

The Impact of Command-Following Task on Human-in-the-Loop Control Behavior

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

Abstract—This paper presents results from an experiment in which 44 human subjects interact with a dynamic system 40 times over a one-week period. The subjects are divided into 4 groups. All groups interact with the same dynamic system, but each group performs a different sequence of command-following tasks. All reference commands have frequency content between 0 and 0.5 Hz. We use a subsystem identification algorithm to estimate the control strategy (feedback and feedforward) that each subject uses on each trial. The experimental and identification results are used to examine the impact of the command-following tasks on the subjects' performance and the control strategies that the subjects learn. Results demonstrate that certain reference commands (e.g., a sum of sinusoids) are more difficult for subjects to learn to follow than others (e.g., a chirp), and the difference in difficulty is related to the subjects' ability to match the phase of the reference command. In addition, the identification results show that differences in command-following performance for different tasks can be attributed to 3 aspects of the subjects' identified controllers: i) compensating for time delay in feedforward; ii) using a comparatively accurate approximation of the inverse dynamics in feedforward; and iii) using a feedback controller with comparatively high gain. Results also demonstrate that subjects generalize their control strategy when the command changes. Specifically, when the command changes, subjects maintain relatively high gain in feedback and retain their feedforward internal model of the inverse dynamics. Finally, we provide evidence that subjects use prediction of the command (if possible) to improve performance but that subjects can learn to improve performance without prediction. Specifically, subjects learn to use feedback controllers with comparatively high gain to improve performance even though the command is unpredictable.

Index Terms—Human control behavior, learning, subsystem identification, task, generalization, prediction.

I. INTRODUCTION

Humans often perform command-following tasks when interacting with dynamic systems. One example is driving an automobile, where the command-following task is to steer the automobile along the path of the road. In this case, the reference command is the road path trajectory. Humans learn to perform command-following tasks for a variety of reference commands—some of which are more challenging than others. In some cases, the reference command is known in advance or at least predictable. Conversely, the reference command may be both unknown and unpredictable. For example, in driving,

the road may be frequently traveled and thus familiar, or it may be unfamiliar and unpredictable (e.g., a winding road).

From a control-system-design perspective, the reference command often influences controller design. Generally, some knowledge of the reference command is required to design a control system that achieves good command-following performance. For example, a control strategy that incorporates an approximation of the inverse system dynamics in feedforward can yield good command-following performance. However, inverse dynamics are generally not causal, and thus, approximating them in feedforward may require that the reference command is known in advance or at least predictable. Alternatively, feedback control can be used to achieve good command-following performance by selecting a controller that makes the magnitude of the loop transfer function large at the frequencies of the reference command. The most common example of this approach is the use of integral control for following step commands. In this case, the controller contains an integrator, which makes the magnitude of the loop transfer function infinite at zero frequency (i.e., the frequency of a constant). This idea can be generalized to a variety of periodic commands by designing feedback controllers that incorporate internal models of the reference command [1]–[6]. In this case, accurate knowledge of the reference command's frequency content is required to design the feedback controller.

This paper examines the strategies that humans use for command-following tasks. Human-in-the-loop (HITL) control behavior and human learning have been studied in a variety of experiments (e.g., [7]–[10]), as well as by comparing proposed models of HITL control strategies with results from HITL experiments (e.g., [11]–[22]). However, [23] demonstrates that different control strategies can yield similar closed-loop responses. Thus, a proposed model that reproduces qualitative features of an HITL experiment is not necessarily an accurate representation of the human's control strategy.

As an alternative, system identification methods can be used to estimate models of the control strategies that humans use in HITL experiments [24]–[48]. In [49], [50], 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 (feedback and feedforward) used by a human in an experiment. This method does not require that a specific control strategy is assumed *a priori*. In [23], this SSID algorithm is used to model the feedforward and feedback control (including feedback time delay) that subjects use in an HITL experiment, where subjects interact with an LTI dynamic system and perform a command-following (i.e., pursuit-tracking) task. The results in [23] demonstrate that

S. A. S. Mousavi, F. Matveeva, X. Zhang, T. M. Seigler, and J. B. Hoagg are with the Department of Mechanical Engineering, University of Kentucky, Lexington, KY, USA. (e-mail: sse243@g.uky.edu, fsma222@g.uky.edu, xzh273@uky.edu, tmseigler@uky.edu, jesse.hoagg@uky.edu).

This work is supported in part by the National Science Foundation (CMMI-1405257, OIA-1849213) and the Kentucky Science and Engineering Foundation (KSEF-3453-RDE-018).

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 in neuroscience [51]–[53].

The experiment in [23] has subjects repeat one command-following task multiple times. This raises questions regarding the impact of the command-following task (i.e., reference command) on HITL control behavior. First, do humans learn to approximate the system’s inverse dynamics in feedforward for a variety of command-following tasks, and if so, how does the specific task impact their ability to learn to implement the approximate inverse dynamics? Second, what happens if the task changes? Do subjects generalize their control strategy from one task to another, and if so, how do they generalize? A control strategy is said to *generalize* if it can be learned in one situation (e.g., one task) and effectively transferred to another situation (e.g., another task) [54]. Studies that provide evidence of generalization include [55]–[59]. Third, what is the impact of reference-command prediction? If the same task (i.e., reference command) is repeated multiple times, then a subject can learn the task, which may allow them to use a control strategy that incorporates a prediction of the reference into the future. Prediction can be interpreted mathematically as using a noncasual (i.e., improper) control strategy. Thus, prediction may be important for approximating the inverse dynamics in feedforward, because if an LTI dynamic system is strictly proper, then its inverse dynamics are improper. Even if a proper approximation of the inverse dynamics are used in feedforward, then the approximation contains phase lead at certain frequencies, and prediction may help humans implement the required phase lead.

This paper provides new insights into the questions in the previous paragraph. We present results from an HITL experiment, where 44 subjects interact with an LTI dynamic system 40 times over a one-week period. The subjects are divided into 4 groups, where each group has 11 subjects. Each group interacts with the same LTI dynamic system but performs a different sequence of command-following tasks over the trials. We extend the SSID algorithm from [23], [49], [50] to allow for identification of the feedforward time delay as well as the feedforward transfer function, feedback transfer function, and feedback time delay. We use this extended SSID method to model the control strategies that each subject uses on each trial, and we examine the impact of different tasks by analyzing the subjects’ command-following performance and the best-fit models of the subjects’ control strategies.

This paper presents several new contributions. First, experimental results show that certain reference commands (e.g., a sum of sinusoids) are more difficult for subjects to learn to follow than others (e.g., a chirp), and the difference in difficulty is related to the subjects’ ability to match the phase of the reference command. Furthermore, the SSID results suggest that differences in command-following performance for different tasks can be attributed to 3 aspects of the subjects’ identified controllers: i) compensating for time delay in feedforward; ii) using a comparatively accurate approximation of the inverse dynamics in feedforward (particularly phase lead); and iii) using a feedback controller with comparatively high

gain across the frequency range of the reference command.

Second, we provide evidence that the subjects generalize their control strategy when the reference command changes, and we identify specific elements of their control that are generalized. Specifically, the SSID results show that subjects generalize by retaining aspects of their feedback and feedforward control strategies—subjects maintain relatively high gain in their feedback control and retain their feedforward internal model of the inverse dynamics.

Third, we provide evidence that subjects use prediction of the reference (if possible) to improve command-following performance but that subjects can learn to improve performance without prediction. The SSID results suggest that if the reference command is unpredictable, then subjects cannot learn to compensate for time delay in feedforward or use a comparatively accurate approximation of the inverse dynamics in feedforward (particularly phase lead). However, in this case, subjects can improve performance by learning to use a feedback controller with comparatively high gain across the frequency range of the reference command. Some preliminary results from this paper appear in [60]; however, this article goes beyond [60] by presenting SSID results that include time delay in the feedforward path, and by presenting significantly extended methods, analyses, and discussion in Sections III–VI.

II. EXPERIMENTAL METHODS

Forty-four people (38 male and 6 female) voluntarily participated in 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. All participants signed an informed consent form prior to participating in the experiment.

In this study, subjects use a single-degree-of-freedom joystick to affect the single-degree-of-freedom 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 a dynamic system. The horizontal position of the controlled object is denoted by y , which is the output of the dynamic system. Another object also moves on the computer screen, and its horizontal position 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 . Figure 1 is a diagram of the experimental setup, and Fig. 2 is a photograph 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. The subjects have no knowledge of the dynamic system relating u and y , or the reference object’s trajectory r prior to the experiment. Furthermore, prior to the experiment, the subjects have no experience with the experimental setup (e.g., the joystick or computer screen interface).




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 a dynamic system, and the controlled object's position y is the output of the dynamic system. A reference object's position r is also shown on the computer screen.




Fig. 2. Photograph of the experimental setup.

A trial is a 60-s period during which a subject manipulates the joystick. Each subject performed 40 trials of the experiment over 7 days. These trials were divided into 4 sessions of 10 trials, and each session was completed in a 20-minute period. Each subject completed no more than one session in a 12-hour period.

For each session, a subject sits in a chair facing the computer screen, which is approximately 60 cm from the subject's eyes. The computer screen is 47.6 cm high and 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. On all trials, subjects chose to use their dominant hand to manipulate the joystick.

The controlled object's position y satisfies the linear time-invariant (LTI) differential equation

$$\ddot{y}(t) + a_2\dot{y}(t) + a_1y(t) + a_0y(t) = b_1\dot{u}(t) + b_0u(t), \quad (1)$$

where $a_0 = 6.4$, $a_1 = 9.76$, $a_2 = 5.2$, $b_0 = 7.04$, and $b_1 = 3.2$, and the initial conditions are $\ddot{y}(0) = \dot{y}(0) = y(0) = 0$. Thus, the transfer function from u to y is given by

$$\mathcal{G}(s) \triangleq \frac{3.2(s + 2.2)}{(s + 1.6)(s^2 + 3.6s + 4)},$$

which has poles at -1.6 and $-1.8 \pm j0.87$, and a zero at -2.2 .

We examine the effects of command-following task (e.g., reference command) by using different reference-command signals r for different subjects and on different trials. We consider 22 different reference-command signals. For all $t \in$

$[0, 60]$, define the reference command

$$c_c(t) \triangleq 2 \sin \frac{\pi t^2}{120},$$

which is an 60-s chirp with frequency content between 0 and 0.5 Hz. For all $t \in [0, 60]$, define the reference command

$$c_s(t) \triangleq \frac{1}{3} \sum_{j=1}^{30} \cos \left(\frac{2\pi j t}{60} - \frac{j(j-1)\pi}{30} \right),$$

which is an 60-s sum of 30 sinusoids with evenly spaced frequencies between 0 and 0.5 Hz and with Schroeder phases [61]. For all $t \in [0, 60]$ and all $i \in \{1, \dots, 20\}$, define the reference command

$$c_i(t) \triangleq \frac{1}{3} \sum_{j=1}^{30} \cos \left(\frac{2\pi j t}{60} + \phi_{i,j} \right),$$

where for all $i \in \{1, \dots, 20\}$ and all $j \in \{1, \dots, 30\}$, $\phi_{i,j}$ is a randomly selected phase such that $c_i(0) = 0$ and the peak magnitude is less than 2.6, that is, $\max_{t \in [0, 60]} |c_i(t)| < 2.6$. Thus, c_i is an 60-s sum of 30 sinusoids with evenly spaced frequencies between 0 and 0.5 Hz and with randomly selected phases $\phi_{i,j}$. The units of the reference command are hash marks (hm), which are vertical lines separated by 2.5 cm on the computer display. The computer screen displays a range of ± 8 hm.

