Adaptive Manual Control: a Predictive Coding Approach

Terenzi, L.; Zaal, P.M.T.; Pool, D.M.; Mulder, Max

DOI

[10.2514/6.2022-2448](https://doi.org/10.2514/6.2022-2448)

Publication date

2022

Document Version

Final published version

Published in

AIAA SCITECH 2022 Forum

Citation (APA)

Terenzi, L., Zaal, P. M. T., Pool, D. M., & Mulder, M. (2022). Adaptive Manual Control: a Predictive Coding Approach. In *AIAA SCITECH 2022 Forum* [AIAA 2022-2448] (AIAA Science and Technology Forum and Exposition, AIAA SciTech Forum 2022). <https://doi.org/10.2514/6.2022-2448>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Adaptive Manual Control: a Predictive Coding Approach

Lorenzo Terenzi*

Delft University of Technology, Delft, Zuid Holland, 2600 HS, Netherlands

Peter M. T. Zaal†

Metis Technology Solutions, Inc., NASA Ames Research Center, Moffett Field, CA, 94035

Daan M. Pool‡ and Max Mulder§

Delft University of Technology, Delft, Zuid-Holland, 2600 HS, Netherlands

Improved understanding of human adaptation can be used to design better (semi-)automated systems that can support the human controller when task characteristics suddenly change. This paper evaluates the effectiveness of a model-based adaptive control technique, Model Reference Adaptive Control (MRAC), for describing the adaptive control policy used by human operators while controlling a time-varying system in a pursuit-tracking task. Ten participants took part in an experiment in which they controlled a time-varying system whose dynamics changed twice between approximate single and double integrator dynamics, and vice versa. Our proposed MRAC controller is composed of a feedforward and a feedback controller and an internal reference model that is used to drive an adaptive control policy. MRAC's adaptive control gains, the internal model parameters, and the learning rates were estimated from the experiment data using non-linear optimization aimed at maximizing the quality-of-fit of participants' control outputs. Participants' control behavior rapidly changed when the dynamics of the controlled system changed, in particular for transitions from single to double integrator dynamics. The MRAC model was indeed able to accurately capture the transient dynamics exhibited by the participants when the system changed from an approximate single to a double integrator, however, for the opposite transition the MRAC gains were always adapted too slowly. Therefore, in its current form, our MRAC model can be used to approximate human adaptation in pursuit tracking tasks when a change in the dynamics of the controlled system requires significant (rate) feedback controller adaptation to maintain satisfactory closed-loop control performance.

Nomenclature

A	=	State space matrix, –
A_m	=	Reference model state space matrix, –
A_n	=	Amplitude of the n^{th} sine of the forcing function, rad
B	=	Control input matrix, –
B_m	=	Reference model control input matrix, –
e	=	Roll attitude error, rad
H_c	=	Controlled dynamics transfer function
H_p	=	Human operator transfer function
H_{mCL}	=	Internal model closed-loop transfer function
H_{mOL}	=	Internal model open-loop transfer function
$k_c(t)$	=	Time varying gain of the controlled element, –

*MSc. Graduate, section Control and Simulation, Faculty of Aerospace Engineering, Delft University of Technology, Delft, Netherlands; lorenzoterenzi96@gmail.com.

†Principal Aerospace Engineer, SimLabs, NASA Ames Research Center, Moffett Field, CA, 94035; peter.m.t.zaal@nasa.gov. Associate Fellow AIAA.

‡Assistant Professor, section Control and Simulation, Faculty of Aerospace Engineering, P.O. Box 5058, 2600 GB Delft, Netherlands; d.m.pool@tudelft.nl. Senior Member AIAA.

§Professor, section Control and Simulation, Faculty of Aerospace Engineering, P.O. Box 5058, 2600 GB Delft, Netherlands; m.mulder@tudelft.nl. Associate Fellow AIAA.

k_r	=	MRAC feedforward gain, –
K_x	=	Vector of state gains, –
k_{x1}	=	MRAC state gain, –
k_{x2}	=	MRAC state derivative gain, s
lr	=	Learning rate, –
r	=	Forcing function, rad/s
P	=	Positive definite matrix, –
t	=	Time, s
x_m	=	Reference model state
u_h	=	Human operator control output, rad
u_m	=	MRAC model control output, rad
VAF	=	Variance accounted for, %
τ	=	Human operator time delay, s
θ	=	MRAC parameter vector, –
ϕ_n	=	Phase of the n^{th} sine of the disturbance signal, rad
γ_r	=	Learning rate of the feedforward gain, –
Γ_x	=	Vector of state gain's learning rate, –
γ_{x1}	=	Learning rate of the state gain, –
γ_{x2}	=	Learning rate of the state derivative gain, –
ω	=	Frequency, rad/s
$\omega_b(t)$	=	Time-varying pole of the controlled element, –
ω_c	=	Crossover frequency of the internal model, rad/s
ω_n	=	Frequency of the n^{th} sine of the disturbance signal, rad/s

I. Introduction

In the last decades, automation has taken an increasingly important role in systems with humans in the loop, such as aircraft or road vehicles. Tasks in relatively controlled environments are easiest to automate. For example, most parts of a commercial flight have been automated because it is possible to rely on precise positioning systems, sensor readings, lack of obstacles, and preestablished navigation plans. The role of airline pilots has become more that of a supervisor of a myriad of automated subsystems. Current automation can achieve a human level of performance in perception, or better, in many tasks. For example, superhuman performance in image classification on the ImageNet dataset was achieved five years ago [1].

Nonetheless humans are still a crucial element in driving cars and piloting aircraft today, because of their ability to quickly adapt to a changing environment [2] and to carry out tasks that are complex and difficult to automate. Whereas the pilot equalization to fixed aircraft dynamics is well known [3], the adaptation to changing dynamics is still unexplored. Mulder et al. [4], in their state-of-the-art in human control behavior modelling paper, conclude that a good understanding of how humans adapt their control policy could serve as inspiration for machine learning researchers and automation engineers to develop more advanced controllers, capable of exhibiting more adaptable behavior. Furthermore, it could help design better training programs for human controllers and more suitable (semi-)automated support systems to increase safety.

This paper aims to increase the understanding of human motor adaptation for simple manual control tasks with a novel approach inspired by a neuroscience framework called ‘*Predictive Coding*’ (PC) [5]. Our approach is centered on an adaptive control technique called Model Reference Adaptive Control (MRAC), which uses an explicit reference ‘internal model’ to detect off-nominal behavior and drive controller adaptation. Inspired by McRuer and Jex’s [6] *crossover model*, which implies that the open-loop dynamics of a manually controlled system approximate a single integrator with a time delay independent of the dynamics of the controlled system, in this paper we implement the crossover model as the reference model in an MRAC controller. Furthermore, we focus on human control adaption in response to changes in controlled system dynamics matching a series of recent experiments [7, 8, 9] in a pursuit tracking task. We present a sensitivity analysis of the MRAC parameters and a human-in-the-loop experiment to determine the feasibility of using MRAC to predict time-varying human control adaptations.

Section II provides a background of research in three fundamental areas used in this paper: internal models, predictive coding and time-varying/adaptive human operators models. In Section III, the MRAC controller implementation is presented in detail. The control task and the setup of the experiment are discussed in Section IV. Section V provides the

stability and sensitivity analyses of the MRAC controller. In Section VI, the experimental results are analysed and presented and their discussion follows in Section VII. Finally, conclusions are drawn in Section VIII.

II. Background

A. Internal Models

This section presents an overview of internal models and we try to justify their use in the development of adaptive and optimal control systems. Internal models are widely recognized as an essential component for motor control in biological systems [10]. Animals have internal models of their body's and environment's dynamics that can be used for motor control and learning [11, 12].

It is commonly accepted that the cerebellum learns and encodes models of the world [13]. These models can encode the dynamics and properties of the human motor system and of objects that humans interact with. Patients with cerebellum damage experience very poor motor control and jittering in their movement and have problems learning new skills [13, 14]. The imprecise motor control in these patients suggests that their brains cannot make use of internal models of the motor system to find optimal motor commands and anticipate probable future states. In this case, the brain will have to rely only on sensory feedback (such as visual feedback), which is severely delayed and does not allow for precise and fast motor control [15].

The exact nature of the models used for sensorimotor control is currently not known, but it is hypothesized they could be *inverse* and/or *forward* dynamical models [16]. Inverse models output the control action needed to achieve a desired output, while forward models predict the output given an efference copy of the control action. On the other hand, given a probability distribution of the current states and a control input, forward internal models can generate a distribution of potential future states. There is substantial anatomical and behavioral evidence that points towards the presence of such forward internal models in the cerebellum [17, 18].

Forward internal models have two main uses. The first one is to plan complex actions without acting them out in the real world, by simulating the results of a control policy through the forward model. This way of controlling is compatible with Optimal Feedback Control (OFC), a control method that makes use of the dynamics of a system to find the set of optimal actions to define a feedback controller [19, 20, 21]. The OFC framework has shown predictive strength in many experiments testing sensorimotor control and in particular hand-eye coordination [10]. Currently, it is the dominant framework to model sensorimotor control [22, 23]. The second main use for internal forward models is state estimation: the distribution over the current states (given the previous state and a control action) is compared with sensory measurements to achieve the best estimate of the current state of the system [10]. A diagram that illustrates how internal models can be used for state estimation and planning is shown in Fig. 1. The state estimator makes use of the sensory measurements, z , and the predicted state x_p by the internal model to obtain a more reliable estimation. The expected state estimate, x_e can be fed back to the internal model. The internal model can be used to generate rollout (sequence of predicted states x_p) given a policy (set of control inputs u) as part of an optimal control scheme.

