , the largest $\mathcal{C}$ is approximately 18.7, whereas the smallest $\mathcal{C}$ is approximately 0.106. For group with G_n , the largest $\mathcal{C}$ is approximately 14.2, whereas the smallest $\mathcal{C}$ is approximately 0.986. For group with G_{rd1} , the largest $\mathcal{C}$ is approximately 12.8, whereas the smallest $\mathcal{C}$ is approximately 0.214. For group with G_{rd2} , the largest $\mathcal{C}$ is approximately 32.0, whereas the smallest $\mathcal{C}$ is approximately 0.066. For group with G_{rd3} , the largest $\mathcal{C}$ is approximately 10.3, whereas the smallest $\mathcal{C}$ is approximately 0.205. Figures 15–19 also show that $\mathcal{C}$ is not closely correlated with trial number.

Figures 20–24 show the Bode plot of $z^{-d}G_{fb}$ for the SSID results from each group with the largest $\mathcal{C}$. For the group with G_m , this SSID result has a large $\mathcal{C}$ because $z^{-d}G_{fb} \approx -G_m^{-1}$ from 0.3 to 0.45 Hz. However, $z^{-d}G_{fb} \not\approx -G_m^{-1}$ below 0.3 Hz. For the group with G_n , this SSID result has a large $\mathcal{C}$ because $z^{-d}G_{fb} \approx -G_n^{-1}$ from 0.2 to 0.3 Hz.

However, $z^{-d}G_{fb} \not\approx -G_n^{-1}$ below 0.2 Hz and above 0.3 Hz. For the group with G_{rd1} , this SSID result has a large $\mathcal{C}$ because $z^{-d}G_{fb} \approx -G_{rd1}^{-1}$ below 0.1 Hz. However, $z^{-d}G_{fb} \not\approx -G_{rd1}^{-1}$ above 0.1 Hz. For the group with G_{rd2} , this SSID result has a large $\mathcal{C}$ because $z^{-d}G_{fb} \approx -G_{rd2}^{-1}$ below 0.15 Hz. However, $z^{-d}G_{fb} \not\approx -G_{rd2}^{-1}$ above 0.15 Hz. For the group with G_{rd3} , this SSID result has a large $\mathcal{C}$ because $z^{-d}G_{fb} \approx -G_{rd3}^{-1}$ below 0.1 Hz. However, $z^{-d}G_{fb} \not\approx -G_{rd3}^{-1}$ above 0.1 Hz.

Figures 20–24 also show the Bode plots of $z^{-d}G_{fb}$ for the SSID results from each group with the 111th largest $\mathcal{C}$ (i.e., 25th percentile). For the groups with G_m , G_n , and G_{rd3} , the SSID results with the 111th largest $\mathcal{C}$ are such that $z^{-d}G_{fb}$ is approximately equal to $-G^{-1}$ at 0 Hz but not approximately equal to $-G^{-1}$ at other frequencies over 0-to-0.5 Hz range. For the group with G_{rd1} , the SSID result with the 111th largest $\mathcal{C}$ is such that $z^{-d}G_{fb}$ is approximately equal to $-G_{rd1}^{-1}$ at 0.4 Hz but not approximately equal to $-G_{rd1}^{-1}$ at other frequencies over 0-to-0.5 Hz range. For the group with G_{rd2} , the SSID result with the 111th largest $\mathcal{C}$ is such that $z^{-d}G_{fb}$ is approximately equal to $-G_{rd2}^{-1}$ at 0.5 Hz but not approximately equal to $-G_{rd2}^{-1}$ at other frequencies over 0-to-0.5 Hz range.

We examine the sensitivity of the SSID results in Sections V and VI to the conditioning of $z^{-d}G_{fb}$ by removing the most ill-conditioned 25% of the SSID results. Figures 25 and 26




Fig. 15. Condition number $\mathcal{C}$ for the group with G_m .




Fig. 16. Condition number $\mathcal{C}$ for the group with G_n .




Fig. 17. Condition number $\mathcal{C}$ for the group with G_{rd1} .




Fig. 18. Condition number $\mathcal{C}$ for the group with G_{rd2} .