The dynamic system (1) is simulated using a dSPACE DS1103 control board. The dSPACE board also measures u , and the ControlDesk software is used to display the controlled object and reference object on the computer screen.

To examine the effects of task (i.e., reference command), the 44 subjects are divided into 4 groups, where each group has 11 subjects. All subjects interact with the dynamic system (1), but each group has a different sequence of tasks (i.e., reference commands). Group 1 performs 40 trials, where for all 40 trials, $r = c_c$. Group 2 performs 40 trials, where for all 40 trials, $r = c_s$. Group 3 performs 40 trials, where for trials $i \in \{1, 2, \dots, 20\}$, $r = c_c$ and for trials $i \in \{21, 22, \dots, 40\}$, $r = c_s$. Group 4 performs 40 trials, where for trials $i \in \{1, 2, \dots, 20\}$, $r = c_i$ and for trials $i \in \{21, 22, \dots, 40\}$, $r = c_s$. Note that the reference commands $c_1, \dots, c_{20}$ for group 4's first 20 trials are unpredictable because the phases $\phi_{i,j}$ are randomly selected.

To examine the effect of the chirp reference command c_c in comparison to the sum-of-sinusoids reference command c_s , we compare experimental and SSID results for group 1 to those for group 2. To examine the effect of changing task from c_c to c_s , we compare experimental and SSID results for group 3 to those for groups 1 and 2. In particular, we examine whether or not the subjects in group 3 generalize their control strategy from the task on the first 20 trials to the task on the last 20 trials. To examine the effect of unpredictable reference commands, we compare experimental and SSID results for group 4 to those for group 2. For example, we examine whether or not the subjects in group 4 learn to improve their performance over the first 20 trials even though the reference command is unpredictable.

III. EXPERIMENTAL DATA

For each of the 1760 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.

A *divergent trial* is a trial, where 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. Group 3 has the most divergent trials—a total of 11, which is 2.5% of the trials. Divergent trials are omitted from the results reported in the rest of this paper.

TABLE I
NUMBER OF DIVERGENT TRIALS.

Group	Trials 1–5	Trials 6–15	Trials 16–20	Trials 21–25	Trials 26–35	Trials 36–40	Total
1	0	0	1	0	0	0	1
2	2	0	0	0	0	0	2
3	3	5	0	1	1	1	11
4	4	3	0	0	0	1	8

A. Time-Domain Data

For each trial, we define the time-averaged error

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

Figures 3 and 4 show y , r , and e for the first and last trials of the subject from each group whose $\|e\|$ on the last trial is the median (i.e., 6th best) of the subjects in the group. For each group, the median subject's $\|e\|$ on the last trial is less than their $\|e\|$ on the first trial. The time-averaged error on the last trial for the subject in group 1 is less than that for the subject in group 2, which is less than that for the subject in group 3, which is less than that for the subject in group 4.




Fig. 3. Output y and reference r on the first and last trials of the subject from each group whose $\|e\|$ on the last trial is the median of the group.




Fig. 4. Error e on the first and last trials of the subject from each group whose $\|e\|$ on the last trial is the median of the group.

Figure 5 shows the mean and standard deviation of $\|e\|$ on each trial for groups 1, 2, and 3. The open-loop (i.e., $u = 0$) time-averaged errors with c_c and c_s are 1.20 and 1.14. By the last trial, the mean $\|e\|$ for each group is at least 50% better than open loop. The mean $\|e\|$ for groups 1 and 2 decreases consistently over the trials. However, on each trial, the mean $\|e\|$ for group 1 is less than that for group 2, which suggests that the chirp c_c is easier to follow than the sum of sinusoids c_s .

In contrast, the mean $\|e\|$ for group 3 increases 54% from trial 20 to 21, which corresponds to the change in task from c_c to c_s . The mean $\|e\|$ for group 3 is comparable to the mean $\|e\|$ for group 1 during the first 20 trials and is comparable to the mean $\|e\|$ for group 2 during the last 20 trials. On trial 21, the mean $\|e\|$ for group 3 is greater than that of group 2 on trial 21 but is still 35% less than the mean $\|e\|$ for group 2 on trial 1. In fact, the mean $\|e\|$ for group 3 on trial 21 is less than the mean $\|e\|$ for group 2 on trials 1 to 13. This




Fig. 5. Mean and standard deviation of $\|e\|$ on each trial for groups 1, 2, and 3. For all groups, the mean $\|e\|$ decreases from trial 1 to 40. However, the mean $\|e\|$ for group 3 increases from trial 20 to 21 when the task changes. The $\circ$ is the mean, and the lines indicate the standard deviation.

the subjects in group 3 generalize the control strategies that they learn during the first 20 trials.

Figure 6 shows the mean and standard deviation of $\|e\|$ on each trial for groups 2 and 4. For each group, the mean $\|e\|$ decreases over the 40 trials. By trial 40, the mean $\|e\|$ for each group is at least 39% better than open loop. Recall that for group 4, the reference is different for each of the first 20 trials. Since the task changes on each trial, the subjects in group 4 are limited in their ability to use prediction of the reference. Nevertheless, the mean $\|e\|$ for group 4 decreases by 31% over the first 20 trials. However, the mean $\|e\|$ for group 4 does not change significantly between trials 11 and 20, suggesting that the subjects reach near-steady-state performance. Note that the




Fig. 6. Mean and standard deviation of $\|e\|$ on each trial for groups 2 and 4. For both groups, the mean $\|e\|$ decreases over the 40 trials. The $\circ$ is the mean, and the lines indicate the standard deviation.

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 the frequencies $\omega_i = 2\pi i/60$ rad/s, where $i \in \{1, 2, \dots, N\}$, which are $N = 30$ 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*

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

$$= \frac{1}{N} \sum_{i=1}^N \left| |y_{\text{dft}}(\omega_i)| - |r_{\text{dft}}(\omega_i)| \right|,$$

which is the frequency-averaged magnitude of the difference between the y_{dft} and r_{dft} assuming that the phase of y_{dft} is equal to the phase of r_{dft} . Similarly, for each trial, define the *frequency-averaged error in the phase of the output*

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

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

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} .

Figures 7 and 8 show the mean and standard deviation of E_m and E_p for each group on each trial. These results are similar to the time-domain results shown in Figs. 5 and 6. However, for each group, the mean E_p is generally greater 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. Tables II and III show the mean E_m and mean E_p for each group on different sets of trials.

For groups 1 and 2, the mean E_m and mean E_p decrease consistently over the trials; however, Tables II and III show that the mean E_p decreases more (in absolute and percent) than the mean E_m . Specifically, for groups 1 and 2, the mean E_m decreases by 50% and 29% from the first 5 trials to the last 5 trials. In contrast, for groups 1 and 2, the mean E_p decreases




Fig. 7. Mean and standard deviation of E_m and E_p on each trial for groups 1, 2, and 3. The mean E_m and E_p decrease from trial 1 to 40, but the mean E_p decreases more (in absolute and percent) than the mean E_m . The $\circ$ is the mean, and the lines indicate the standard deviation.




Fig. 8. Mean and standard deviation of E_m and E_p on each trial for groups 2 and 4. The mean E_m and E_p decrease from trial 1 to 40, but the mean E_p decreases more (in absolute and percent) than the mean E_m . The $\circ$ is the mean, and the lines indicate the standard deviation.

TABLE II
MEAN E_m .

Group	Trials 1–5	Trials 6–15	Trials 16–20	Trials 21–25	Trials 26–35	Trials 36–40
1	0.06	0.05	0.04	0.04	0.04	0.03
2	0.07	0.06	0.05	0.05	0.05	0.05
3	0.07	0.05	0.05	0.06	0.05	0.05
4	0.08	0.07	0.07	0.05	0.05	0.05

TABLE III
MEAN E_p .

Group	Trials 1–5	Trials 6–15	Trials 16–20	Trials 21–25	Trials 26–35	Trials 36–40
1	0.11	0.07	0.05	0.05	0.04	0.04
2	0.15	0.12	0.09	0.09	0.07	0.06
3	0.12	0.08	0.07	0.10	0.08	0.07
4	0.19	0.17	0.16	0.13	0.11	0.09

by 64% and 60% from the first 5 trials to the last 5 trials. This observation suggests that these groups' improvement in $\|e\|$ is attributed more to improvement in matching the phase of the reference than improvement in matching the magnitude of the reference. In addition, the mean E_m for groups 1 and 2 are comparable on many trials, whereas the mean E_p for group 2 is significantly greater than that for group 1 on every trial. Thus, the fact that the mean $\|e\|$ for group 2 is greater than that for group 1 is attributed more to error in phase than error in magnitude. This suggests that the sum of sinusoids c_s is more difficult to follow than the chirp c_c , because it is more difficult to match the phase of the reference than its magnitude.

For group 3, the mean E_m and mean E_p decrease over the first 20 trials where the task is the same (i.e., predictable); increase from trial 20 to 21 where the task changes (i.e., is unpredictable); and decrease over the last 20 trials where the task is again the same. Tables II and III show that mean E_m and mean E_p increase by 20% and 43% from the last 5 trials before the task changes (i.e., trials 16–20) to the first 5 trials after the task changes (i.e., trials 21–25). Thus, the increase in $\|e\|$ when the task changes from c_c to c_s is attributed more to error in phase than error in magnitude. On trials 21–25, the mean E_m and mean E_p for group 3 is less than those for groups 2 and 3 on trials 1–5. This observation suggests that when the task changes on trial 21, the subjects in group 3 generalize the control strategy that they learn during the first 20 trials.

For group 4, the mean E_m and mean E_p decrease over the first 20 trials; however, these metrics for group 4 are consistently greater than and decrease less than for the other groups. Furthermore, these metrics for group 4 do not change significantly between trials 11 and 20. Notably, the mean E_p for group 4 is significantly greater than that of the other groups over the first 20 trials. Thus, group 4's mean $\|e\|$ over the first 20 trials is greater than that for the other groups primarily because of error in phase as opposed to error in magnitude. One possible explanation for Group 4's poor command-following performance over the first 20 trials (relative to the other groups) is that the reference is different on each trial, and thus, the subjects are limited in their

ability to predict the reference into the future. This potential explanation is supported by the fact that the group's command-following error is attributed more to error in phase than error in magnitude because limited predictive capability could manifest itself in phase lag (i.e., the phase of y lagging the phase of r). This hypothesis is discussed further with the SSID results presented in Section V.

IV. MODELING CONTROL STRATEGIES USING SSID

We discretize $\mathcal{G}$ 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) implies that

$$\hat{y}(z) = G(z)\hat{u}(z), \quad (2)$$

where $\hat{u}$ and $\hat{y}$ are the z -transforms of u_k and y_k .