The last piece of evidence we present here in favour of internal models is related to the experiment with a compensatory manual control task described by Young [2]. Participants were asked to control and detect the change of dynamics in a time-varying system. The participants of the experiment were divided into three groups: active controllers, passive controllers and observers. The active controllers actively controlled the system. The passive controllers were led to believe that they controlled the system, but the system was in fact controlled by a different participant. The observers just observed the screen. All participants had to press a button when they recognized a change in controlled system dynamics. It was observed that the detection time across all conditions for the observers was significantly higher (mean of 3.74 s) than the other two classes of participants, who showed similar detection times (mean of 1.30 s for active controllers and of 1.50 s for passive ones). The most probable cause for this delay for the observer is the absence of an internal model for the dynamics of the controlled system. The participants that only observed were not engaged in executing motor actions and therefore had a harder time estimating how the system would respond to input. This highlights the role that internal models have in motor control tasks.

B. Predictive Coding

Predictive coding (PC) is a general framework to explain how the brain processes information [5]. PC states that the human brain continuously generates predictions about the stream of information that is coming from the world, and its objective is to minimize the prediction error [21], and thus surprise. The minimization of the prediction error can be

Fig. 1 Illustration of how sensory data and internal models can be used for state estimation, control and planning.

achieved by either updating the internal model that provides the predictions (‘learning’) or by adapting the actions of the agent. The PC theory, which is very formal in nature, is further explained by Friston in a series of papers [24, 25]. It is postulated that the brain has an internal generative model of the world. This is a model that can generate a rich and hierarchical representation of the expected/predicted sensory information from a latent, non-observable variable [25]. For example, from a single “idea” of an object (the latent variable), it is possible to generate richer representations in the lower hierarchical cortical area, such as its shape, sound, and expected dynamic properties. The process that generates the sensory predictions is “top-down”, from the higher level of the cortical areas to the lower ones, while the sensory information is processed “bottom-up” as illustrated in Fig. 2.

The sensory predictions at the very bottom of the abstraction hierarchy are compared with the actual sensory data. The unexplained sensory inputs, i.e., the “errors”, are propagated up the hierarchy for further processing. These prediction errors are minimized either by acting on the world or by changing the internal model over time [27]. An essential part of the PC theory is *precision weighting*, i.e., the estimation of the error’s reliability, which also depends on the level of signal noise [5].

This theory of neural computation is supported by anatomical evidence and there are canonical microcircuit models of PC [28]. As PC requires, there are different cells that are responsible for feedforward and feedback information [5, 28], and these cells show interactions at different frequencies [29]. Computational models of the visual cortex explain many of the unsolved visual phenomena like end-stopping and non-classical surround effects [30]. Similar results are reported for the auditory cortex [31, 32]. Nonetheless, while it is recognized that the ability to predict is a key element of the human type of intelligence, there is still no consensus in the scientific community about PC: criticism highlights that the theory is quite imprecise, difficult to test, and therefore susceptible to *ad hoc* changes [33].

C. Adaptive Human Control Models

Studies on modeling adaptive human manual control have mostly focused on direct *identification* of time-varying human control dynamics, or developing rule-based *prediction* models of when and how humans will adapt. In time-varying human operator *identification*, methods have been developed that enable identifying the time-varying adaptive policies that human controllers use in response to a change in controlled system dynamics. A very flexible approach is based on Kalman filters, which are used to recursively estimate parameters of human operator models [34, 35, 36]. The main drawback of this method is the speed of convergence of the recursive estimator: in the event of a sudden change in the operator, the estimated parameters may be unreliable due to lag in the Kalman filter’s readjustment of its internal covariance matrices.

Another promising line of research uses ARX-based identification techniques [8, 9]. ARX models rely on minimal assumptions about human operator behavior and can be extended for multiple modalities, such as control in the presence of motion feedback [37, 38]. As an alternative to such recursive estimators, Zaal [7] has proposed to use a time-varying parametric human control model, whose parameter variations are all described by an assumed time-varying function (e.g., sigmoid centred at the time of the dynamics transition). The model parameters were found by maximum likelihood

Fig. 2 Simplified scheme of a predictive coding framework, adapted from [26]. The input enters a network region into the superficial cortical layers (in red). There the input is processed together with the prediction, coming from a deep cortical layer (blue) next in the hierarchy. The resulting error is further propagated up the hierarchy and the process repeats.

estimation (including the transition time of the dynamics) augmented with a genetic algorithm [39]. The main drawback of this approach is that the transient dynamics are a priori assumed and thus predetermined, and cannot handle more than one transition in dynamics.

Since the 1960s, several qualitative and quantitative rule-based models have been proposed for *predicting* human control adaptations, see, for example, the review articles by Young [2] and Xu et al. [40]. For example, Phatak and Bekey [41] proposed their ‘*Supervisory Control Model*’, which can model how operators may detect a change in controlled dynamics from off-nominal tracking errors and error rates in compensatory tracking and trigger adaptation of their control strategy. Stark and Young [42] proposed a similar qualitative model for the human controller adaptation that also proposes operators rely on mismatches in tracking error and its rate, but in comparison with an explicit (and adaptive) internal model of the dynamics of the controlled system. Most recently, Hess [43] has proposed an extension to his ‘*Structural Model*’ that relies on a heuristically developed scheme that compares current tracking performance to a nominal state to drive a rule-based detection of changing conditions and adaptive learning rules that change the human operator’s control gains to new, appropriate, settings.

III. Model Reference Adaptive Control

A. Motivation

This study uses Model Reference Adaptive Control (MRAC), an adaptive control technique that relies on an explicit internal model, to predict time-varying human operator adaptations in tracking tasks with a pursuit display. MRAC is first introduced, then the selection of the internal model, the architecture of the controller, and the considered parameter estimation method. MRAC was thought to be a good candidate to model time-varying human-operator behavior for the following reasons:

- 1) Humans adapt in the presence of mismatched predictions: PC states that any prediction errors, i.e., differences

between predicted and observed states, are propagated up the neuronal hierarchy to drive actions and/or change the encoded internal models. MRAC works similarly, since the adaptation of controller parameters is driven by the difference between the predicted output of the internal model and the observed output [2, 41, 43].

- 2) Like a human controller, MRAC has an internal model: while controlling the system, and receiving mostly visual feedback from the screen, humans can learn the dynamics of the controlled element. This internal model is used for control. Similarly, MRAC uses its internal model for control by being the main drive behind the adaptation.
- 3) Humans in the loop show approximately equal open-loop dynamics when controlling a wide range of dynamical systems: McRuer and Jex [6] found that the open-loop dynamics in the crossover region typically resemble those of a single integrator with a time delay. In an analogous way, MRAC can use a constant internal reference model to define the ideal control policy independently of the controlled dynamics.

B. Mathematical Development

The mathematical details of MRAC and its formulation are explained here. For more detailed examples and further derivations of MRAC controllers, please refer to [44]. Fig. 3 shows a diagram of the implementation of the MRAC controller as considered in this paper. The gains of the controller are $K_x = [k_{x1}, k_{x2}]$ (feedback gains on the controlled element output x and rate $\dot{x}$, consistent with the human feedback control of single and double integrator-like systems [7, 9]) and k_r (a feedforward gain on the reference input r , which is appropriate for a pursuit task [45]). The controlled element dynamics are indicated as H_c , the closed-loop reference model is H_{mCL} and the prediction error that drives MRAC's adaptation is indicated as e_p .

Fig. 3 Diagram of the MRAC human control model.

For deriving the adaptive control laws, as also indicated in Fig. 3, we first derive an expression for the dynamics of the error between the output of the reference model and the output of the controlled system, see Fig. 3. First, it is assumed that the reference model can be written in state-space form as:

$$\dot{x}_m = A_m x_m + B_m r \quad (1)$$

where x_m is the state of the reference model and r is the reference signal. The controlled dynamics can also be expressed in state-space form as:

$$\dot{x} = Ax + Bu \quad (2)$$

where u is the control input, see Fig. 3. A direct observation equation is assumed such that the output $y = x$. The state-space matrix A and the control input matrix B are unknown, which reflects the uncertainty about the dynamics of the system. The aim of the control action is to minimize the prediction error between the internal model's output and real output, i.e., $e_p = x_m - x$. The feedback control signal is assumed to have the following form:

$$u = K_x(t)x + k_r(t)r \quad (3)$$

where $K_x(t) = [k_{x1}(t), k_{x2}(t)]$ and $k_r(t)$ are all time-varying gains. MRAC assumes the existence of ideal gains, K_x^* and k_r^* , that can be found if Eq. (3) is substituted into Eq. (2) and the result is compared to Eq. (1). These ideal gains

must satisfy:

$$\begin{aligned} A + BK_x^* &= A_m \\ Bk_r^* &= B_m \end{aligned} \quad (4)$$

To simplify the derivation, two new quantities $\tilde{K}_x$ and $\tilde{k}_r$ can be defined, which are the gain deviations from their optimal values:

$$\begin{aligned} \tilde{K}_x &= K_x(t) - K_x^* \\ \tilde{k}_r &= k_r(t) - k_r^* \end{aligned} \quad (5)$$

By combining Eqs. (2), (3), and (5) the following expression for the error can be found:

$$\dot{e}_p = \dot{x}_m - \dot{x} = A_m e_p - B\tilde{K}_x x - B\tilde{k}_r r \quad (6)$$

MRAC relies on the Lyapunov stability theory to prove the stability of the system. To ensure tracking of the reference model and therefore that $\lim_{t \rightarrow \infty} e_p(t) = 0$ the following conditions need to be satisfied:

- 1) The existence of a Lyapunov function $V(t, e_p, \tilde{K}_x, \tilde{k}_r) > 0, \forall t > 0$
- 2) $\frac{dV(t, e_p, \tilde{K}_x, \tilde{k}_r)}{dt} < 0$ for all $t > 0$
- 3) $\frac{dV(t, e_p, \tilde{K}_x, \tilde{k}_r)}{dt} \in L_\infty$ norm, i.e., $\frac{d^2 V(t, e_p, \tilde{K}_x, \tilde{k}_r)}{dt^2}$ must be bounded