Fig. 19. Condition number $\mathcal{C}$ for the group with G_{rd3} .

show the mean and standard deviation of $\|G_{ff}G - 1\|_1$ and $\|z^{-d}G_{fb}\|_1$ on each trial for each group. The trends observed in Figs. 25 and 26 match those observed in Figs. 6 and 14. The same trends hold if the most ill-conditioned 10%, 15%, 20%, 25%, or 30% of the results are omitted.




Fig. 20. Bode plot of $z^{-d}G_{fb}$ for the SSID results of the group with G_m with the largest condition number and the 111th largest number.




Fig. 21. Bode plot of $z^{-d}G_{fb}$ for the SSID results of the group with G_n with the largest condition number and the 111th largest number.




Fig. 22. Bode plot of $z^{-d}G_{fb}$ for the SSID results of the group with G_{rd1} with the largest condition number and the 111th largest number.

APPENDIX C VALIDATION OF SSID RESULTS

For each trial, we obtain the validation data $\{y_{v,k}\}_{k=1}^n$ by simulating $\hat{y}_v(z) = \hat{G}_{yr}(z)\hat{r}(z)$, where all initial conditions are zero; $\hat{G}_{yr}$ is the closed-loop transfer function (5) obtained from the identified G_H , G_{fb} , and d ; and $\hat{y}_v(z)$ is the z -




Fig. 23. Bode plot of $z^{-d}G_{fb}$ for the SSID results of the group with G_{rd2} with the largest condition number and the 111th largest number.




Fig. 24. Bode plot of $z^{-d}G_{fb}$ for the SSID results of the group with G_{rd3} with the largest condition number and the 111th largest number.

transform of $y_{v,k}$. For each trial, we compute the variance accounted for (VAF), which is a measure of the accuracy of the identified closed-loop transfer function and is given by

$$\text{VAF} \triangleq 1 - \frac{\sum_{k=n_1}^n |y_k - y_{v,k}|^2}{\sum_{k=n_1}^n |y_k|^2},$$

where $n_1 = 26$. The initial 0.5 s of data are omitted to reduce the impact of nonzero initial conditions because the experimental data may have nonzero initial conditions, whereas the validation data is computed with zero initial conditions.

Figure 27 shows the mean and standard deviation of the VAF for each trial. For each group, the mean VAF increases over the trials, which implies that as the subjects learn, their control behavior is more accurately modeled by the relatively low-order LTI controller (4) used in this paper. This observation suggests that the subjects' control behavior tends to become more linear as they learn, which is consistent with the results in [32], [41].

Table XIII shows the mean VAF on different sets of trials for each group. A one-way ANOVA comparing each group's mean VAF over the last 5 trials yields $F_{4,11} = 16.6$ and $p < 0.001$, thus confirming a statistical difference between groups. A Tukey post-hoc pairwise comparison yields $p < 0.05$ for




Fig. 25. Mean and standard deviation of $\|G_{\text{ff}}G - 1\|_1$ on each trial. Plots omit the most ill-conditioned 25% of trials. The $\circ$ is the mean, and the lines indicate standard deviation.




Fig. 26. Mean and standard deviation of $\|z^{-d}G_{\text{fb}}\|_1$ on each trial. Plots omit the most ill-conditioned 25% of trials. The $\circ$ is the mean, and the lines indicate standard deviation.

each of G_m or G_{rd1} paired with either G_n or G_{rd3} . Thus, the mean VAF over the last 5 trials with G_m and G_{rd1} is greater than that with G_n and G_{rd3} . Furthermore, Table II shows that the mean $\|e\|$ over the last 5 trials with G_m and G_{rd1} is less than that with G_n and G_{rd3} .




Fig. 27. Mean and standard deviation of VAF on each trial. The mean VAF tends to increase over the trials. The $\circ$ is the mean, and the lines indicate standard deviation.

TABLE XIII
MEAN VAF AND CHANGE IN MEAN VAF FROM FIRST 5 TO LAST 5 TRIALS.

	Trials 1–5	Trials 6–10	Trials 11–20	Trials 21–30	Trials 31–35	Trials 36–40	Change
G_m	0.85	0.93	0.93	0.95	0.96	0.96	0.11
G_n	0.53	0.62	0.72	0.76	0.80	0.84	0.31
G_{rd1}	0.91	0.92	0.92	0.92	0.94	0.94	0.03
G_{rd2}	0.72	0.76	0.82	0.85	0.88	0.90	0.18
G_{rd3}	0.52	0.60	0.62	0.68	0.71	0.74	0.22