Each subject's control strategy is modeled by the LTI control architecture shown in Fig. 9, which is given by

$$\hat{u}(z) = z^{-\tau_{fb}} G_{fb}(z) \hat{e}(z) + z^{-\tau_{ff}} G_{ff}(z) \hat{r}(z), \quad (3)$$

where $\hat{r}(z)$ and $\hat{e}(z)$ are the z -transforms of r_k and e_k ; G_{fb} and G_{ff} are the transfer functions of the feedback and feedforward controllers; and the nonnegative integers τ_{fb} and τ_{ff} are the feedback and feedforward delays. 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 . Define $T_{fb} \triangleq 10^3 \tau_{fb} T_s$ and $T_{ff} \triangleq 10^3 \tau_{ff} T_s$, which are the feedback and feedforward time delays in milliseconds. It follows from (2) and (3) that the closed-loop transfer function from r_k to y_k is

$$\tilde{G}_{yr}(z) \triangleq \frac{G(z) [z^{-\tau_{ff}} G_{ff}(z) + z^{-\tau_{fb}} G_{fb}(z)]}{1 + z^{-\tau_{fb}} G_{fb}(z) G(z)}, \quad (4)$$

and the closed-loop transfer function from r_k to e_k is

$$\tilde{G}_{er}(z) \triangleq 1 - \tilde{G}_{yr}(z) = \frac{1 - z^{-\tau_{ff}} G_{ff}(z) G(z)}{1 + z^{-\tau_{fb}} G_{fb}(z) G(z)}. \quad (5)$$




Fig. 9. Model of the control strategy. The control strategy is modeled using feedforward transfer function G_{ff} , feedforward delay τ_{ff} , feedback transfer function G_{fb} , and feedback delay τ_{fb} .

For each trial, we use SSID to determine the control strategy (i.e., G_{ff} , τ_{ff} , G_{fb} , τ_{fb}) of the form (3) that is the best-fit to the experimental data. For each trial and for $i \in \{1, 2, \dots, N\}$, define $H(\omega_i) \triangleq y_{dft}(\omega_i)/r_{dft}(\omega_i)$, which is the closed-loop frequency-response data from r_k to y_k .

Since y_k is bounded, we assume that $\tilde{G}_{yr}$ is asymptotically stable (i.e., its poles are contained in the open unit disk of

the complex plane). Thus, it follows from (4) that G_{ff} is asymptotically stable. The SSID algorithm presented in this section can be implemented in a computationally efficient manner if G_{ff} is finite impulse response (FIR). Thus, we let G_{ff} be FIR. Since G_{ff} is asymptotically stable, the assumption that G_{ff} is FIR does not significantly restrict the class of feedforward behavior; see [49] for more information.

The SSID objective is to determine G_{ff} , τ_{ff} , G_{fb} , and τ_{fb} such that the modeled frequency response $\{\tilde{G}_{yr}(e^{j\omega_i T_s})\}_{i=1}^N$ approximates the frequency-response data $\{H(\omega_i)\}_{i=1}^N$. More specifically, we seek to find the quadruple $(G_{\text{ff}}, \tau_{\text{ff}}, G_{\text{fb}}, \tau_{\text{fb}})$ that minimizes the cost function

$$\begin{aligned} J(G_{\text{ff}}, \tau_{\text{ff}}, G_{\text{fb}}, \tau_{\text{fb}}) &\triangleq \sum_{i=1}^N \left| \tilde{G}_{yr}(e^{j\omega_i T_s}) - H(\omega_i) \right|^2 \\ &= \sum_{i=1}^N \left| \frac{[e^{-j\omega_i T_s \tau_{\text{ff}}} G_{\text{ff}}(e^{j\omega_i T_s}) + e^{-j\omega_i T_s \tau_{\text{fb}}} G_{\text{fb}}(e^{j\omega_i T_s})]}{1 + e^{-j\omega_i T_s \tau_{\text{fb}}} G_{\text{fb}}(e^{j\omega_i T_s}) G(e^{j\omega_i T_s})} \right. \\ &\quad \left. \times G(e^{j\omega_i T_s}) - H(\omega_i) \right|^2, \end{aligned} \quad (6)$$

subject to the constraint that $\tilde{G}_{yr}$ is asymptotically stable.

To identify G_{ff} , τ_{ff} , G_{fb} , and τ_{fb} , we develop an SSID algorithm based on the method in [23]. However, [23] does not address identification of τ_{ff} . Thus, we extend the method in [23] to identify not only G_{ff} , G_{fb} , and τ_{fb} but also τ_{ff} . The details of this SSID algorithm are in Appendix A of the Supplemental Material, and the method is summarized as follows. First, we generate two candidate pools. The feedback candidate pool contains possible models of G_{fb} and τ_{fb} . Every element in this candidate pool is such that $\tilde{G}_{yr}$ is asymptotically stable. The feedforward-delay candidate pool contains possible values of τ_{ff} . For each possible model in the candidate pools, the cost J is convex in the coefficients of G_{ff} . Thus, for each model in the feedback candidate pool, we solve a sequence of convex optimizations to find the best-fit G_{ff} and τ_{ff} . Then, we search the feedback candidate pool to determine the quadruple $(G_{\text{ff}}, \tau_{\text{ff}}, G_{\text{fb}}, \tau_{\text{fb}})$ that minimizes J . This SSID algorithm has provable properties; see [49] for an analysis of the algorithm without delay (i.e., $\tau_{\text{fb}} = \tau_{\text{ff}} = 0$).

To achieve good command-following performance (i.e., $y_k \approx r_k$, or equivalently, $e_k \approx 0$), it follows that the control strategy should be such that $|\tilde{G}_{er}(e^{j\omega T_s})|$ is small for ω in the frequency range of the command r_k . Many different control strategies can be used to achieve good command-following performance—two such strategies are high gain in feedback and approximate inverse dynamics G^{-1} in feedforward. For high gain in feedback, $|G_{\text{fb}}(e^{j\omega T_s})|$ is large at the frequencies of r_k . In contrast, approximate inverse dynamics in feedforward requires $e^{-j\omega T_s \tau_{\text{ff}}} G_{\text{ff}}(e^{j\omega T_s}) \approx G^{-1}(e^{j\omega T_s})$ at the frequencies of r_k . With either of these strategies, it follows from (5) that $|\tilde{G}_{er}(e^{j\omega T_s})|$ is small at the frequencies of r_k . Thus, $|e_k|$ is small if $\tilde{G}_{er}$ is asymptotically stable. Other control strategies that achieve good command-following performance include combinations of approximating G^{-1} in feedforward and using high gain in feedback. See [23, Sec. IV] for a more detailed discussion of these different strategies.

Note that if $y_k \equiv r_k$, then $\tilde{G}_{yr}(e^{j\omega T_s}) = 1$ over the

frequency range of r_k . In this case, the SSID problem is not well posed because there are infinitely feedback $z^{-\tau_{\text{fb}}} G_{\text{fb}}$ and feedforward $z^{-\tau_{\text{ff}}} G_{\text{ff}}$ controller pairs that minimize the cost function (6). For example, (6) is minimized by any stabilizing feedback controller $z^{-\tau_{\text{fb}}} G_{\text{fb}}$ if the feedforward controller satisfies $z^{-\tau_{\text{ff}}} G_{\text{ff}} = G^{-1}$ over the frequency range of r_k . However, this degenerate case does not arise in the SSID analysis presented in this paper, because for all trials, $y_k \neq r_k$.

We use the SSID algorithm in Appendix A of the Supplemental Material as opposed to traditional system identification techniques (e.g., [62]–[66]), because the unknown subsystem (3) is connected in feedback with G , and traditional system identification algorithms applied to closed-loop architectures can yield trivial solutions [67], [68]. For example, $z^{-\tau_{\text{ff}}} G_{\text{ff}} = G^{-1}$ and $z^{-\tau_{\text{fb}}} G_{\text{fb}} = -G^{-1}$ is a trivial solution with many traditional system identification approaches. In contrast, $z^{-\tau_{\text{ff}}} G_{\text{ff}} = G^{-1}$ and $z^{-\tau_{\text{fb}}} G_{\text{fb}} = -G^{-1}$ does not minimize the SSID cost (6) because $H(\omega_k) \neq 1$. Thus, $z^{-\tau_{\text{ff}}} G_{\text{ff}} = G^{-1}$ and $z^{-\tau_{\text{fb}}} G_{\text{fb}} = -G^{-1}$ is not a solution with the SSID algorithm in Appendix A of the Supplemental Material. However, if the identified feedback controller satisfies $z^{-\tau_{\text{fb}}} G_{\text{fb}} \approx -G^{-1}$, then the denominator of (4) is approximately zero. In this case, the SSID results are ill conditioned. The feedback candidate pool is selected to prevent $z^{-\tau_{\text{fb}}} G_{\text{fb}} = -G^{-1}$. Appendix B in the Supplemental Material provides details on the SSID candidate pools used in this paper. Although the feedback candidate pool prevents $z^{-\tau_{\text{fb}}} G_{\text{fb}} = -G^{-1}$, the SSID algorithm could yield ill-conditioned results if $z^{-\tau_{\text{fb}}} G_{\text{fb}} \approx -G^{-1}$ at some frequencies. Appendix C in the Supplemental Material 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.

V. SSID RESULTS

For each trial, we identify the second-order exactly proper FIR feedforward transfer function G_{ff} , feedforward delay τ_{ff} , second-order strictly proper feedback transfer function G_{fb} , and feedback delay τ_{fb} that minimize J . The feedforward transfer function order is selected to allow G_{ff} to approximate G^{-1} with approximately 0.1% error over the 0-to-0.5 Hz range of the reference commands used in the experiment. See Appendix B in the Supplemental Material for the details of the candidate pools. The SSID algorithm is implemented on a supercomputer using parallel processing. Appendix D in the Supplemental Material presents a validation analysis of the SSID results.

Figures 10–13 show the Bode plots of the identified controllers $z^{-\tau_{\text{ff}}} G_{\text{ff}}$ and $z^{-\tau_{\text{fb}}} G_{\text{fb}}$, and the resulting closed-loop transfer function $\tilde{G}_{yr}$ on trial 1 and 40 for the subject whose $\|e\|$ on the last trial is the median of the subjects in that group. For these subjects, the closed-loop transfer function $\tilde{G}_{yr}$ is closer to one (i.e., unity magnitude and 0° phase) on trial 40 than on trial 1. This observation agrees with the time-domain results in Figs. 3 and 4, which show that y approximates r more closely on trial 40 than on trial 1. For these subjects, $\tilde{G}_{yr}$ is closer to one on the last trial, in part, because the




Fig. 10. Identified $z^{-\tau_{ff}}G_{ff}$ and $z^{-\tau_{fb}}G_{fb}$, and resulting $\tilde{G}_{yr}$ on trial 1 and 40 for the median subject from group 1. The identified $z^{-\tau_{ff}}G_{ff}$ is closer to G^{-1} and $\tilde{G}_{yr}$ is closer to one on trial 40 than on trial 1.




Fig. 11. Identified $z^{-\tau_{ff}}G_{ff}$ and $z^{-\tau_{fb}}G_{fb}$, and resulting $\tilde{G}_{yr}$ on trial 1 and 40 for the median subject from group 2. The identified $z^{-\tau_{ff}}G_{ff}$ is closer to G^{-1} and $\tilde{G}_{yr}$ is closer to one on trial 40 than on trial 1.




Fig. 12. Identified $z^{-\tau_{ff}}G_{ff}$ and $z^{-\tau_{fb}}G_{fb}$, and resulting $\tilde{G}_{yr}$ on trial 1 and 40 for the median subject from group 3. The identified $z^{-\tau_{ff}}G_{ff}$ is closer to G^{-1} and $\tilde{G}_{yr}$ is closer to one on trial 40 than on trial 1.




Fig. 13. Identified $z^{-\tau_{ff}}G_{ff}$ and $z^{-\tau_{fb}}G_{fb}$, and resulting $\tilde{G}_{yr}$ on trial 1 and 40 for the median subject from group 4. The identified $z^{-\tau_{ff}}G_{ff}$ is closer to G^{-1} and $\tilde{G}_{yr}$ is closer to one on trial 40 than on trial 1.

identified $z^{-\tau_{ff}}G_{ff}$ for the last trial approximates the inverse dynamics G^{-1} better than it does on the first trial.