There are no direct guidelines to define the Lyapunov function, but usually it is chosen as a quadratic function with respect to the variables of interest [44]. It can also be interpreted as an energy function. In this case the following function was chosen:

$$V(e_p, \tilde{K}_x, \tilde{k}_r) = e_p^T P e_p + |b|(\tilde{K}_x \Gamma_x^{-1} \tilde{K}_x^T + \frac{\tilde{k}_r^2}{\gamma_r}) > 0 \quad (7)$$

where b is the only entry of the control effectiveness matrix B . With $P, \Gamma_x^{-1}, \gamma_r > 0$, the function V is also above zero at all times. Note that this derivation is valid for a second-order system as we assumed the control matrix is of the form $B = [0; b]$. By taking the time derivative of the Lyapunov function and using Eq. (6), the following is obtained:

$$\begin{aligned} \dot{V}(e_p, \tilde{K}_x, \tilde{k}_r) &= -e^T (PA_m + A_m P) e_p + 2|b|\tilde{K}_x (-x e_p^T \bar{P} \text{sign}(b) + \Gamma_x^{-1} \dot{\tilde{K}}_x^T) \\ &\quad + 2|b|\tilde{k}_r (-r e_p^T \bar{P} \text{sign}(b) + \frac{\dot{\tilde{k}}_r}{\gamma_r}) \end{aligned} \quad (8)$$

and by selecting P to satisfy the Lyapunov equation:

$$PA_m + A_m^T P = -Q \quad (9)$$

it can be found that:

$$\begin{aligned} \dot{V}(e_p, \tilde{K}_x, \tilde{k}_r) &= -e_p^T Q e_p + 2|b|\tilde{K}_x (-x e_p^T \bar{P} \text{sign}(b) + \Gamma_x^{-1} \dot{\tilde{K}}_x^T) \\ &\quad + 2|b|\tilde{k}_r (-r e_p^T \bar{P} \text{sign}(b) + \frac{\dot{\tilde{k}}_r}{\gamma_r}) \end{aligned} \quad (10)$$

where Q is a negative definite matrix and $\bar{P}$ the second column of the matrix P (since it is assumed that matrix B has only one entry equal to b). For the Lyapunov function derivative to be negative at all times the following conditions need to be imposed:

$$\begin{aligned} -x e_p^T \bar{P} \text{sign}(b) + \Gamma_x^{-1} \dot{\tilde{K}}_x^T &= 0 \\ -r e_p^T \bar{P} \text{sign}(b) + \frac{\dot{\tilde{k}}_r}{\gamma_r} &= 0 \end{aligned} \quad (11)$$

This directly implies:

$$\begin{aligned} \dot{\tilde{K}}_x &= \Gamma_x x e_p^T \bar{P} \text{sign}(b) \\ \dot{\tilde{k}}_r &= \gamma_r r e_p^T \bar{P} \text{sign}(b) \end{aligned} \quad (12)$$

In this way an expression is derived for the rate of change of the feedback gain that would ensure the tracking of the reference model. Finally, it is necessary to show that the derivative of the Lyapunov function is bounded. The derivative of the Lyapunov function satisfies the following inequality:

$$\dot{V}(e_p, \tilde{K}_x, \tilde{k}_r) = -e_p^T Q e_p \leq -\lambda_{\min} \|e_p\|_2^2 \quad (13)$$

where $\lambda_{\min}$ is the smallest eigenvalue of the matrix Q . Since $\|e_p\|_2^2 \in L_\infty$ it can be shown that $\dot{V}(e_p, \tilde{K}_x, \tilde{k}_r)$ is bounded. By using Barbalat's lemma [44], it can be concluded that $\lim_{t \rightarrow \infty} e_p(t) = 0$.

In the previous derivation, it was assumed that there was no delay in the control input to derive a simple adaptive control law. Humans in reality control with considerable delays in perception and actuation that should be accounted for. The presence of delays can severely affect the performance and stability of a linear state feedback controller. An adaptive state feedback controller, such as MRAC, is also affected by such delays. The system controlled by the human operator can therefore be written more formally as:

$$\dot{x}(t) = A(t)x(t) + B(t)u(t - \tau) \quad (14)$$

$$- + u(t) = K_x x(t) + k_r r(t) \quad (15)$$

where τ stands for the input delay. This paper does not provide any formal guarantees on the stability of the considered MRAC controller, but techniques exist that establish its bounded (not asymptotic) stability, such as Bounded Linear Stability Analysis (BLAS) [46] and other techniques for time delay margin estimation [47, 48]. More generally, the design of stable controllers and their analysis in presence of delays and uncertainties has been reported with Lyapunov-Krasovskii functional techniques [49]. A comprehensive tutorial and review of Lyapunov-methods for time-delayed systems was written by Fridman [50].

C. Internal Model Selection

A key choice for every MRAC controller is the internal model. While MRAC assumes what in manual control research is known as a 'pursuit' control structure (where the reference signal is available as a separate input variable), the selected reference model proposed in this paper is the *crossover model* proposed by McRuer and Jex [6], which describes the combined human-operator controlled-dynamics open-loop dynamics in a compensatory task. It is assumed that this model is also valid to describe the open-loop dynamics in pursuit tasks. Further information on how pursuit tracking differs from compensatory tracking can be found in the review by Mulder et al. [51]. The crossover model in the Laplace domain has two free parameters: the effective time delay, τ , and the crossover frequency, ω_c :

$$H_{mOL}(s) = \frac{\omega_c}{s} e^{-\tau s} \quad (16)$$

For its internal reference model, our MRAC human controller model (see Fig. 3) makes use of the corresponding closed-loop dynamics:

$$H_{mCL}(s) = \frac{H_{mOL}(s)}{1 + H_{mOL}(s)} \quad (17)$$

The time delay τ and the crossover frequency ω_c of the model are defined *a priori*, i.e., based on estimates obtained from available human controller experiment data, e.g., from those obtained by McRuer and Jex [6].

D. Parameter Estimation

The MRAC human control model introduced in Fig. 3 and this section has a number of parameters that can be tuned to approximate human adaptation as closely as possible. The following parameters need to be determined to parameterize the MRAC controller:

$$\theta = [k_{x1}, k_{x2}, k_r, \omega_c, \tau, \gamma_{x1}, \gamma_{x2}, \gamma_r] \quad (18)$$

Here, k_{x1} is the state gain, k_{x2} is the state derivative gain, k_r is the feedforward gain, while ω_c and τ are the parameters of the internal reference model, see Eq. (16). The learning rate parameters γ_{x1} , γ_{x2} , and γ_r correspond to the adaptation of k_{x1} , k_{x2} and k_r , respectively. These parameters can be found solving the following non-linear optimization problem:

$$\arg \min_{\theta} J(\theta) = \frac{1}{N} \left(\sum_{i=0}^{N-1} (u_h - u_m)^2 \right) \quad (19)$$

where u_m stands for the MRAC controller output and u_h the measured output of the human controller. Similar time-domain procedures have been applied in the past to estimate the gains of human operators, e.g. [39, 7]. Since this is a non-linear optimization scheme, care must be taken when initializing the parameters to minimize the chance of finding a local minimum. For each case, the optimization was run fifteen times, each time with a different initial condition. The set of parameters that resulted in the minimum value of the objective function were selected. Additionally, since the parameters have different scales, they were normalized to improve the conditioning of the optimization problem. Finally, an upper bound on the learning rate (in this case a value of 50 was used for all learning rates) improves convergence as the optimization may diverge if very high learning rates are selected by the optimizer.

IV. Experiment Setup

A. Control Task

Fig. 4 shows the block diagram of the pitch attitude pursuit control task considered in this paper. This task was similar to earlier time-varying compensatory control experiments [7, 8, 9]. The human operator's objective is to minimize the error between the reference pitch attitude r and the current pitch attitude x , which could be inferred from the visual display, see Fig. 4. A pursuit task was chosen because the MRAC model uses the reference signal as an input and thus its adaptive mechanism could directly mirror how human operators adapt in this task. The pitch dynamics of the controlled system are given by $H_c(s, t)$. The quasi-linear human operator control dynamics in the task are represented by the time-varying control dynamics $H_p(s, t)$ and a remnant signal n [4, 6].

Fig. 4 Control task block diagram.

B. Controlled Dynamics

Participants controlled the system with time-varying dynamics $H_c(s, t)$ across multiple clear changes in the system dynamics matching those also tested in earlier experiments [7, 8, 9]. The controlled system dynamics $H_c(s, t)$ had the following transfer function form:

$$H_c(s, t) = \frac{k_c(t)}{s(s + \omega_b(t))} \quad (20)$$

where the parameters $k_c(t)$ and $\omega_b(t)$ were varied over time according to sigmoid functions. By modifying the parameter ω_b the system response changed either towards a single-integrator-like response (for $\omega_b \gg 1$ rad/s) or a double-integrator-like response (for $\omega_b \ll 1$ rad/s). The gain k_c varied over time to keep the level of control activity approximately constant [7]. In the current experiment, time-varying tracking runs all included two controlled system transitions, for which the mathematical definition of the ω_c and k_c variations was the following:

$$\omega_b(t) = \begin{cases} \omega_{b1} + \frac{\omega_{b2} - \omega_{b1}}{1 + e^{-G(t-M_1)}}, & \text{for } t \leq \frac{T}{2} \\ \omega_{b2} + \frac{\omega_{b1} - \omega_{b2}}{1 + e^{-G(t-M_2)}}, & \text{for } t \geq \frac{T}{2} \end{cases} \quad (21)$$

$$k_c(t) = \begin{cases} k_{c1} + \frac{k_{c2} - k_{c1}}{1 + e^{-G(t-M_1)}}, & \text{for } t \leq \frac{T}{2} \\ k_{c2} + \frac{k_{c1} - k_{c2}}{1 + e^{-G(t-M_2)}}, & \text{for } t \geq \frac{T}{2} \end{cases} \quad (22)$$

where $T = 90$ s is the length of a measurement run, and $M_1 = \frac{T}{3}$ and $M_2 = \frac{2T}{3}$ are the times at which the dynamics change. The value of G , which controls the transition speed, was kept constant at a value of 100 s^{-1} , as also used in [7, 9], to simulate a step-like change in the controlled system's parameters.