A. Feedback and Feedforward Time Delay

Figures 14 and 15 show the mean and standard deviation of the identified feedback time delay T_{fb} and the identified feedforward time delay T_{ff} on each trial for each group. The average identified T_{fb} over all 40 trials for groups 1, 2, 3, and 4 are 232 ms, 215 ms, 213 ms, and 277 ms, respectively. These results for human feedback time delay with visual feedback are consistent with the results in [23]. Figure 14 also shows that there is no consistent trend in the mean T_{fb} over the trials.

In contrast, the mean T_{ff} tends to decrease over the 40 trials for groups 1 and 2. This observation suggests that groups 1 and 2 learn to predict the reference r into the future and use this prediction of the reference for feedforward control. This prediction is possible for groups 1 and 2 because the




Fig. 14. Mean and standard deviation of T_{fb} on each trial. The mean T_{fb} does not exhibit a consistent trend. The $\circ$ is the mean, and the lines indicate the standard deviation.

reference is the same for all 40 trials. However, the mean T_{ff} for group 1 is consistently less than the mean T_{ff} for group 2, which suggests that c_c is easier to learn to predict than c_s . Table IV shows the mean T_{ff} for each group on different sets




Fig. 15. Mean and standard deviation of T_{ff} on each trial. The mean T_{ff} tends to decrease for trials where the task (i.e., reference) is predictable. The $\circ$ is the mean, and the lines indicate the standard deviation.

TABLE IV
MEAN T_{ff} .

Group	Trials 1–5	Trials 6–15	Trials 16–20	Trials 21–25	Trials 26–35	Trials 36–40
1	118	40	22	15	12	10
2	292	183	119	94	55	43
3	98	43	37	114	81	36
4	350	311	350	299	195	147

For group 3, the mean T_{ff} decreases over the first 20 trials where the task is the same (i.e., predictable); increases significantly from trial 20 to 21 where the task changes (i.e., is unpredictable); and decreases over the last 20 trials where the task is again the same. Table IV shows that the mean T_{ff} increases by 208% from trials 16–20 to trials 21–25, which corresponds to the change in task from c_c to c_s . On trials 21–25, the mean T_{ff} for group 3 is greater than that of group 2 on trials 21–25 but is still 61% less than the mean T_{ff} for group 2 on trials 1–5. In fact, the mean T_{ff} for group 3 on trial 21 is less than the mean T_{ff} for group 2 on trials 1 to 7.

For group 4, the mean T_{ff} is relatively large and does not have a consistent trend over the first 20 trials where the task changes (i.e., is unpredictable). However, from trials 21 to 40, the task for group 4 does not change. During these last 20 trials, the mean T_{ff} for group 4 decreases. These trends for groups 3 and 4 are consistent with the observation that when the reference is predictable, subjects learn to predict the reference command into the future and use this prediction of the reference for feedforward control.

B. Feedforward Control

For each identified feedforward controller, we define

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

$$-G^{-1}(e^{j\omega T_s})| d\omega,$$

which is the frequency-averaged magnitude of the difference between the identified $z^{-\tau_{ff}} G_{ff}$ and the inverse dynamics G^{-1} over the 0-to-0.5 Hz range (i.e., the 0-to- π rad/s range). Figure 16 shows the mean and standard deviation of $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ on each trial for each group, and Table V shows the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ for each group on different




Fig. 16. Mean and standard deviation of $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ on each trial. The difference between the identified $z^{-\tau_{ff}} G_{ff}$ and G^{-1} decreases over the trials. The $\circ$ is the mean, and the lines indicate the standard deviation.

TABLE V
MEAN $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$.

Group	Trials 1–5	Trials 6–15	Trials 16–20	Trials 21–25	Trials 26–35	Trials 36–40
1	0.77	0.43	0.39	0.33	0.33	0.27
2	1.46	1.01	0.83	0.70	0.64	0.59
3	0.83	0.52	0.49	0.76	0.60	0.58
4	2.05	2.09	2.03	1.34	1.03	0.85

For groups 1 and 2, the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ decreases consistently over the trials. Figures 17 and 18 are the Bode plots of the average identified feedforward controller $z^{-\tau_{ff}} G_{ff}$ over all 11 subjects on trials 1, 20, 21, and 40 for groups 1 and 2, respectively. For both groups, the average identified $z^{-\tau_{ff}} G_{ff}$ on trial 40 approximates G^{-1} better than on trials 1, 20, and 21. Thus, by the last trial, the subjects learn an approximation of G^{-1} and use this approximation in feedforward. This result agrees with the results reported in [23] for a different experiment. This result also supports the internal model hypothesis in neuroscience [51]–[53].

For group 3, the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ decreases over the first 20 trials where the task is the same; increases from trial 20 to 21 where the task changes; and decreases over the last 20 trials where the task is again the same. Table V shows that mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ increases by 55% from trials 16–20 to trials 21–25, which corresponds to the change in task from c_c to c_s . However, the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ for group 3 on trials 21–25 is 8% less than that on trials 1–5 for group 3, and 48% less than that on trials 1–5 for group 2. This observation helps to explain the mechanism that the subjects in group 3 used to generalize from the c_c task to the c_s task—the subjects in group 3 retain some of their internal model




Fig. 17. Average identified feedforward controller for group 1 on trials 1, 20, 21, and 40. The shaded region shows the standard deviation.




Fig. 19. Average identified feedforward controller for group 3 on trials 1, 20, 21, and 40. The shaded region shows the standard deviation.




Fig. 18. Average identified feedforward controller for group 2 on trials 1, 20, 21, and 40. The shaded region shows the standard deviation.




Fig. 20. Average identified feedforward controller for group 4 on trials 1, 20, 21, and 40. The shaded region shows the standard deviation.

of G^{-1} , which did not change from trial 20 to 21. Figure 19 shows the average identified $z^{-\tau_{ff}} G_{ff}$ over all 11 subjects in group 3 on trials 1, 20, 21, and 40. The average identified $z^{-\tau_{ff}} G_{ff}$ on trial 21 approximates G^{-1} better than on trial 1. In particular, the phase of $z^{-\tau_{ff}} G_{ff}$ is closer to the phase of G^{-1} on trial 21 than on trial 1.

For group 4, the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ does not decrease over the first 20 trials where the task is changing, and decreases over the last 20 trials where the task is the same. The observation that the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ does not decrease over the first 20 trials is consistent with the task changing on each trial. Approximating G^{-1} in feedforward requires subjects to use phase lead over the 0-to-0.5 Hz range. Since the subjects cannot predict the reference, it is difficult to implement phase lead in feedforward. Figure 20 is the average identified $z^{-\tau_{ff}} G_{ff}$ over all 11 subjects in group 4 on trials 1, 20, 21, and 40. The average identified $z^{-\tau_{ff}} G_{ff}$ on trial 20 does not approximate G^{-1} as well as on trial 40. In fact, the

average identified $z^{-\tau_{ff}} G_{ff}$ on trial 20 does not have phase lead at frequencies greater than 0.1 Hz. Thus, some mechanism other than learning to approximate G^{-1} in feedforward is responsible for group 4 improving their performance over the first 10 trials. This improvement in performance over the first 10 trials is most likely a result of the feedback control used during these trials, which is discussed in Section V-C.

For group 4, one impediment to learning to approximate G^{-1} in feedforward over the first 20 trials is the feedforward time delay, which causes phase lag in the feedforward controller $z^{-\tau_{ff}} G_{ff}$. Since the reference changes on each of the first 20 trials, the subjects in group 4 cannot learn to predict the reference and compensate for feedforward time delay. Figure 21 is average identified feedforward transfer function G_{ff} over all 11 subjects in group 4 on trials 1, 20, 21, and 40. Note that G_{ff} is equal to the feedforward controller if the feedforward delay is zero. Thus, we can interpret G_{ff} as the feedforward controller that the subjects could achieve

if they could compensate completely for feedforward time delay. Figure 21 shows that by trial 20, G_{ff} approximates G^{-1} comparatively well from 0 to 0.25 Hz. This observation suggests that over the first 20 trials the subjects in group 4 are attempting to approximate G^{-1} in feedforward; however, the significant feedforward time delay (see Fig. 15) prevents an accurate approximation of the phase of G^{-1} .




Fig. 21. Average identified feedforward transfer function for group 4 on trials 1, 20, 21, and 40. The shaded region shows the standard deviation.

C. Feedback Control

For each identified feedback controller, we define

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

which is the frequency-averaged magnitude of $z^{-\tau_{fb}} G_{fb}$ over the 0-to-0.5 Hz range. Note that $\|z^{-\tau_{fb}} G_{fb}\|_1 = \frac{1}{\pi} \int_0^\pi |G_{fb}(e^{j\omega T_s})| d\omega$, and thus $\|z^{-\tau_{fb}} G_{fb}\|_1$ does not depend on the feedback delay τ_{fb} . Figure 22 shows the mean and standard deviation of $\|z^{-\tau_{fb}} G_{fb}\|_1$ on each trial for each group, and Table VI shows the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ for each group on different sets of trials.

For all groups, the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ tends to increase over the first 20 trials. In particular, for each group the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ increases from trials 1–5 to trials 16–20. These increases suggest that over the first 20 trials, the subjects learn to use more feedback gain in a frequency-averaged sense. One potential explanation is that through repeated interactions with G , the subjects learn to increase the feedback gain without causing closed-loop instability. To examine this conjecture, we compute the stability margins (i.e., upward gain margin and phase margin) associated with each identified feedback controller. Tables VII and VIII show the mean upward gain margins and mean phase margins for each group on different sets of trials during the first 20 trials. For each group, the mean upward gain margin decreases 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 upward gain margins. This observations suggests that through




Fig. 22. Mean and standard deviation of $\|z^{-\tau_{fb}} G_{fb}\|_1$ on each trial. The $\circ$ is the mean, and the lines indicate the standard deviation.

TABLE VI
MEAN $\|z^{-\tau_{fb}} G_{fb}\|_1$.

Group	Trials 1–5	Trials 6–15	Trials 16–20	Trials 21–25	Trials 26–35	Trials 36–40
1	1.85	1.82	2.48	1.82	2.32	2.27
2	1.22	1.64	1.77	1.61	1.72	1.98
3	1.49	1.94	2.91	1.70	1.61	1.42
4	1.14	1.25	1.36	1.06	1.06	1.40

TABLE VII
MEAN UPWARD GAIN MARGIN (ABSOLUTE).

Group	Trials 1–5	Trials 6–15	Trials 16–20	Trials 21–25	Trials 26–35	Trials 36–40
1	1.27	1.51	1.18	1.32	1.23	1.33
2	2.25	1.62	1.63	1.71	1.88	1.56
3	1.54	1.35	1.24	1.95	1.81	2.58
4	2.67	2.90	1.92	2.70	3.20	2.15

TABLE VIII
MEAN PHASE MARGIN (DEGREES).

Group	Trials 1–5	Trials 6–15	Trials 16–20	Trials 21–25	Trials 26–35	Trials 36–40
1	20	15	14	14	12	12
2	44	29	25	37	25	25
3	18	17	17	45	37	42
4	36	61	43	54	57	48

repeated interaction with G , the subjects learn the dynamics G well enough to use less conservative feedback controllers (i.e., feedback controllers with smaller upward gain margins) without causing closed-loop instability. Table VIII shows that for groups 1, 2, and 3, the mean phase margin also decreases from trials 1–5 to trials 16–20. In contrast, the mean phase margin for group 4 does not have a clear trend over the first 20 trials. Note that the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ and stability margins do not have a consistent trend over the last 15 trials. However, during these trials, the subjects are using a comparatively accurate approximation of G^{-1} in feedforward. In this case, (4) implies that the closed-loop transfer function $\hat{G}_{yr}$ (and thus, closed-loop response) is insensitive to the feedback controller G_{fb} . Thus, for trials where $z^{-\tau_{ff}} G_{ff}$ approximates G^{-1} , the identified G_{fb} may not be an accurate representation of the feedback used by the subjects.

Table VI shows that the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ for group 1

is consistently greater than that for group 2. Furthermore, Tables VII and VIII show that these larger feedback gains for group 1 relative to group 2 correspond with smaller upward gain margins and phase margins. Thus, group 1's mean $\|e\|$ is less than that of group 2, in part, because group 1 learns to use feedback controllers with larger gain and less conservative stability margins.

For group 3, the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ increases over the first 20 trials where the task is the same (i.e., predictable), and decreases from trial 20 to 21 where the task changes (i.e., is unpredictable). Table VI shows that the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ decreases by 42% from the last 5 trials before the task changes (i.e., trials 16–20) to the first 5 trials after the task changes (i.e., trials 21–25). Furthermore, Tables VII and VIII show that this decrease in feedback gain corresponds with a significant increase in upward gain margin and phase margin. This suggests that the subjects in group 3 compensated for the change in task by using more conservative feedback controllers. However, the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ for group 3 on trials 21–25 is 14% greater than the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ for group 3 on trials 1–5, and 39% greater than the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ for group 2 on trials 1–5. This observation helps to explain another mechanism that the subjects in group 3 use to generalize from the c_c task to the c_s task—when the task changes, group 3 do not revert to using feedback controllers with frequency-averaged gain as small as those used initially by the subjects in either group 2 or group 3. In other words, the subjects in group 3 retain some of their learned feedback control strategy when the task changes.

For group 4, the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ increases over the first 20 trials where the task changes (i.e., is unpredictable). Specifically, Table VI shows that the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ increases by 19% from trials 1–5 to trials 16–20. The observation that subjects learn to increase the feedback gain over the first 20 trials helps explain why the mean $\|e\|$ for group 4 decreases over the first 20 trials (as shown in Fig. 6) even though the reference is unpredictable.

VI. SUMMARY AND DISCUSSION

This paper presented several new contributions. First, we showed that the sum of sinusoids c_s is harder for humans to learn to follow than the chirp c_c ; furthermore, the associated SSID results provided insights into why c_s is harder to follow than c_c . The time-domain results (Fig. 5) show that c_s is harder to follow than c_c , that is, the time-averaged error $\|e\|$ for group 2 is greater than that for group 1 on every trial. The frequency-domain analysis (Fig. 7) shows that the difference in performance between groups 1 and 2 is attributed more to error in phase than error in magnitude. The SSID results show that the mean feedforward time delay T_{ff} for group 2 is greater than that for group 1 on every trial (see Fig. 15). Thus, group 2 does not compensate for feedforward time delay as well as group 1, which suggests that c_s is harder to learn to predict than c_c . Similarly, group 2 does not learn to approximate G^{-1} in feedforward as well as group 1 (see Fig. 16). The Bode plots of the average identified feedforward controllers (Figs. 17 and 18) show that group 2's approximation of G^{-1} in feedforward is worse than group 1's approximation, in large

part, because group 2 does not learn to use sufficient phase lead in feedforward (particularly at higher frequencies). This can be attributed, in part, to the fact that group 2 cannot predict the reference as well as group 1, which results in larger feedforward time delay, which, in turn, decreases the phase lead of the feedforward controller. Group 1 also uses consistently larger frequency-averaged gain in feedback (Fig. 22) than group 2, and these higher gains correspond to smaller stability margins. Thus, learning to use less-conservative feedback controllers with larger gain is another factor that contributes to smaller time-averaged error $\|e\|$ for group 1 relative group 2.

Second, we provided evidence that the subjects in group 3 generalize aspects of their control strategy when the reference changes from c_c to c_s , and we identified specific mechanisms of generalization. The time-domain data (Fig. 5) demonstrates that group 3 generalizes the control strategy they learn during the first 20 trials to the last 20 trials after the reference changes. The SSID results provide insights into how group 3 generalizes their control strategies. The frequency-averaged magnitude of the identified feedback controllers for group 3 decrease when the task changes (see Fig. 22), which suggests that group 3 compensates for the change in reference by using more conservative feedback controllers. However, the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ for group 3 after the reference changes is still greater than the mean $\|z^{-\tau_{fb}} G_{fb}\|_1$ for group 2 and group 3 on trials 1–5. Thus, when the reference changes, group 3 does not revert to using feedback controllers with frequency-averaged gain as small as those used initially by either group 2 or group 3. This observation suggests that when the task changes, group 3 retains aspects of their learned feedback control strategy. In addition, although the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ for group 3 increases when the reference changes, the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ for group 3 after the reference changes is still less than the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ for groups 2 and 3 on trials 1–5 (see Fig. 16). Thus, when the reference changes, group 3 retains their feedforward internal model of G^{-1} . It is also noteworthy that the mean feedforward time delay T_{ff} for group 3 increases when the reference changes, because the reference is unpredictable for the trials immediately after the change. This increase in mean T_{ff} is one factor that causes the the mean $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ to increase when the reference changes.

Third, we provided evidence that humans use prediction of the reference (if possible) to improve command-following performance, but that humans can learn to improve performance without prediction. The time-domain data (Fig. 6) shows that the mean $\|e\|$ for group 4 decreases over the first 20 trials, where the reference changes on each trial and is unpredictable. However, the mean $\|e\|$ for group 4 is greater than that for group 2 during this first 20 trials and does not change significantly between trials 11 and 20, suggesting that there is a limit to the subjects' achievable performance when the reference is unpredictable. The SSID results provide insights into how group 4 improves performance over the first 20 trials, where the reference is unpredictable. Specifically, the frequency-averaged magnitude of the identified feedback controllers for group 4 increases over the first 20 trials (see Fig. 22). This increase in feedback gain over the first 20 trials

helps explain why the mean $\|e\|$ for group 4 decreases over the first 20 trials even though the reference is unpredictable. In contrast to the other groups, group 4 does not learn to compensate for and decrease feedforward time delay over the first 20 trials (see Fig. 15). This observation suggests that humans use prediction of the reference into the future for feedforward control if and only if the reference is predictable. Similarly, group 4 does not learn to approximate G^{-1} in feedforward over the first 20 trials, which is in contrast to the other groups (see Fig. 16). Notably, the SSID results (see Figs. 20 and 21) show that group 4 attempts to approximate G^{-1} in feedforward; however, the significant feedforward time delay prevents an accurate approximation of the phase of G^{-1} .

The results of this paper provide new insights into the impact of command-following task on HITL control behavior. These results could have application to design and analysis for a variety of HITL technologies, including: active prostheses and exoskeletons; co-robotic systems (e.g., robotic-therapy devices for motor rehabilitation); and human-operated devices and vehicles (e.g., automobiles and aircraft). Nevertheless, open questions remain. For example, how do the results in this paper extend to more complex dynamic systems (e.g., higher-order systems, nonlinear systems, systems with multiple inputs and outputs)? As one example, it would be interesting to examine how haptic feedback (e.g., [69]) would impact the results described in this paper.

REFERENCES

- [1] P. C. Young and J. C. Willems. An approach to the multivariable servomechanism problem. *Int. J. Contr.*, 15:961–979, 1972.
- [2] E. J. Davison. The output control of linear time-invariant multivariable systems with unmeasurable arbitrary disturbances. *IEEE Trans. Autom. Contr.*, 17:621–630, 1972.
- [3] E. J. Davison and A. Goldenberg. Robust control of a general servomechanism problem: The servo compensator. *Automatica*, 11:461–471, 1975.
- [4] B. A. Francis, A. Sebakh, and W. M. Wonham. Synthesis of multivariable regulators: The internal model principle. *J. Appl. Math. Optim.*, 1:64–86, 1974.
- [5] B. A. Francis and W. M. Wonham. The internal model principle for linear multivariable regulators. *J. Appl. Math. Optim.*, 2:170–194, 1975.
- [6] J. B. Hoagg, M. A. Santillo, and D. S. Bernstein. Internal model control in the shift and delta domains. *IEEE Trans. Autom. Contr.*, 53:1066–1072, 2008.
- [7] R. Shadmehr and F. A. Mussa-Ivaldi. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14:3208–3224, 1994.
- [8] F. Gandolfo, F. A. Mussa-Ivaldi, and E. Bizzi. Motor learning by field approximation. *Proc. National Academy of Sciences*, 93:3843–3846, 1996.
- [9] J. V. Cohn, P. Dizio, and J. R. Lackner. Reaching during virtual rotation: Context specific compensations for expected coriolis forces. *Journal of Neurophysiology*, 83:3230–3240, 2000.
- [10] 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. *Journal of Neuroscience*, 17:1519–1528, 1997.
- [11] T. Flash and N. Hogan. The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5:1688–1703, 1985.
- [12] Y. Wada and M. Kawato. A neural network model for arm trajectory formation using forward and inverse dynamics models. *Neural Networks*, 6:919–932, 1993.
- [13] M. Katayama and M. Kawato. Virtual trajectory and stiffness ellipse during multijoint arm movement predicted by neural inverse models. *Biol. Cybern.*, 69:353–362, 1993.
- [14] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan. An internal model for sensorimotor integration. *Science*, 269:1880–1882, 1995.
- [15] N. Bhusan and R. Shadmehr. Computational nature of human adaptive control during learning of reaching movements in force fields. *Biol. Cybern.*, 81:39–60, 1999.
- [16] J. B. Dingwell, C. D. Mah, and F. A. Mussa-Ivaldi. Manipulating objects with internal degrees of freedom: evidence for model-based control. *Journal of Neurophysiology*, 88:222–235, 2002.
- [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. *Progress in Brain Research*, 142:203–222, 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. *Journal of Neurophysiology*, 91:1158–1170, 2004.
- [19] E. Guigon, P. Baraduc, and M. Desmurget. Computational motor control: redundancy and invariance. *Journal of Neurophysiology*, 97:331–347, 2007.
- [20] D. Liu and E. Todorov. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *Journal of Neuroscience*, 27:9354–9368, 2007.
- [21] H. Chen-Harris, W. M. Joiner, V. Ethier, D. S. Zee, and R. Shadmehr. Adaptive control of saccades via internal feedback. *Journal of Neuroscience*, 28:2804–2813, 2008.
- [22] A. J. Nagengast, D. A. Braun, and D. M. Wolpert. Optimal control predicts human performance on objects with internal degrees of freedom. *PLOS Computational Biology*, 5:e1000419, 2009.
- [23] 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. Cybernetics*, 48(2):543–555, 2018.
- [24] D. T. McRuer and E. S. Krendel. The human control operator as a servo system element. *Journal of the Franklin Institute*, 6(267):511–536, 1959.
- [25] D. McRuer, D. Graham, E. Krendel, and W. Reisener. Human pilot dynamics in compensatory systems: Theory, models, and experiments with controlled element and forcing function variations. Technical Report AFFDL-TR-65-15, Air Force Flight Dynamics Laboratory, Research and Technology Division, Air Force Systems Command, Wright-Patterson Air Force Base, 1965.
- [26] R. J. Wasicko, D. T. McRuer, and R. E. Magdaleno. Human pilot dynamics in single-loop systems with compensatory and pursuit displays. Technical Report AFFDL-TR-66-137, Air Force Flight Dynamics Laboratory, Research and Technology Division, Air Force Systems Command, Wright-Patterson Air Force Base, 1966.
- [27] D. T. McRuer, D. Graham, and E. S. Krendel. Manual control of a single-loop system: Part i. *Journal of the Franklin Institute*, 1(283):1–29, 1967.
- [28] D. T. McRuer, D. Graham, and E. S. Krendel. Manual control of a single-loop system: Part ii. *Journal of the Franklin Institute*, 2(283):145–168, 1967.
- [29] R. L. Stapleford, D. T. McRuer, and R. E. Magdaleno. Pilot describing function measurements in a multiloop task. *IEEE Trans. Human Factors in Electronics*, 8(2):113–125, 1967.
- [30] D. T. McRuer and D. H. Weir. Theory of manual vehicle control. *IEEE Trans. Man-Machine Systems*, 10(4):257–291, 1969.
- [31] R. A. Hess. Pursuit tracking and higher levels of skill development in the human pilot. *IEEE Trans. Systems, Man, and Cybernetics*, 11(4):262–273, 1981.
- [32] A. Fagergren, O. Ekeberg, and H. Forssberg. Precision grip force dynamics: A system identification approach. *IEEE Trans. Biomed. Eng.*, 47:1366–1375, 2000.
- [33] A. Ghoreyshi and H. L. Galiana. Simultaneous identification of oculomotor subsystems using a hybrid system approach: Introducing hybrid extended least squares. *IEEE Trans. Biomed. Eng.*, 57:1089–1098, 2010.
- [34] 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. Cybernetics*, 43(6):1936–1949, 2013.
- [35] J. Venrooij, M. M. van Paassen, M. Mulder, D. A. Abbink, M. Mulder, F. C. T. van der Helm, and H. H. Bultthoff. A framework for biodynamic feedthrough analysis—part i: Theoretical foundations. *IEEE Trans. Cybernetics*, 44(9):1686–1698, 2014.
- [36] J. Venrooij, M. M. van Paassen, M. Mulder, D. A. Abbink, M. Mulder, F. C. T. van der Helm, and H. H. Bultthoff. A framework for biodynamic feedthrough analysis—part ii: Validation and application. *IEEE Trans. Cybernetics*, 44(9):1699–1710, 2014.
- [37] J. Venrooij, D. A. Abbink, M. Mulder, M. M. van Paassen, M. Mulder, F. C. T. van der Helm, and H. H. Bultthoff. A biodynamic feedthrough