As listed in Table 1, four different conditions were tested in the experiment. To reduce predictability, all four conditions were presented in a mixed order throughout the experiment, meaning that every next tracking run could be either a steady-state condition (DYN 1 or DYN 2) or a time-varying run (DYN 121 or DYN 212). For DYN 121, the participants controlled a time-varying system that for the first 30 s behaved approximately like a single integrator, after the transition at $M_1 = \frac{T}{3}$ approximately as a double integrator, and after the final transition at M_2 again as a single integrator in the final 30 s. In condition DYN 212 the order of the dynamics transitions was reversed.

Table 1 Different tested controlled dynamics settings.

DYN	ω_{b1} [rad/s]	ω_{b2} [rad/s]	k_{c1} [-]	k_{c2} [-]
1	6.0	6.0	90	90
2	0.2	0.2	30	30
121	6.0	0.2	90	30
212	0.2	6.0	30	90

C. Forcing Function

The target attitude was defined by the forcing function $r(t)$, see Fig. 4, which was defined as a sum of ten sinusoids to make the task challenging and the signal unpredictable. The quasi-random forcing function signal was defined by:

$$r(t) = \sum_{n=1}^{10} A_n \sin(\omega_n t + \phi_n) \quad (23)$$

where A_n is the amplitude of the n^{th} sine wave, ϕ_n the phase, and ω_n the frequency. While for this experiment, for which no spectral methods were used for human operator identification, the use of a sinusoidal tracking signal was not strictly required, it was still important to have an unpredictable signal to prevent participants from learning $r(t)$. The signal was also chosen to be periodic, with a period equal to $T_r = \frac{T}{3}$, one-third of the measurement time T , to be able to directly compare behavior and time traces across the three different 30-second run segments. The periodicity preserved the task difficulty, as participants were exposed to the same signal when transitioning from one dynamics to the other and vice versa. The parameters of the sinusoidal waves that make up the signal $r(t)$ are listed in Table 2.

Table 2 Parameters of the sinusoidal reference signal.

n	A_n [rad]	ω_n [rad/s]	Testing	Validation
			ϕ_n [rad]	ϕ_n [rad]
1	$2.905 \cdot 10^{-2}$	0.419	2.841	3.006
2	$1.916 \cdot 10^{-2}$	1.047	3.319	6.037
3	$1.020 \cdot 10^{-2}$	1.885	0.718	4.544
4	$6.032 \cdot 10^{-3}$	2.722	0.768	2.811
5	$3.356 \cdot 10^{-3}$	3.979	2.925	5.917
6	$1.983 \cdot 10^{-3}$	5.655	5.145	1.842
7	$1.230 \cdot 10^{-3}$	8.188	2.085	3.401
8	$9.331 \cdot 10^{-4}$	10.681	0.383	2.998
9	$7.541 \cdot 10^{-4}$	14.032	0.763	4.614
10	$6.674 \cdot 10^{-4}$	17.383	3.247	2.888

For the current experiment, measurement runs were divided into two groups to create two datasets: 1) a *testing* dataset used to estimate the parameters of the MRAC model and 2) a *validation* dataset for independent evaluation of fit quality and to prevent over-fitting. As shown in Table 2, the target signals used for the testing and validation data had the same frequencies and amplitudes, but used different sinusoid phases.

Finally, to ensure participants were not tracking exactly the same signals in all testing and validation tracking runs, for the DYN 212 condition (see Table 1) the tracking signal was mirrored, i.e., $-r(t)$ was used.

D. Apparatus

The experiment was performed in the Human-Machine Interaction Laboratory (HMI Lab) simulator at the Faculty of Aerospace Engineering at Delft University of Technology, see Fig. 5. The simulation software for the experiment was implemented in the in-house developed DUECA/DUSIME framework for real-time simulations [52]. The tracking display was presented on the simulator's primary flight display, directly in front of the participants, as shown in Fig. 5. The display was updated at a frequency of 60 Hz and had a latency of approximately 20-25 ms [53]. The control inputs were provided using an electro-hydraulic side stick located at the right hand side of the participant, which was configured to rotate only around the pitch axis. The stick torsional stiffness was 25 N m/rad, the damping coefficient was 0.22 N s/rad and its inertia was 0.01 kg m².

Fig. 5 Human-in-the-loop experiment setup.

E. Participants

Ten participants (7M, 3F) volunteered to perform the experiment. All participants were students or staff at Delft University of Technology. Since this study focused on understanding the general adaptability of human controllers, participants were not required to be (active) pilots. Only two participants had extensive prior experience with tracking tasks. All participants provided written informed consent prior to their participation. The study was approved by TU Delft's Human Research Ethics Committee under application number 1409.

F. Procedures

The participants were given a briefing on the experiment's procedures and the manual control task prior to performing the experiment. They were told that the goal of the experiment was to understand how humans adapt their control behavior when controlling time-varying dynamics. The participants were encouraged to provide continuous control inputs and at all times to follow the always moving target attitude indicator.

The experiment started with a familiarization phase, which was short and ended when the participants were able to stabilize the system across the different conditions. It was followed by a training phase, with the objective to bring participants to asymptotic performance in the tracking task. During the training phase the participants were exposed first to the steady-state conditions DYN 1 and DYN 2 and then to the conditions with the time-varying dynamics (DYN 121 and DYN 212). The training phase ended when participants achieved consistent performance for three consecutive runs in the time-varying conditions.

The measurement phase of the experiment consisted of a total of 22 runs. For the testing dataset, five runs each were collected for the DYN 121 and DYN 212 conditions; an additional three runs each were performed for the validation dataset. Furthermore, three additional runs for the steady-state conditions DYN 1 and DYN 2 were added to make the experiment less predictable, i.e., to not always expose participants to the time-varying dynamics. The order of runs was determined using an incomplete Latin square: considering each of the 22 runs as separate entities, a full Latin square of 22 different experimental sequences was obtained. From this Latin square ten test condition orders were randomly chosen for use in the experiment.

G. Data Analysis

During the experiment, all measurable control loop signals as indicated in Fig. 4 were recorded as the measurement data. As explained in subsection III.D, the MRAC model of Fig. 3 was fit to the average measured human control output data (u_h) to determine its parameters: the internal reference model parameters (ω_c, τ), the MRAC learning rates ($\gamma_{x1}, \gamma_{x2}, \gamma_r$), and the time-varying control gains (k_{x1}, k_{x2}, k_r). For this estimation, the 5-run *testing* datasets for conditions DYN 121 and DYN 212 were used for each participant. It should be noted that, for reference, also transition-specific model fits were obtained, by using only the first transition data and excluding the last thirty seconds of each run.

To quantify the attained MRAC model quality-of-fit, the Variance Accounted For (VAF) was used:

$$\text{VAF} = \max \left(0, \left(1 - \frac{\sum_{k=1}^N |u_h[k] - u_m[k]|^2}{\sum_{k=1}^N |u_h[k]|^2} \right) \right) \quad (24)$$

Please note from Eq. (24) that the VAF as considered here varies between 0 (poor model fit) to 1 (perfect corresponding between measured u_h and modelled u_m) and that we clip the VAF at 0 to avoid negative VAF values. For all fitted models, the VAF was calculated both for the testing (5 runs) and validation (3 runs) datasets, to assess the generalization ability of the MRAC model, especially in conditions DYN 121 and DYN 212. Furthermore, the VAF across complete average tracking runs of N data points as defined in Eq. (24) was calculated, but also a sliding-window equivalent (10-second window) to analyze local variations in fit quality.

For comparing the MRAC model's quality-of-fit and parameter variations across the tested conditions, a statistical analysis was performed with a linear mixed-effects model [54] that included the participant number as the random factor.

V. MRAC Sensitivity Analysis

A sensitivity analysis was performed to analyse the relation between the selected MRAC learning rates (i.e., $\gamma_{x1}, \gamma_{x2}, \gamma_r$) and the adaptation of the MRAC control gains. The learning rates control how quickly the model adapts to the changes in the controlled dynamics as detected through errors between the expected and the observed output of the plant (e_p , see Fig. 3). In the considered MRAC formulation the three learning rates γ_{x1}, γ_{x2} and γ_r , i.e., one for each control gain, can be set and varied independently. The results shown here, however, were obtained with the same learning rate lr for all three parameters and are only meant to provide a feel for the effect of the learning rate setting. For brevity, only simulated results for condition DYN 121, where the controlled dynamics change between an initial approximate single-integrator system and a double-integrator system, are shown for learning rates settings of $lr = [0.5, 1, 5, 10, 30]$.

Fig. 6 shows both the time traces of the system pitch attitude x and the operator's control input u for different learning rates. Note that in the top figure (x) also the target signal r is shown for reference. Fig. 6 shows that the lower the learning rate the sharper the oscillations observed after the change in dynamics at 30 seconds. Furthermore, when the dynamics of the plant are changed back to an approximate single integrator at 60 seconds, higher learning rates cause an undershoot of the reference signal. Fig. 7 depicts the corresponding time traces of k_r, k_{x1} and k_{x2} . The most pronounced change is observed in k_{x2} , which sharply increases after the dynamics are changed from a single to a double integrator at $t = 30$ s, as expected. The expected drop of the k_{x2} gain after the dynamics are reverted back to the approximate single integrator at $t = 60$ s is, however, less pronounced and for low lr even almost negligible. The gains change faster when the dynamics change from 1 to 2 since the controller initially is not able to stabilize the system and large tracking errors occur that are inconsistent with reference model's predictions. On the other hand, for a transition from double to single integrator controlled dynamics, the closed-loop system retains stability and errors between the internal model output and the system output are thus much smaller, which leads to notably less gain adaptation, see Fig. 7. Finally, a notable increase of k_r is also observed in Fig. 7 at 60 seconds for larger values of learning rate. The increase in k_r is caused by the still elevated value of k_{x2} , which decreases the control output of the controller.

Fig. 6 Output and control output sensitivity to the learning rate lr for condition DYN 121.

Fig. 7 Sensitivity of MRAC gains to the learning rate lr for condition DYN 121.