- model based on neuromuscular principles. *IEEE Trans. Cybernetics*, 44(7):1141–1154, 2014.
- [38] J. Venrooij, M. Mulder, D. A. Abbink, M. M. van Paassen, M. Mulder, F. C. T. van der Helm, and H. H. Bultmann. Mathematical biodynamic feedthrough model applied to rotorcraft. *IEEE Trans. Cybernetics*, 44(7):1025–1038, 2014.
- [39] 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. Cybernetics*, 45(2):253–265, 2015.
- [40] 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 Transactions on Cybernetics*, 46(11):2609–2621, 2016.
- [41] J. Venrooij, M. Mulder, M. Mulder, D. A. Abbink, M. M. van Paassen, F. C. T. van der Helm, and H. H. Bultmann. Admittance-adaptive model-based approach to mitigate biodynamic feedthrough. *IEEE Trans. Cybernetics*, 47(12):4169–4181, 2017.
- [42] 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 Transactions on Cybernetics*, 48(4):1242–1252, 2018.
- [43] 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 Transactions on Human-Machine Systems*, 48(5):496–508, 2018.
- [44] 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 Transactions on Human-Machine Systems*, 48(5):486–495, 2018.
- [45] W. Fu, A. Landman, M. M. van Paassen, and M. Mulder. Modeling human difference threshold in perceiving mechanical properties from force. *IEEE Transactions on Human-Machine Systems*, 48(4):359–368, 2018.
- [46] F. M. Drop, D. M. Pool, M. R. M. van Paassen, M. Mulder, and H. H. Bultmann. Objective model selection for identifying the human feedforward response in manual control. *IEEE Transactions on Cybernetics*, 48(1):2–15, 2018.
- [47] F. M. Drop, D. M. Pool, M. M. van Paassen, M. Mulder, and H. H. Bultmann. Effects of target signal shape and system dynamics on feedforward in manual control. *IEEE Transactions on Cybernetics*, 49(3):768–780, 2019.
- [48] W. Fu, M. M. van Paassen, D. A. Abbink, and M. Mulder. Framework for human haptic perception with delayed force feedback. *IEEE Transactions on Human-Machine Systems*, 49(2):171–182, 2019.
- [49] X. Zhang and J. B. Hoagg. Subsystem identification of multivariable feedback and feedforward systems. *Automatica*, 72:131–137, 2016.
- [50] X. Zhang and J. B. Hoagg. Frequency-domain subsystem identification with application to modeling human control behavior. *Sys. Contr. Lett.*, 87:36–46, 2016.
- [51] D. M. Wolpert, R. C. Miall, and M. Kawato. Internal models in the cerebellum. *Trends in Cognitive Sciences*, 2:338–347, 1998.
- [52] M. Kawato. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9:718–727, 1999.
- [53] R. Shadmehr, M. A. Smith, and J. W. Krakauer. Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, 33:89–108, 2010.
- [54] D. J. Reinkensmeyer, J. L. Emken, and S. C. Cramer. Robotics, motor learning, and neurologic recovery. *Annual Review of Biomedical Engineering*, 6:497–521, 2004.
- [55] M. A. Conditt and F. A. Mussa-Ivaldi. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *Journal of Neurophysiology*, 78:554–560, 1997.
- [56] M. A. Conditt and F. A. Mussa-Ivaldi. Central representations of time during motor learning. *Proc. National Academy of Sciences*, 96:11625–11630, 1999.
- [57] A. Karniel and F. A. Mussa-Ivaldi. Does the motor control system use multiple models and context switching to cope with a variable environment? *Experimental Brain Research*, 143:520–524, 2002.
- [58] N. Malfait, D. M. Sherrill, and D. J. Ostry. Transfer of motor learning across arm configurations. *Journal of Neuroscience*, 22:2956–2960, 2002.
- [59] N. Censor, D. Sagi, and L. G. Cohen. Common mechanisms of human perceptual and motor learning. *Nature Reviews Neuroscience*, 13:658–664, 2012.
- [60] F. Matveeva, S. A. Seyyed Mousavi, X. Zhang, T. M. Seigler, and J. B. Hoagg. On the effects of changing reference command as humans learn to control dynamic systems. In *Proc. Conf. Dec. Contr.*, pages 1211–1216, Las Vegas, NV, December 2016.
- [61] M. A. Schroeder. Synthesis of low peak-factor signals and binary sequences of low autocorrelation. *IEEE Trans. Info. Theory*, 16:85–89, 1970.
- [62] T. Soderstrom and P. Stoica. *System Identification*. Prentice-Hall, Upper Saddle River, NJ, 1989.
- [63] J. N. Juang. *Applied System Identification*. Prentice-Hall, Upper Saddle River, NJ, 1993.
- [64] P. Van Overschee and B. De Moor. A unifying theorem for three subspace system identification algorithms. *Automatica*, 31(12):1853–1864, 1995.
- [65] L. Ljung. *System Identification: Theory for the User, 2nd ed.* Prentice-Hall Information and Systems Sciences, Upper Saddle River, NJ, 1999.
- [66] R. Pintelon and J. Schoukens. *System Identification: A Frequency Domain Approach*. IEEE Press, Piscataway, NJ, 2001.
- [67] L. Ljung I. Gustavsson and T. Soderstrom. Identification of processes in closed loop—Identifiability and accuracy aspects. *Automatica*, 13(1):59–75, 1977.
- [68] U. Forsell and L. Ljung. Closed-loop identification revisited. *Automatica*, 35:1215–1241, 1999.
- [69] L. Scalera, S. Seriani, P. Gallina, M. Di Luca, and A. Gasparetto. An experimental setup to test dual-joystick directional responses to vibrotactile stimuli. *IEEE Trans. Haptics*, 11(3):378–387, 2018.
- [70] K. Amano, N. Goda, S. Nishida, Y. Ejima, T. Takeda, and Y. Ohtani. Estimation of the timing of human visual perception from magnetoencephalography. *Journal of Neuroscience*, 26(15):3981–3991, 2006.
- [71] J.-J. Orban de Xivry and P. Lefèvre. Saccades and pursuit: two outcomes of a single sensorimotor process. *Journal Physiology*, 584:11–23, 2007.



S. Alireza Seyyed Mousavi received the B.S. degree in Mechanical Engineering from Islamic Azad University of Mashhad in 2011, and the M.S. degree in Mechanical Engineering from Northern Illinois University in 2013, and the Ph.D. degree in Mechanical Engineering from University of Kentucky in 2019. He is currently a postdoctoral fellow with the Department of Intelligent Systems Engineering at Indiana University Bloomington.



Faina Matveeva received the B.S. degree in Mechanical Engineering from University of Kentucky in 2016.



Xingye Zhang received the B.S. degree in Systems and Control from Beihang University in 2009, the M.Eng. degree in Fundamental Mechanics (Dynamics and Control) from Beihang University in 2011, the M.A. degree in Mathematics from the University of Kentucky in 2015, and the Ph.D. degree in Mechanical Engineering from the University of Kentucky in 2015. He is currently a senior research engineer at Inceptio Technology.



Michael Seigler received the B.S. degree in Mechanical Engineering from Clemson University in 2000, and the M.S. and Ph.D. degrees in Mechanical Engineering from Virginia Tech in 2002 and 2005. In 2006, he joined the Mechanical Engineering department at the University of Kentucky, where he is currently an Associate Professor.



Jesse B. Hoagg received the B.S.E. degree in Civil and Environmental Engineering from Duke University in 2002, the M.S.E. degree in Aerospace Engineering from the University of Michigan in 2003, the M.S. degree in Mathematics from the University of Michigan in 2005, and the Ph.D. degree in Aerospace Engineering from the University of Michigan in 2006. He is currently the Donald and Gertrude Lester Professor of Mechanical Engineering at the University of Kentucky.

APPENDIX A
SUBSYSTEM IDENTIFICATION ALGORITHM

This appendix presents the SSID algorithm used to identify the feedforward transfer function, feedforward delay, feedback transfer function, and feedback delay used by subjects in the HITL experiment. This SSID algorithm is based on the method in [23], [49], which can identify multivariable LTI feedback and feedforward subsystems. However, the method in [23], [49] does not address the identification of feedforward time delay. In this appendix, we extend the method in [23], [49] to allow for the identification of feedforward time delay.

Let $G: \mathbb{C} \rightarrow \mathbb{C}$ be a real rational discrete-time transfer function, which is assumed to be known and is interconnected with an unknown dynamic subsystem as shown in Fig. 23. The input $\{r_j\}_{j=1}^n$ and output $\{y_j\}_{j=1}^n$ sequences are assumed to be measured with sample time T_s . Note that n is the number of samples.




Fig. 23. A known LTI subsystem G interconnected with an unknown subsystem. The external input r_k and the output y_k are measured.