VI. Experimental Results

A. Quality of Fit

As explained in Section IV.G, the MRAC model of Fig. 3 was fit to the data collected from all experiment participants. Fig. 8 shows a Box plot of the Variance Accounted For (VAF) of the aggregated control output for the estimated MRAC controllers. As can be verified from Fig. 8, no significant changes in VAF across the different steady-state (DYN 1 and 2) and time-varying conditions (DYN 121 and 212) were found. Consistent with earlier experiments [55], the VAF is significantly lower for control of a single integrator system (DYN 1, $M = 0.601$) compared to double integrator control data (DYN 2, $M = 0.722$), $b = -0.053$ ($SE = 0.013$), $t(22) = -3.901$, $p < 0.001$. The fact that single-integrator control data are less accurately modeled also explains the significantly lower VAF for DYN 121 ($M = 0.644$) compared to DYN 212 ($M = 0.71$), $b = -0.027$, $t(22) = -2.105$, $p = 0.047$. Please note that while the lower VAF for both time-varying conditions on the validation dataset (orange boxes) compared to the testing dataset (blue boxes) seems to suggest overfitting of the MRAC model, these lower validation VAFs can in fact be partially explained by the lower number of tracking runs used for averaging for the validation dataset (testing: 5 runs, validation: 3 runs).

Fig. 8 Variance Accounted For (VAF) of the fitted MRAC model for different controlled dynamics conditions.

Because of the time-varying nature of the adaptive controller, next to the overall VAF it is relevant to consider time variations in fit quality to assess the ability of MRAC to capture how humans adapt. Fig. 9a and 9b show the corresponding aggregated time-series data of the windowed VAF computed over a moving 10-second time window for the DYN 121 and DYN 212 conditions, respectively. A time window of ten seconds is able to show the ability of MRAC to capture the transient dynamics without being too strongly affected by stochastic variations the measured signals. Two observations can be made from Fig. 9. First, the model fits the data better during the blue shaded windows when the double integrator (DYN 2) dynamics are controlled, with a mean of approximately 0.75 compared to the average VAF of around 0.6 observed in the (white shaded) DYN 1 windows. This observation is consistent with the VAF data in Fig. 8. Second, while the VAF tends to increase for transitions from DYN 1 to DYN 2, it drops noticeably more for the reverse transitions from DYN 2 to DYN 1, see for example the two DYN 1 windows in Fig. 9a. This difference suggests that the MRAC controller can capture how participants adapt their control strategy when transitioning from effective single integrator dynamics to approximate double integrator dynamics, while the opposite is not true.

Based on consideration of time-windowed VAF data as shown in Fig. 9, a number of collected datasets were found to be insufficiently consistent for reliable MRAC model fitting. Two examples of such outlying time-windowed VAF data, for which the VAF even drops to 0 for certain periods, are shown in Fig. 9b. As in such cases also the estimated MRAC parameters are likely not representative, the following data have been excluded from the aggregated analysis of the fitted MRAC model: data from Participants 2 and 7 for DYN 2 and DYN 212 and from Participant 5 for DYN 2. While these exclusions are attributable to the participants comparatively inconsistent tracking for these conditions, they also indicate that MRAC model is not very robust when fitting it to non-expert human controller data.

(a) Time-windowed VAF for condition DYN 121.

(b) Time-windowed VAF for condition DYN 212. VAFs of the two outliers are shown with separate solid lines.

Fig. 9 Time-windowed VAF of the MRAC model for a 10-second window size. The solid red lines represent the mean values, while the shaded red areas indicate one standard deviation.

B. MRAC Prediction Error

To further investigate why the MRAC model does not seem to capture transitions from single to double integrator control and vice versa at the same accuracy, we compare the control output u , controlled system output x , and the prediction error e_p for the MRAC controller aggregated across all participants. For brevity, the former comparison data are included in Appendix A, which shows comparisons of the aggregated measured controlled system output x and human operator output u with our MRAC predictions, for reference.

Fig. 10 shows the aggregated prediction error, e_p , for the estimated controllers. The prediction error is the difference between the observed output generated by the internal model and the actual output of the controlled system, see Fig. 3. This parameter, together with the state x , target input r , and the learning rates drives the adaptation of the adaptive MRAC gains. As is clear from Fig. 10, the prediction error is largest after the two controlled system transitions at $t = 30$ s and $t = 60$ s. As expected, due to the fact that the internal reference model assumes a perfect linear human without remnant, the prediction error never converges to zero. For DYN 121, Fig. 10a shows that the magnitude of e_p after the first transition from DYN 1 to DYN 2 is larger than for the second opposite transition at $t = 60$ s. The DYN 212 data in Fig. 10b show more comparable prediction error magnitude for the transitions from DYN 2 to DYN 1 and vice versa. Still, with larger magnitudes of for the state x and its derivative $\dot{x}$ (see Appendix A), the MRAC gain adaptations are consistently stronger for DYN 1 to DYN 2 transitions.

C. Adaptive Controller Gains

Before presenting MRAC control gain adaptation results for the time-varying conditions DYN 121 and DYN 212, first the gains obtained for the steady-state conditions DYN 1 and DYN 2 are discussed here for reference. As expected [7, 8, 9], the gain k_{x1} was found to be significantly higher in condition DYN 1 ($M = 0.179$) than for double integrator control (DYN 2, $M = 0.145$), $b = 0.037$ ($SE = 0.01$), $t(6) = 3.696$, $p = 0.01$. Also the gain k_r was slightly higher in condition DYN 1 on average ($M = 0.163$) compared to condition DYN 2 ($M = 0.153$), $b = 0.033$ ($SE = 0.015$), $t(6) = 2.24$, $p = 0.06$. Finally, as expected due to the increased requirement for lead compensation, the gain k_{x2} was significantly higher in condition DYN 2 ($M = 0.070$) than for DYN 1 ($M = 0.045$), $b = 0.024$ ($SE = 0.002$), $t(6) = 11.81$, $p < 0.001$. These results, i.e., higher state derivative gains and lower state feedback and feedforward gains for condition DYN 2 compared to DYN 1, are in line with the pursuit task data of [55].

For both time-varying conditions, Fig. 11 shows the time series of the MRAC gains for each individual participant, while Fig. 12 shows the corresponding aggregated average results. Especially from Fig. 11 it can be observed that the feedforward gain k_r (in black) and the state gain k_{x1} (in blue) did not show significant changes due to the controlled element transitions; the statistical variability between participants in these gain values is much stronger. The state derivative feedback gain k_{x2} (in red), however, was found to adapt significantly after a change in the controlled dynamics.

Fig. 10 Mean and standard deviation of the prediction error for the experimental data of conditions DYN 121 and DYN 122.

In condition DYN 121 at $t = 30$ s and condition DYN 122 at $t = 60$ s, so for a transition from our single to double integrator dynamics, k_{x2} increased from 0.05 to 0.085 on average. This means that the contribution of ‘lead’ in the feedback control increased, as expected [4, 6]. On the other hand, for the reverse transitions in both conditions, k_{x2} decreased less, from a value of 0.085 to 0.06, and also much less sharply. As also noted from our sensitivity analysis in Section V, the rate of change of the MRAC gains is not symmetric: a change from DYN 1 to DYN 2 leads to a much faster adaptation than the opposite change from DYN 2 to DYN 1.

D. Reference Model Parameters and Learning Rates

In addition to the time-varying control gains, for the MRAC model of Fig. 3 also a number of different parameters were estimated from the experiment data. Fig. 13a and Fig. 13b show the estimated crossover frequency ω_c and time delay τ of the internal model of the MRAC controller, respectively, across all participants. For ω_c , which is found to be between 2-2.5 rad/s on average for all conditions, no statistically significant changes across conditions were found. For the delay τ , only the difference between DYN 1 ($M = 0.228$ s) and DYN 2 ($M = 0.207$ s) is statistically significant, $b = 0.009$ ($SE = 0.003$), $t(22) = 2.64$, $p = 0.014$. Still, the corresponding effect size is very small. Overall these results are in line with the findings of [7] and suggest that a constant reference model for our MRAC controller should provide reasonable predictions.

The final estimated MRAC parameters, which are important to have a complete picture of how its gains can adapt over time, are the learning rates γ_{x1} , γ_{x2} , and γ_r . Fig. 14 shows these learning rates associated with the different gains, each indicated with its own box color, for conditions DYN 121 and DYN 122. On average, γ_{x2} is found to be around 5, while the learning rates for k_{x1} and k_r are higher and between 10 and 15. Note that these learning rates correspond also to the higher lr values tested in the sensitivity analysis in Section V and suppress strong transients at the controlled system transitions. Finally, as expected due to the fact that in both the DYN 121 and DYN 122 both types of controlled dynamics transitions are included, Fig. 14 shows that the learning rates are not significantly different across the two conditions.

Fig. 11 Estimated time series of MRAC gains for each participant.

Fig. 12 Estimated mean and standard deviation of MRAC gains for conditions DYN 121 and DYN 212.

Fig. 13 Estimated internal crossover model parameters across the different dynamics conditions.

Fig. 14 Estimated learning rates for conditions DYN 121 and DYN 212.

Fig. 15 Estimated learning rates for the reduced 60-second DYN 12 and DYN 21 datasets.

E. Transition-Specific MRAC Fits

In the previous subsections, we observed that MRAC does not fit the adaptive human control data equally well in transitions from DYN 1 to DYN 2 and vice versa. For all previously presented data, we have compared both types of transitions with the same learning rate settings. In this subsection, we further investigate whether fitting MRAC to each transition separately, and thus having a potentially transition-specific learning rate, can increase the quality-of-fit for DYN 2 to DYN 1 transitions. For this analysis, the MRAC fitting procedure was repeated based on only the first 60 s of the tracking runs for conditions DYN 121 and DYN 212. To emphasize that only a single transition in controlled dynamics is considered, here we refer to these reduced datasets as DYN 12 and DYN 21, respectively.

Fig. 15 shows a Box plot of the MRAC learning rates estimated on the reduced DYN 12 and DYN 21 datasets. Comparison with Fig. 14 shows that indeed the median learning rates for DYN 21 are estimated to be consistently higher than those found for DYN 12. In particular, for DYN 12 the median $\gamma_{x1} = 12 \text{ rad}^{-1}$ and $\gamma_{x2} = 5 \text{ rad}^{-1} \text{ s}^{-1}$, while for DYN 21 the median $\gamma_{x1} = 22 \text{ rad}^{-1}$ and $\gamma_{x2} = 18 \text{ rad}^{-1} \text{ s}^{-1}$. Furthermore, these values are also higher than those found for the full condition tracking runs (DYN 121 and DYN 212, see Fig. 14), for which the median $\gamma_{x1} = 13.8 \text{ rad}^{-1}$ and $\gamma_{x2} = 6 \text{ rad}^{-1} \text{ s}^{-1}$. Overall, this is the expected result, as the transition from DYN 1 to DYN 2 naturally induces a stronger adaptation due to larger changes in x and thus e_p compared to the DYN 2 to DYN 1 transitions. An increased MRAC learning rate, to some extent, can compensate for this effect.

The means and standard deviations of the MRAC gains for the reduced DYN 12 and DYN 21 datasets are shown in Fig. 16. For DYN 12, a direct comparison with Fig. 12 shows no significant differences in the post-transition magnitude of the gains, nor in the rate of change of the k_{x2} gain, which still transitions from a value of 0.05 to 0.09 in approximately 4 seconds, as can also be observed in Fig. 12. The main observed difference due to the transition-specific learning rate is seen in the standard deviation of the gain estimates, which for example is reduced from 0.04 to 0.02 for k_{x1} . For the reverse transition in dataset DYN 21, a larger difference in the adaptation of k_{x2} is observed compared to DYN 12. For both cases, the initial value of $k_{x2} \approx 0.08$ while participants controlled DYN 2. After the transition, k_{x2} was found to stabilize around 0.05 for DYN 212, while a stronger drop to approximately 0.03 is found for the learning rates fitted to the DYN 21 data.

Finally, Figs. 17a and 17b show the corresponding windowed VAF of the MRAC models fitted on the reduced DYN 12 and DYN 21 datasets. These results show that notwithstanding a substantial increase in learning rate, the mean windowed VAF data after the transition for DYN 12 and DYN 21 are not improved compared to their values for conditions DYN 121 and DYN 212, respectively. Especially for DYN 21, Fig. 17b still shows the initial drop to a VAF of around 0.4 directly after the transition. Overall, this indicates that the asymmetry in the human control adaptation predictions provided by MRAC cannot be compensated with transition-specific learning rate parameters.

VII. Discussion

This paper evaluated whether Model Reference Adaptive Control (MRAC), an adaptive control technique that uses an ‘internal model’ to drive controller adaptation when task conditions change, is a viable candidate for modeling and predicting human adaptive control behavior due to changes in controlled element dynamics. Representative transitions between controlled dynamics approximating single and double integrator dynamics also tested in earlier experiments were chosen to assess the validity of MRAC’s adaptive law in simulations and a human-in-the-loop experiment performed with ten participants.

We first assessed the ability of MRAC to capture steady-state control behavior of the participants. The median VAF for our MRAC model was found to be 0.64 and 0.75 for conditions DYN 1 and DYN 2, respectively. The consistently lower VAFs for single-integrator control data are consistent with earlier research [7, 55]. The overall quality-of-fit of our MRAC model seems to be slightly lower for these steady-state cases compared to, for example, the parametric time-varying model fits with 5-run average VAFs of 0.77 and 0.85 reported in [7] for the same controlled dynamics. Still, given the truly adaptive nature of our MRAC model, the quality-of-fit is still found to be sufficient to obtain reasonably accurate human control behavior predictions.

Our results as presented in this paper show that the ability of MRAC to approximate human time-varying control behavior strongly depends on the dynamics of the transition that occurs in the controlled dynamics. When the controlled system was suddenly changed from an (approximate) single integrator system to dynamics that approximate a double integrator, the significant differences that occur between the ‘stable’ reference model and ‘unstable’ control loop enable our MRAC model to match well with experimental human control data. In contrast, for the opposite transition a sharp drop in VAF indicates less accurate prediction of human adaptation by MRAC. This was shown to result from insufficiently strong adaptation of the MRAC model’s control gains (k_r, k_{x1}, k_{x2}) due to the fact that less dramatic prediction errors occurred for this transition to a more stable controlled system. Overall, the results show that the current MRAC adaptive law can be a good approximation of the adaptive mechanism that humans operators use to adapt their control strategy, but that in its current form it cannot predict all time-varying adaptations well.

In this paper, we also investigated whether the asymmetry in MRAC’s prediction for controlled element transitions from single to double integrator controlled dynamics and vice versa could be alleviated with different learning rates for the MRAC model’s adaptive control gains for both types of transitions. This was done by fitting transition-specific learning rates using only the single-transition data (the first 60 seconds) for conditions DYN 121 and DYN 212. While indeed increased learning rates and a factor two faster adaptation of the k_{x2} gain were found for the transition from double to single integrator dynamics, the MRAC model’s accuracy (VAF) in describing human control adaptation directly following this transition was not improved. From our analysis, it was also found that even higher learning rates are not possible with an MRAC controller meant to approximate human control behavior, due to the presence of a (human) delay in the control loop. While a nominal MRAC controller has theoretical convergence guarantees, these no longer hold in the presence of a delay and too high learning rates may cause MRAC to become unstable.

Our current implementation of the MRAC model is centered on a scalar prediction error e_p , defined as the difference between a predicted controlled element output x_m and its true value x . This definition of e_p ensures that the adaptive control laws are activated the moment that a difference in tracking performance occurs. Hence, we can accurately

Fig. 16 Mean and standard deviation of the aggregated MRAC gains, estimated using only the first sixty seconds of the each run.

Fig. 17 Time-windowed VAF of the MRAC model for the reduced DYN 12 and DYN 21 datasets for a 10-second window size. The solid red lines represent the mean values, while the shaded red areas indicate one standard deviation.

predict the expected adaptation for transitions from DYN 1 to DYN 2 for which such a performance discrepancy occurs immediately due to reduced loop stability. In reality, human operators would likely not only change their control behavior if a performance discrepancy occurs, but also when their required control effort becomes noticeably different (even if performance is not affected). The lacking adaptation we predict for DYN 2 to DYN 1 transitions results in an MRAC model prediction that ‘over-controls’ the DYN 1 dynamics post-transition, see Fig. 18 and 19. Hence, a crucial next research step for improving our current MRAC model is to investigate the effects of adding a control effort prediction error to the definition of e_p .

In our current implementation of an MRAC model for predicting human control adaptation, a fixed internal reference model was used. For our application to control adaptation in a pursuit task, with changes in the controlled system dynamics this is partially correct, as human controllers will adapt to achieve largely invariant open-loop system dynamics even when the controlled dynamics change [4, 6]. However, as McRuer and Jex [6] show for compensatory tracking tasks, the parameters (crossover frequency, effective time delay) of the open loop dynamics do change, which is not accounted for in our model. In addition, for application to human control adaptations for which this invariance does not hold, such as changes in tracking display type from pursuit to preview displays [55], our current MRAC implementation can perhaps be coupled with a separate, on-line estimator that also uses occurring prediction errors to drive the adaptation of MRAC’s internal model. As human controllers also update their internal models through learning [2, 4], a good predictive model of human control performance would also be expected to include an adaptive internal model.

The use of an MRAC controller to model human time-varying pursuit tracking behavior, as investigated in this paper, is still a limited view on Predictive Coding (PC) in human control. For example, unlike more sophisticated models of PC, our presented MRAC framework cannot by itself generate the complex strategies, possibly subject to constraints, that human controllers can carry out, such as landing without instruments, performing an obstacle avoidance maneuver, and reaching a hovering target under disturbances with an helicopter. A promising research direction towards this goal are techniques generally referred to as ‘model-based optimal control’, which can capture complex dynamical behavior shown by animals [20], including open-loop planning [10]. Alternatively, ‘inverse reinforcement learning’ [56] can perhaps be used to estimate the reward function (or intention) underlying observed complex control behavior and used to derive an explicit control policy for predicting human control behavior. Overall, we see a potential for PC to enable increased insight in complex adaptive human behavior as exhibited throughout the aerospace field.

VIII. Conclusion

This paper evaluated the use of a model-based adaptive control algorithm, Model Reference Adaptive Control (MRAC), for predicting the adaptive control policies of human operators in control of a system with time-varying dynamics in a pursuit tracking task. MRAC, with its internal reference model that is used to drive controller adaptation, was chosen as biological systems are known to make use of internal models for sensorimotor control, as also formalized in the neuroscience theory of ‘*predictive coding*’. In our implementation, the internal reference model of the MRAC controller was chosen to be McRuer and Jex’s *crossover model*, which captures the invariance of the open-loop dynamics in closed-loop manual control tasks across a wide range of controlled system dynamics.

An experiment was conducted where ten participants performed a pursuit tracking task in which the controlled dynamics transitioned between approximated single and double integrator systems twice during each tracking run. The parameters of the MRAC model were estimated by fitting it to the collected experiment data by minimizing the error between the MRAC output and the response of each participant, averaged across multiple runs.

The MRAC model was found to accurately predict the transient control policy of the participants for transitions from single to double integrator controlled system dynamics, a well-known adaptation that requires an increase in rate feedback compensation. On the other hand, MRAC was not able to approximate the change in the control policy of the participants for the inverse transition. Even with transition-specific learning rate settings, for the transition to more stable controlled system dynamics, which inherently results in smaller internal model prediction errors, the MRAC model’s gains were updated too slowly compared to the speed of human adaption. While MRAC was only partially successful at predicting human control behavior under different changes of the controlled system, the intuitive control scheme can be easily extended in future work and contribute to innovations in machine learning and predictive coding towards creating model-based controllers with human-like adaptive control policies.