Let N be a positive integer, and let $\omega_1 < \omega_2 < \dots < \omega_N$ be nonnegative. For $k \in \mathcal{N} \triangleq \{1, 2, \dots, N\}$, let $r_{\text{dft}}(\omega_k)$ and $y_{\text{dft}}(\omega_k)$ denote the discrete Fourier transform of $\{r_j\}_{j=1}^n$ and $\{y_j\}_{j=1}^n$ at ω_k ; define the closed-loop frequency response data $H(\omega_k) \triangleq y_{\text{dft}}(\omega_k)/r_{\text{dft}}(\omega_k)$; and define $\sigma_k \triangleq e^{j\omega_k T_s}$.

Let $G_{\text{ff}}, G_{\text{fb}}: \mathbb{C} \rightarrow \mathbb{C}$ denote real rational discrete-time transfer functions; let the nonnegative integers τ_{ff} and τ_{fb} denote the feedforward delay and feedback delay, respectively; and let G_{ff} be FIR. The unknown subsystem in Fig. 23 is modeled using the LTI control (3) shown in Fig. 9.

Our objective is to determine $G_{\text{ff}}, \tau_{\text{ff}}, G_{\text{fb}}$, and τ_{fb} such that the modeled frequency response $\{\tilde{G}_{yr}(\sigma_k)\}_{k=1}^N$ approximates the data $\{H(\omega_k)\}_{k=1}^N$, where $\tilde{G}_{yr}$ is given by (4). To achieve this objective, we seek to find $G_{\text{ff}}, \tau_{\text{ff}}, G_{\text{fb}}$, and τ_{fb} that minimize the cost $J(G_{\text{ff}}, \tau_{\text{ff}}, G_{\text{fb}}, \tau_{\text{fb}})$ given by (6), subject to the constraint that $\tilde{G}_{yr}$ is asymptotically stable.

We parameterize the feedback and feedforward controllers by their numerator and denominator coefficients and cast the SSID problem in terms of these coefficients. Let n_{ff} and d_{ff} be nonnegative integers that denote the degree of the numerator and denominator of G_{ff} , respectively, and define $a \triangleq d_{\text{ff}} + 1$. Let n_{fb} and d_{fb} be nonnegative integers that denote the degrees of the numerator and denominator of G_{fb} , respectively, and define $b \triangleq n_{\text{fb}} + d_{\text{fb}} + 1$.

Consider the functions $\mathcal{N}_{\text{ff}}: \mathbb{C} \times \mathbb{R}^a \rightarrow \mathbb{C}$, $\mathcal{N}_{\text{fb}}: \mathbb{C} \times \mathbb{R}^b \rightarrow \mathbb{C}$, and $\mathcal{D}_{\text{fb}}: \mathbb{C} \times \mathbb{R}^b \rightarrow \mathbb{C}$ given by

$$\begin{aligned}\mathcal{N}_{\text{ff}}(z, \alpha) &\triangleq [z^{n_{\text{ff}}} \quad z^{n_{\text{ff}}-1} \quad \dots \quad z \quad 1] \alpha, \\ \mathcal{N}_{\text{fb}}(z, \beta) &\triangleq [z^{n_{\text{fb}}} \quad \dots \quad z \quad 1 \quad 0_{1 \times d_{\text{fb}}}] \beta, \\ \mathcal{D}_{\text{fb}}(z, \beta) &\triangleq z^{d_{\text{fb}}} + [0_{1 \times (n_{\text{fb}}+1)} \quad z^{d_{\text{fb}}-1} \quad \dots \quad z \quad 1] \beta.\end{aligned}$$

Define

$$\mathcal{G}_{\text{ff}}(z, \alpha) \triangleq \frac{\mathcal{N}_{\text{ff}}(z, \alpha)}{z^{d_{\text{ff}}}}, \quad \mathcal{G}_{\text{fb}}(z, \beta) \triangleq \frac{\mathcal{N}_{\text{fb}}(z, \beta)}{\mathcal{D}_{\text{fb}}(z, \beta)},$$

where α contains the numerator coefficients of $\mathcal{G}_{\text{ff}}$, and β contains the numerator and denominator coefficients of $\mathcal{G}_{\text{fb}}$.

The real rational transfer function G can be expressed as $G = Z/P$, where Z and P are coprime polynomials. Next, consider the cost function $\mathcal{J}: \mathbb{R}^a \times \mathbb{N} \times \mathbb{R}^b \times \mathbb{N} \rightarrow [0, \infty)$ given by

$$\begin{aligned}\mathcal{J}(\alpha, \psi, \beta, \gamma) &\triangleq J(\mathcal{G}_{\text{ff}}(z, \alpha), \psi, \mathcal{G}_{\text{fb}}(z, \beta), \gamma) \\ &= \alpha^T \Omega_2(\beta, \gamma) \alpha + \Omega_0(\beta, \gamma) \\ &\quad + \text{Re } \Upsilon_1^*(\beta, \gamma) (\text{diag } \Gamma(\psi)) \Upsilon_2(\beta, \gamma) \alpha,\end{aligned}$$

where $\psi \in \mathbb{N}$ represents the feedforward delay, $\gamma \in \mathbb{N}$ represents the feedback delay, and

$$\begin{aligned}\Omega_2(\beta, \gamma) &\triangleq \text{Re} \sum_{k=1}^N A_k^*(\beta, \gamma) A_k(\beta, \gamma) \in \mathbb{R}^{a \times a}, \\ \Omega_0(\beta, \gamma) &\triangleq \sum_{k=1}^N |B_k(\beta, \gamma)|^2 \in \mathbb{R}, \\ \Upsilon_1(\beta, \gamma) &\triangleq 2 \begin{bmatrix} B_1(\beta, \gamma) & \dots & B_N(\beta, \gamma) \end{bmatrix}^T \in \mathbb{C}^N, \\ \Upsilon_2(\beta, \gamma) &\triangleq \begin{bmatrix} A_1(\beta, \gamma) & \dots & A_N(\beta, \gamma) \end{bmatrix}^T \in \mathbb{C}^{N \times a}, \\ \Gamma(\psi) &\triangleq \begin{bmatrix} \sigma_1^{-\psi} & \dots & \sigma_N^{-\psi} \end{bmatrix}^T \in \mathbb{C}^N,\end{aligned}$$

where for all $k \in \mathcal{N}$,

$$\begin{aligned}A_k(\beta, \gamma) &\triangleq \frac{\sigma_k^{\gamma-n_{\text{ff}}} Z(\sigma_k) \mathcal{D}_{\text{fb}}(\sigma_k, \beta)}{\tilde{\mathcal{D}}(\sigma_k, \beta, \gamma)} v(\sigma_k) \in \mathbb{C}^{1 \times a}, \\ B_k(\beta, \gamma) &\triangleq \frac{Z(\sigma_k) \mathcal{N}_{\text{fb}}(\sigma_k, \beta)}{\tilde{\mathcal{D}}(\sigma_k, \beta, \gamma)} - H(\omega_k) \in \mathbb{C},\end{aligned}$$

and

$$\begin{aligned}\tilde{\mathcal{D}}(z, \beta, \gamma) &\triangleq z^\gamma \mathcal{D}_{\text{fb}}(z, \beta) P(z) + \mathcal{N}_{\text{fb}}(z, \beta) Z(z), \\ v(z) &\triangleq \begin{bmatrix} z^{n_{\text{ff}}} & \dots & z & 1 \end{bmatrix}.\end{aligned}$$

We restrict our attention to $(\beta, \gamma) \in \mathbb{R}^b \times \mathbb{N}$ contained in

$$\begin{aligned}\mathcal{S} &\triangleq \{(\beta, \gamma) \in \mathbb{R}^b \times \mathbb{N}: \beta \in \mathbb{R}^b, \gamma \in \mathbb{N}, \text{ and if } \lambda \in \mathbb{C} \\ &\quad \text{and } \tilde{\mathcal{D}}(\lambda, \beta, \gamma) = 0, \text{ then } |\lambda| < 1\},\end{aligned}$$

which is the set of parameters that yield asymptotically stable closed-loop transfer functions.

Let m be a positive integer, and let $\Phi \subset \mathcal{S}$ be a set with m elements. We call Φ the *feedback candidate pool*. For all $i, j \in \mathcal{M} \triangleq \{1, 2, \dots, m\}$, let $\phi_i, \phi_j \in \Phi$ be such that if $i \neq j$, then $\phi_i \neq \phi_j$.

Let

$$E_\beta \triangleq \begin{bmatrix} I_b & 0_{b \times 1} \end{bmatrix}, \quad E_\gamma \triangleq \begin{bmatrix} 0_{1 \times b} & 1 \end{bmatrix}.$$

Then, for all $i \in \mathcal{M}$, define the cost function

$$\begin{aligned}\mathcal{J}_i(\alpha, \psi) &\triangleq \mathcal{J}(\alpha, \psi, E_\beta \phi_i, E_\gamma \phi_i) \\ &= \alpha^T \Omega_2(E_\beta \phi_i, E_\gamma \phi_i) \alpha + \Omega_0(E_\beta \phi_i, E_\gamma \phi_i) \\ &\quad + \text{Re } \Gamma^T(\psi) \Omega_1^T(E_\beta \phi_i, E_\gamma \phi_i) \alpha,\end{aligned}$$

where

$$\Omega_1(\beta, \gamma) \triangleq \Upsilon_2^T(\beta, \gamma) (\text{diag } \Upsilon_1(\beta, \gamma))^* \in \mathbb{C}^{a \times N}.$$

Note that for all $\psi \in \mathbb{N}$, $\mathcal{J}_i(\alpha, \psi)$ is convex in α . If the number N of frequency response data is sufficiently large, then it can be shown that $\Omega_2(E_\beta \phi_1, E_\gamma \phi_1), \dots, \Omega_2(E_\beta \phi_m, E_\gamma \phi_m)$ are positive definite and thus nonsingular. In this case, for each $i \in \mathcal{M}$,

$$\alpha_i(\psi) \triangleq -\frac{1}{2} \Omega_2^{-1}(E_\beta \phi_i, E_\gamma \phi_i) \text{Re } \Omega_1(E_\beta \phi_i, E_\gamma \phi_i) \Gamma(\psi)$$

exists, and for all $\psi \in \mathbb{N}$ and all $i \in \mathcal{M}$, $\alpha_i(\psi)$ is the unique global minimizer of $\mathcal{J}_i(\alpha, \psi)$. Define the auxiliary cost

$$\begin{aligned} \mathcal{Q}_i(\psi) &\triangleq \mathcal{J}_i(\alpha_i(\psi), \psi) \\ &= - \begin{bmatrix} \text{Re } \Gamma(\psi) \\ -\text{Im } \Gamma(\psi) \end{bmatrix}^T \mathcal{F}_i \begin{bmatrix} \text{Re } \Gamma(\psi) \\ -\text{Im } \Gamma(\psi) \end{bmatrix} \\ &\quad + \Omega_0(E_\beta \phi_i, E_\gamma \phi_i), \end{aligned}$$

where

$$\begin{aligned} \mathcal{F}_i &\triangleq \frac{1}{4} \begin{bmatrix} \text{Re } \Omega_1^T(E_\beta \phi_i, E_\gamma \phi_i) \\ \text{Im } \Omega_1^T(E_\beta \phi_i, E_\gamma \phi_i) \end{bmatrix} \Omega_2^{-1}(E_\beta \phi_i, E_\gamma \phi_i) \\ &\quad \times \begin{bmatrix} \text{Re } \Omega_1^T(E_\beta \phi_i, E_\gamma \phi_i) \\ \text{Im } \Omega_1^T(E_\beta \phi_i, E_\gamma \phi_i) \end{bmatrix}^T. \end{aligned}$$

Let p be a positive integer, and let $\Psi \subset \mathbb{N}$ be a set with p elements. We call Ψ the *feedforward-delay candidate pool*. For all $i, j \in \mathcal{P} \triangleq \{1, 2, \dots, p\}$, let $\psi_i, \psi_j \in \Psi$ be such that if $i \neq j$, then $\psi_i \neq \psi_j$.