A. MRAC and Experimental Time Trace Data

Fig. 18 shows the reference signal r and the system's output x aggregated across participants for each of the four tested conditions (testing dataset). The aggregated output shows the mean response as a solid line and the standard deviation with a shaded area. The control output of the participants is shown in blue while the corresponding MRAC model predictions are shown in red. Fig. 18 shows that participants controlled the steady-state system for conditions DYN 1 and DYN 2 consistently. For the time-varying conditions, participants' tracking ability decreased just after a change in controlled element dynamics. In particular, participants tend to overshoot the target by 35% on average when the dynamics change from DYN 1 to DYN 2 and undershoot it by around 30% when the dynamics change from DYN 2 to DYN 1. After the 30-second mark in condition DYN 121 in Fig. 18c, when the dynamics change from DYN 1 to DYN 2, the next two peaks in the reference signal are both significantly overshoot. The same can be observed for DYN 212 in Fig. 18d after the transition at $t = 60$ s. The output standard deviation increased after the dynamics transition. Tracking performance gradually improved after the transition, as participants adapted to the new dynamics.

Fig. 18 Mean and standard deviation of the aggregated output of the participants, y_h , and of the MRAC controllers, y_m , for the four tested conditions.

Fig. 19 compares the aggregated measured control output u and the MRAC controller predictions. For condition DYN 121, between $t = 30$ and 40 s, the MRAC model accurately captures the transient control behavior shown by participants, who have a higher control activity and tend to overshoot the target reference signal as shown in Fig. 18. On

the other hand, MRAC is not equally able to capture the transition from DYN 2 back to DYN 1 at $t = 60$ s. Fig. 19c and d show that the control output is lower and has higher frequency components than the measured u signals.

Fig. 19 Mean and standard deviation of the aggregated control output of the participants, u_h , and the MRAC models, u_m , for the four tested conditions.

References

- [1] K. He et al. “Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification”. In: *2015 IEEE International Conference on Computer Vision (ICCV)*. 2015 IEEE International Conference on Computer Vision (ICCV). Dec. 2015, pp. 1026–1034. doi: 10.1109/ICCV.2015.123.
- [2] L. R. Young. “On Adaptive Manual Control”. In: *IEEE Transactions on Man-Machine Systems* 10.4 (Dec. 1969), pp. 292–331. ISSN: 2168-2860. doi: 10.1109/TMMS.1969.299931.
- [3] D. M. Pool et al. “Pilot Equalization in Manual Control of Aircraft Dynamics”. In: *Proc. of the IEEE Conference on Systems, Man, & Cybernetics (IEEE - SMC), San Antonio (TX), October 11-14 (2009)*, pp. 2554–2559.
- [4] M. Mulder et al. “Manual Control Cybernetics: State-of-the-Art and Current Trends”. In: *IEEE Transactions on Human-Machine Systems* 48.5 (2018), pp. 468–485. ISSN: 2168-2305. doi: 10.1109/THMS.2017.2761342.
- [5] K. J. Friston. “Does Predictive Coding Have a Future?” In: *Nature Neuroscience* 21.8 (Aug. 1, 2018), pp. 1019–1021. ISSN: 1546-1726. doi: 10.1038/s41593-018-0200-7.

- [6] D. T. McRuer and H. R. Jex. "A Review of Quasi-Linear Pilot Models". In: *IEEE Transactions on Human Factors in Electronics* HFE-8.3 (Sept. 1967), pp. 231–249. ISSN: 2168-2852. DOI: 10.1109/THFE.1967.234304.
- [7] P. M. T. Zaal. "Manual Control Adaptation to Changing Vehicle Dynamics in Roll–Pitch Control Tasks". In: *Journal of Guidance, Control, and Dynamics* 39.5 (2016), pp. 1046–1058. DOI: 10.2514/1.G001592. URL: <https://doi.org/10.2514/1.G001592>.
- [8] A. van Grootheest et al. "Identification of Time-Varying Manual Control Adaptations with Recursive ARX Models". In: *2018 AIAA Modeling and Simulation Technologies Conference*. AIAA SciTech Forum. American Institute of Aeronautics and Astronautics, Jan. 7, 2018. DOI: 10.2514/6.2018-0118. URL: <https://arc.aiaa.org/doi/10.2514/6.2018-0118>.
- [9] W. Plaetinck et al. "Online Identification of Pilot Adaptation to Sudden Degradations in Vehicle Stability". In: *IFAC-PapersOnLine* 51.34 (2019), pp. 347–352. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2019.01.020. URL: <https://dx.doi.org/10.1016/j.ifacol.2019.01.020>.
- [10] D. McNamee and D. M. Wolpert. "Internal Models in Biological Control". In: (Oct. 3, 2018). ISSN: 2573-5144. DOI: 10.17863/CAM.30459. URL: <https://www.repository.cam.ac.uk/handle/1810/283097>.
- [11] R. C. Miall and D. M. Wolpert. "Forward Models for Physiological Motor Control". In: *Four Major Hypotheses in Neuroscience* 9.8 (Nov. 1, 1996), pp. 1265–1279. ISSN: 0893-6080. DOI: 10.1016/S0893-6080(96)00035-4. URL: <http://www.sciencedirect.com/science/article/pii/S0893608096000354>.
- [12] M. Kawato. "Internal Models for Motor Control and Trajectory Planning". In: *Current Opinion in Neurobiology* 9.6 (Dec. 1, 1999), pp. 718–727. ISSN: 0959-4388. DOI: 10.1016/S0959-4388(99)00028-8. URL: <http://www.sciencedirect.com/science/article/pii/S0959438899000288>.
- [13] A. J. Bastian. "Moving, Sensing and Learning with Cerebellar Damage". In: *Current opinion in neurobiology* 21.4 (2011), pp. 596–601. ISSN: 0959-4388. DOI: 10.1016/j.conb.2011.06.007. pmid: 21733673. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3177958/> (visited on 01/14/2021).
- [14] C. J. Stoodley. "The Cerebellum and Neurodevelopmental Disorders". In: *Cerebellum (London, England)* 15.1 (Feb. 2016), pp. 34–37. ISSN: 1473-4222. DOI: 10.1007/s12311-015-0715-3. pmid: 26298473. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4811332/> (visited on 01/13/2021).
- [15] D. M. Wolpert, Z. Ghahramani, and M. I. Jordan. "An Internal Model for Sensorimotor Integration". In: *Science* 269.5232 (Sept. 29, 1995), p. 1880. DOI: 10.1126/science.7569931. URL: <http://science.sciencemag.org/content/269/5232/1880.abstract>.
- [16] D. M. Wolpert, R. C. Miall, and M. Kawato. "Internal Models in the Cerebellum". In: *Trends in Cognitive Sciences* 2.9 (Sept. 1, 1998), pp. 338–347. ISSN: 1364-6613. DOI: 10.1016/S1364-6613(98)01221-2. URL: <http://www.sciencedirect.com/science/article/pii/S1364661398012212> (visited on 01/14/2021).
- [17] H. Tanaka et al. "The Cerebro-Cerebellum as a Locus of Forward Model: A Review". In: *Frontiers in Systems Neuroscience* 14 (2020). ISSN: 1662-5137. DOI: 10.3389/fnsys.2020.00019. URL: <https://www.frontiersin.org/articles/10.3389/fnsys.2020.00019/full> (visited on 01/14/2021).
- [18] S. Shipp. "Neural Elements for Predictive Coding". In: *Frontiers in Psychology* 7 (2016). ISSN: 1664-1078. DOI: 10.3389/fpsyg.2016.01792. URL: <https://www.frontiersin.org/articles/10.3389/fpsyg.2016.01792/full>.
- [19] E. Todorov and M. I. Jordan. "Optimal Feedback Control as a Theory of Motor Coordination". In: *Nature Neuroscience* 5.11 (Nov. 1, 2002), pp. 1226–1235. ISSN: 1546-1726. DOI: 10.1038/nn963. URL: <https://doi.org/10.1038/nn963>.
- [20] E. Todorov. "Optimality Principles in Sensorimotor Control". In: *Nature Neuroscience* 7.9 (Sept. 1, 2004), pp. 907–915. ISSN: 1546-1726. DOI: 10.1038/nn1309. URL: <https://doi.org/10.1038/nn1309>.
- [21] K. J. Friston. "What Is Optimal about Motor Control?" In: *Neuron* 72.3 (Nov. 3, 2011), pp. 488–498. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2011.10.018. URL: <http://www.sciencedirect.com/science/article/pii/S0896627311009305>.
- [22] J. Diedrichsen, R. Shadmehr, and R. B. Ivry. "The Coordination of Movement: Optimal Feedback Control and Beyond". In: *Trends in Cognitive Sciences* 14.1 (Jan. 1, 2010), pp. 31–39. ISSN: 1364-6613. DOI: 10.1016/j.tics.2009.11.004. URL: <http://www.sciencedirect.com/science/article/pii/S1364661309002587>.