For all $i \in \mathcal{M}$, let $q_i \in \mathcal{P}$ be the smallest integer such that

$$\mathcal{Q}_i(\psi_{q_i}) = \min_{j \in \mathcal{P}} \mathcal{Q}_i(\psi_j).$$

Next, let $\ell \in \mathcal{M}$ be the smallest integer such that $\mathcal{Q}_\ell(\psi_{q_\ell}) = \min_{i \in \mathcal{M}} \mathcal{Q}_i(\psi_{q_i})$. Thus, the identified parameters are $\alpha_\ell(\psi_{q_\ell})$, $\tau_{\text{ff}} \triangleq \psi_{q_\ell}$, $\beta_\ell \triangleq E_\beta \phi_\ell$, and $\tau_{\text{fb}} \triangleq E_\gamma \phi_\ell$, which implies that the identified transfer functions are

$$G_{\text{ff}}(z) \triangleq \mathcal{G}_{\text{ff}}(z, \alpha_\ell(\psi_{q_\ell})), \quad G_{\text{fb}}(z) \triangleq \mathcal{G}_{\text{fb}}(z, \beta_\ell). \quad (7)$$

We now summarize this SSID method. For an analysis of its properties without delay, see [49].

Algorithm 1. Consider the closed-loop transfer function (4) and the frequency-response data $\{H(\omega_k)\}_{k=1}^N$.

- Step 1. Generate the feedback candidate pool $\Phi \subset \mathcal{S}$ and feedforward-delay candidate pool $\Psi \subset \mathbb{N}$.
- Step 2. For each $i \in \mathcal{M}$, compute $\mathcal{Q}_i(\psi)$ and find smallest integer $q_i \in \mathcal{P}$ such that $\mathcal{Q}_i(\psi_{q_i}) = \min_{j \in \mathcal{P}} \mathcal{Q}_i(\psi_j)$.
- Step 3. Find the smallest integer $\ell \in \mathcal{M}$ such that $\mathcal{Q}_\ell(\psi_{q_\ell}) = \min_{i \in \mathcal{M}} \mathcal{Q}_i(\psi_{q_i})$.
- Step 4. The identified parameters are $\alpha_\ell(\psi_{q_\ell})$, $\tau_{\text{ff}} \triangleq \psi_{q_\ell}$, $\beta_\ell \triangleq E_\beta \phi_\ell$, and $\tau_{\text{fb}} \triangleq E_\gamma \phi_\ell$.
- Step 5. The identified transfer functions are given by (7).

APPENDIX B

DESCRIPTION OF CANDIDATE POOLS FOR SSID

For each trial, we use Algorithm 1 to identify the best-fit model of the subject's control (3). 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, G_{fb} is modeled as a second-order strictly proper transfer function (i.e., $n_{\text{fb}} = 1$ and $d_{\text{fb}} = 2$). We select $n_{\text{ff}} = d_{\text{ff}}$ large enough to allow G_{ff} to approximate G^{-1} with approximately 0.1% error over the 0-to-0.5 Hz range. Thus, $n_{\text{ff}} = d_{\text{ff}} = 2$.

The feedback candidate pool Φ is designed to capture a wide range of behavior over the 0-to-0.5 Hz range and contains approximately one billion elements. The feedback 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) $\max_{\omega \in [0, \pi]} |G_{\text{fb}}(e^{j\omega T_s})| \leq 30.5$.
- C4) If $\lambda \in \mathbb{C}$ is a pole of $\tilde{G}_{\text{yr}}$, then $|\lambda| < 0.998$.
- C5) For all $\phi \in \Phi$, $E_\gamma \phi \in \{4, 5, 6, \dots, 25\}$.

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. 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. A continuous-time pole or zero with magnitude greater than 31.5 rad/s has negligible effect on the Bode plot over the range $(0, \pi]$ rad/s. Thus, we restrict Φ to include only 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 the frequency range $(0, \pi]$ 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 [23] 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. This range is consistent with [23], [70], [71].

The feedforward-delay candidate pool is $\Psi = \{0, 1, 2, \dots, 25\}$, which restricts the feedforward time delay to the range of [0, 500] ms.

Algorithm 1 is coded in C++ for parallel computation and implemented on the University of Kentucky High Performance Computing Cluster. For each trial, it takes approximately 3 h to run Algorithm 1 on one compute node; each node has a 16 Intel E5-2670 @ 2.6 GHz cores.

APPENDIX C CONDITIONING OF SSID RESULTS

Define the condition number

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

which is a measure of the conditioning of the identified feedback controller $z^{-\tau_{fb}} G_{fb}$. A larger value of $\mathcal{C}(G_{fb}, \tau_{fb})$ indicates a more poorly conditioned identified $z^{-\tau_{fb}} G_{fb}$. Figures 24–27 show $\mathcal{C}$ for each SSID result of groups 1, 2, 3, and 4, respectively; the SSID results are organized from the largest to smallest $\mathcal{C}$. For group 1, the largest $\mathcal{C}$ is approximately 32.0, whereas the smallest $\mathcal{C}$ is approximately 0.066. For group 2, the largest $\mathcal{C}$ is approximately 15.5, whereas the smallest $\mathcal{C}$ is approximately 0.352. For group 3, the largest $\mathcal{C}$ is approximately 16.0, whereas the smallest $\mathcal{C}$ is approximately 0.073. For group 4, the largest $\mathcal{C}$ is approximately 11.4, whereas the smallest $\mathcal{C}$ is approximately 0.260. Figures 24–27 also show that there is limited correlation between the trial




Fig. 24. Condition number $\mathcal{C}$ for the group 1.




Fig. 25. Condition number $\mathcal{C}$ for the group 2.

Figures 28–31 show the Bode plots of $z^{-\tau_{fb}} G'_{fb}$ for the SSID result from each group with the largest $\mathcal{C}$. For each group, this SSID result has a large $\mathcal{C}$ because $z^{-\tau_{fb}} G_{fb} \approx$




Fig. 27. Condition number $\mathcal{C}$ for the group 4.

$-G^{-1}$ at frequencies below approximately 0.1 Hz. However, $z^{-\tau_{fb}} G_{fb}$ is not approximately equal to $-G^{-1}$ at frequencies above 0.1 Hz.

Figures 28–31 also show the Bode plots of $z^{-\tau_{fb}} G_{fb}$ for the SSID result from each group with the 133rd largest $\mathcal{C}$ (i.e., the 30th percentile). For the group 1 SSID result with the 133rd largest $\mathcal{C}$, $z^{-\tau_{fb}} G_{fb}$ is not approximately equal to $-G^{-1}$ at any frequencies over the 0-to-0.5 Hz range. For the group 2 SSID result with the 133rd largest $\mathcal{C}$, $z^{-\tau_{fb}} G_{fb}$ is approximately equal to $-G^{-1}$ at frequencies below 0.02 Hz but is not approximately equal to $-G^{-1}$ at other frequencies over the 0-to-0.5 Hz range. For the group 3 SSID result with the 133rd largest $\mathcal{C}$, $z^{-\tau_{fb}} G_{fb}$ is approximately equal to $-G^{-1}$ at frequencies below 0.02 Hz but is not approximately equal to $-G^{-1}$ at other frequencies over the 0-to-0.5 Hz range. For the group 4 SSID result with the 133rd largest $\mathcal{C}$, $z^{-\tau_{fb}} G_{fb}$ is not approximately equal to $-G^{-1}$ at any frequencies over the 0-to-0.5 Hz range.

We examine the sensitivity of the SSID results in Section V to the conditioning of the identified $z^{-\tau_{fb}} G_{fb}$ by removing the most ill-conditioned 30% of the SSID results (i.e., 132 trials). Figures 32–35 show the mean and standard deviation of T_{fb} , T_{ff} , $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$, and $\|z^{-\tau_{fb}} G_{fb}\|_1$ on each trial for each group. The trends observed in Figs. 32–35 are the same as




Fig. 28. Bode plots of $z^{-\tau_{fb}}G_{fb}$ for the group 1 SSID results with the largest condition number and the 133rd largest condition number.




Fig. 29. Bode plots of $z^{-\tau_{fb}}G_{fb}$ for the group 2 SSID results with the largest condition number and the 133rd largest condition number.




Fig. 30. Bode plots of $z^{-\tau_{fb}}G_{fb}$ for the group 3 SSID results with the largest condition number and the 133rd largest condition number.

those observed in Figs. 14–16 and 22. The same observations hold if the most ill-conditioned 10%, 15%, 20%, 25%, or 30% of the results are omitted.




Fig. 31. Bode plots of $z^{-\tau_{fb}}G_{fb}$ for the group 4 SSID results with the




Fig. 32. Mean and standard deviation of T_{fb} on each trial. Plots omit the most ill-conditioned 30% of trials. The $\circ$ is the mean, and the lines indicate




Fig. 33. Mean and standard deviation of T_{ff} on each trial. Plots omit the most ill-conditioned 30% of trials. The $\circ$ is the mean, and the lines indicate the standard deviation.

APPENDIX D VALIDATION OF SSID RESULTS

For each trial, we simulate the identified closed-loop system, where the input to the simulation is $\{r_k\}_{k=1}^n$, and the output of the simulation is the validation data $\{y_{v,k}\}_{k=1}^n$. Specifically, we simulate

$$\hat{y}_v(z) = \tilde{G}_{yr}(z)\hat{r}(z),$$




Fig. 34. Mean and standard deviation of $\|z^{-\tau_{ff}} G_{ff} - G^{-1}\|_1$ on each trial. Plots omit the most ill-conditioned 30% of trials. The $\circ$ is the mean, and the




Fig. 36. Mean and standard deviation of VAF on each trial. The $\circ$ is the mean, and the lines indicate the standard deviation.




Fig. 35. Mean and standard deviation of $\|z^{-\tau_{fb}} G_{fb}\|_1$ on each trial. Plots omit the most ill-conditioned 30% of trials. The $\circ$ is the mean, and the lines indicate the standard deviation.

where all initial conditions are zero, $\hat{y}_v(z)$ is the z -transform of the validation data $y_{v,k}$, and $\hat{G}_{yr}$ is the closed-loop transfer function (4) obtained from the identified G_{ff} , τ_{ff} , G_{fb} , and τ_{fb} .

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 $n_1 = 26$. Note that VAF is calculated using data from the time interval $(0.5, 60]$ s. We omit the interval $[0, 0.5]$ s to reduce the impact of nonzero initial conditions. The validation data is computed with zero initial conditions; however, the experimental data may have nonzero initial conditions.

Figure 36 shows the mean and standard deviation of the VAF for each trial. For all groups, the mean VAF over the last 5 trials is greater than that over the first 5 trials. Specifically, the mean VAF over the last 5 trials for groups 1, 2, 3, and 4 is 0.90, 0.80, 0.75, and 0.73, respectively. The mean VAF over the first 5 trials for groups 1, 2, 3, and 4 is 0.73, 0.63, 0.62, and 0.48, respectively. Thus, the subjects' control behavior can be modeled more accurately by the low-order LTI controller (3) on the later trials than the earlier trials. This observation