- [23] S. H. Scott. “The Computational and Neural Basis of Voluntary Motor Control and Planning”. In: *Trends in Cognitive Sciences* 16.11 (Nov. 1, 2012), pp. 541–549. ISSN: 1364-6613. DOI: 10.1016/j.tics.2012.09.008. URL: <https://doi.org/10.1016/j.tics.2012.09.008> (visited on 09/02/2020).
- [24] K. J. Friston et al. “Action and Behavior: A Free-Energy Formulation”. In: *Biological Cybernetics* 102.3 (Mar. 2010), pp. 227–260. ISSN: 1432-0770. DOI: 10.1007/s00422-010-0364-z. PMID: 20148260.
- [25] K. J. Friston. “The Free-Energy Principle: A Unified Brain Theory?”. In: *Nature Reviews Neuroscience* 11.2 (Feb. 1, 2010), pp. 127–138. ISSN: 1471-0048. DOI: 10.1038/nrn2787. URL: <https://doi.org/10.1038/nrn2787>.
- [26] W. Jiahui. *Python Interactive Network Visualization Using NetworkX, Plotly and Dash*. URL: <https://towardsdatascience.com/python-interactive-network-visualization-using-networkx-plotly-and-dash-e44749161ed7>.
- [27] M. Heilbron and M. Chait. “Great Expectations: Is There Evidence for Predictive Coding in Auditory Cortex?”. In: *Sensory Sequence Processing in the Brain* 389 (Oct. 1, 2018), pp. 54–73. ISSN: 0306-4522. DOI: 10.1016/j.neuroscience.2017.07.061. URL: <http://www.sciencedirect.com/science/article/pii/S030645221730547X>.
- [28] A. Bastos et al. “Canonical Microcircuits for Predictive Coding”. In: *Neuron* 76 (Nov. 21, 2012), pp. 695–711. DOI: 10.1016/j.neuron.2012.10.038.
- [29] J. F. Mejias et al. “Feedforward and Feedback Frequency-Dependent Interactions in a Large-Scale Laminar Network of the Primate Cortex”. In: *Science Advances* 2.11 (Nov. 1, 2016), e1601335. DOI: 10.1126/sciadv.1601335. URL: <http://advances.sciencemag.org/content/2/11/e1601335.abstract>.
- [30] R. P. N. Rao and D. H. Ballard. “Predictive Coding in the Visual Cortex: A Functional Interpretation of Some Extra-Classical Receptive-Field Effects”. In: *Nature Neuroscience* 2.1 (1 Jan. 1999), pp. 79–87. ISSN: 1546-1726. DOI: 10.1038/4580. URL: https://www.nature.com/articles/nn0199_79.
- [31] M. I. Garrido et al. “Dynamic Causal Modeling of the Response to Frequency Deviants”. In: *Journal of Neurophysiology* 101.5 (May 1, 2009), pp. 2620–2631. ISSN: 0022-3077. DOI: 10.1152/jn.90291.2008. URL: <https://doi.org/10.1152/jn.90291.2008>.
- [32] A. Todorovic and F. P. de Lange. “Repetition Suppression and Expectation Suppression Are Dissociable in Time in Early Auditory Evoked Fields”. In: *The Journal of Neuroscience* 32.39 (Sept. 26, 2012), p. 13389. DOI: 10.1523/JNEUROSCI.2227-12.2012. URL: <http://www.jneurosci.org/content/32/39/13389.abstract>.
- [33] N. Kogo and C. Trengove. “Is Predictive Coding Theory Articulated Enough to Be Testable?”. In: *Frontiers in computational neuroscience* 9 (Sept. 8, 2015), pp. 111–111. ISSN: 1662-5188. DOI: 10.3389/fncom.2015.00111. PMID: 26441621. URL: <https://pubmed.ncbi.nlm.nih.gov/26441621>.
- [34] A. Popovici, P. M. T. Zaal, and D. M. Pool. “Dual Extended Kalman Filter for the Identification of Time-Varying Human Manual Control Behavior”. In: *AIAA Modeling and Simulation Technologies Conference*. American Institute of Aeronautics and Astronautics, June 2, 2017. DOI: 10.2514/6.2017-3666. URL: <https://arc.aiaa.org/doi/10.2514/6.2017-3666>.
- [35] J. Rojer et al. “UKF-Based Identification of Time-Varying Manual Control Behaviour”. In: *IFAC-PapersOnLine*. 14th IFAC Symposium on Analysis, Design, and Evaluation of Human Machine Systems HMS 2019 52.19 (Jan. 1, 2019), pp. 109–114. ISSN: 2405-8963. DOI: 10.1016/j.ifacol.2019.12.120. URL: <http://www.sciencedirect.com/science/article/pii/S240589631931955X> (visited on 01/18/2021).
- [36] E. R. Boer and R. V. Kenyon. “Estimation of Time-Varying Delay Time in Nonstationary Linear Systems: An Approach to Monitor Human Operator Adaptation in Manual Tracking Tasks”. In: *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 28.1 (Jan. 1998), pp. 89–99. ISSN: 1558-2426. DOI: 10.1109/3468.650325.
- [37] F. M. Nieuwenhuizen et al. “Modeling Human Multichannel Perception and Control Using Linear Time-Invariant Models”. In: *Journal of Guidance, Control, and Dynamics* 31.4 (2008), pp. 999–1013. DOI: 10.2514/1.32307.
- [38] M. Linssen. “Identifying Time-Varying Multimodal Manual Control Using Recursive ARX Model Techniques”. In: (2020). URL: <https://repository.tudelft.nl/islandora/object/uuid%3A442f4308-0ea2-41a5-b38c-ee6b1a289f78> (visited on 01/18/2021).

- [39] P. M. T. Zaal et al. "Modeling Human Multimodal Perception and Control Using Genetic Maximum Likelihood Estimation". In: *Journal of Guidance, Control, and Dynamics* 32.4 (July 1, 2009), pp. 1089–1099. doi: 10.2514/1.42843. URL: <https://arc.aiaa.org/doi/10.2514/1.42843> (visited on 09/22/2020).
- [40] S. Xu et al. "Review of Control Models for Human Pilot Behavior". In: *Annual Reviews in Control* 44 (Jan. 1, 2017), pp. 274–291. ISSN: 1367-5788. doi: 10.1016/j.arcontrol.2017.09.009. URL: <http://www.sciencedirect.com/science/article/pii/S136757881730024X>.
- [41] A. V. Phatak and G. A. Bekey. "Model of the Adaptive Behavior of the Human Operator in Response to a Sudden Change in the Control Situation". In: *IEEE Transactions on Man-Machine Systems* 10.3 (Sept. 1969), pp. 72–80. ISSN: 2168-2860. doi: 10.1109/TMMS.1969.299886.
- [42] L. Stark and L. R. Young. *Biological Control System - a Critical Review and Evaluation, Developments in Manual Control*. In collab. with NASA. Mar. 1, 1965. URL: http://archive.org/details/nasa_techdoc_19650009660.
- [43] R. A. Hess. "Modeling Human Pilot Adaptation to Flight Control Anomalies and Changing Task Demands". In: *Journal of Guidance, Control, and Dynamics* 39.3 (2016), pp. 655–666. doi: 10.2514/1.G001303. URL: <https://doi.org/10.2514/1.G001303>.
- [44] N. T. Nguyen. *Model-Reference Adaptive Control: A Primer*. Advanced Textbooks in Control and Signal Processing. Springer International Publishing, 2018. ISBN: 978-3-319-56392-3. doi: 10.1007/978-3-319-56393-0. URL: <https://www.springer.com/gp/book/9783319563923>.
- [45] A. V. Efremov, M. S. Tiaglik, and I. I. K. "Pilot behavior model in pursuit and preview tracking tasks". In: *AIAA Scitech 2021 Forum*. AIAA-2021-1135. 2021. doi: 10.2514/6.2021-1135.
- [46] N. T. Nguyen et al. "Bounded Linear Stability Analysis - A Time Delay Margin Estimation Approach for Adaptive Control". In: Aug. 1, 2009. doi: 10.2514/6.2009-5968.
- [47] N. T. Nguyen and E. Summers. "On Time Delay Margin Estimation for Adaptive Control and Robust Modification Adaptive Laws". In: Aug. 1, 2011. doi: 10.2514/6.2011-6438.
- [48] Y. Huang et al. "Time-Delay Margin and Robustness of Incremental Nonlinear Dynamic Inversion Control". In: *Journal of Guidance, Control, and Dynamics* (2021). doi: 10.2514/1.G006024.
- [49] E. Fridman and S.-I. Niculescu. "On Complete Lyapunov–Krasovskii Functional Techniques for Uncertain Systems with Fast-Varying Delays". In: *International Journal of Robust and Nonlinear Control* 18.3 (Feb. 2008), pp. 364–374. ISSN: 10498923, 10991239. doi: 10.1002/rnc.1230. URL: <http://doi.wiley.com/10.1002/rnc.1230> (visited on 01/12/2021).
- [50] E. Fridman. "Tutorial on Lyapunov-Based Methods for Time-Delay Systems". In: *European Journal of Control* 20.6 (Nov. 1, 2014), pp. 271–283. ISSN: 0947-3580. doi: 10.1016/j.ejcon.2014.10.001. URL: <http://www.sciencedirect.com/science/article/pii/S0947358014000764>.
- [51] M. Mulder et al. "Manual Control with Pursuit Displays: New Insights, New Models, New Issues". In: *IFAC-PapersOnLine*. 14th IFAC Symposium on Analysis, Design, and Evaluation of Human Machine Systems HMS 2019 52.19 (Jan. 1, 2019), pp. 139–144. ISSN: 2405-8963. doi: 10.1016/j.ifacol.2019.12.125. URL: <http://www.sciencedirect.com/science/article/pii/S2405896319319603>.
- [52] M. M. Van Paassen, O. Stroosma, and J. Delatour. "DUECA - Data-Driven Activation in Distributed Real-Time Computation". In: *AIAA Modeling and Simulation Technologies Conference*. Aug. 14, 2000. doi: 10.2514/6.2000-4503.
- [53] O. Stroosma et al. "Measuring Time Delays in Simulator Displays". In: *AIAA Modeling and Simulation Technologies Conference*. Aug. 20, 2007. ISBN: 978-1-62410-160-1. doi: 10.2514/6.2007-6562.
- [54] P. M. T. Zaal, D. M. Pool, and M. Mulder. "Linear Mixed-Effects Models for Human-in-the-Loop Tracking Experiment Data". In: *AIAA Scitech 2021 Forum*. AIAA-2021-1014. 2021. doi: 10.2514/6.2021-1014.
- [55] K. van der El et al. "An Empirical Human Controller Model for Preview Tracking Tasks". In: *IEEE Transactions on Cybernetics* 46.11 (Nov. 2016), pp. 2609–2621. doi: 10.1109/TCYB.2015.2482984.
- [56] A. Y. Ng and S. Russell. "Algorithms for Inverse Reinforcement Learning". In: *In Proc. 17th International Conf. on Machine Learning*. Morgan Kaufmann, 2000. pp. 663–670.