

Modeling the Human Operator's Detection of a Change in Controlled Element Dynamics

Martin Barragan

Modeling the Human Operator's Detection of a Change in Controlled Element Dynamics

by

Martin Barragan

to obtain the degree of Master of Science
at the Delft University of Technology,

Student number: 4832191
Project duration: October 2022 – October 2023
Supervisors: Prof. dr. ir. M. Mulder
Dr. ir. M. M. van Paassen
Dr. ir. D. M. Pool

Preface

This report contains the work I have done over the past year, starting with the scientific article and followed by the preliminary thesis report. It has undoubtedly been the most difficult project I have worked on in my life, with many challenges along the way, but the end result is very rewarding. I would like to say a big thank you to my supervisors Daan, Max, and Rene. Without your guidance and great ideas, this work would not have been possible. Also, a big thank you to my family, who have always been very supportive of me and helped motivate me during the challenging times. Finally, a big thank you to my friends here in Delft, who made my time as a student a lot of fun.

I hope you enjoy reading this report!

Martin Barragan
Delft, October 2023

Contents

Preface	iii
Nomenclature	xiii
I Scientific Article	1
II Preliminary Thesis Report	35
1 Introduction	37
2 Relevant Concepts for Manual Control	39
2.1 Internal Models	39
2.2 Adaptive Model Theory	40
2.3 McRuer's Crossover Model	41
3 Review of Existing Models	43
3.1 Models Investigating the Detection Phase	43
3.1.1 Young and Stark Model	43
3.1.2 Phatak and Bekey Model	46
3.1.3 Niemela and Krendel Model	50
3.1.4 Hess Model	50
3.2 Models Without Explicit Boundaries for Detection	53
3.2.1 Recursive Autoregressive-Exogeneous Model	53
3.2.2 Artificial Neural Networks	55
3.2.3 Model Reference Adaptive Control	55
3.3 Conclusions	57
4 Simulation Results	59
4.1 Simulation Design	59
4.1.1 Forcing Function	59
4.1.2 Controlled Element Dynamics	60
4.1.3 Human Operator Dynamics	62
4.1.4 Noise Dynamics	63
4.1.5 Simulink Implementation	64
4.1.6 Simulation Verification	64
4.2 Simulation Results	67
4.2.1 Methodology	67
4.2.2 DYN12 Results	68
4.2.3 DYN21 Results	69
4.3 Conclusions	73
5 Future Work	77
5.1 Experiment Design	77
5.1.1 Control Task	77
5.1.2 CE Dynamics	77
5.1.3 Forcing Functions	78
5.1.4 Participants	78
5.1.5 Hypotheses	79
5.2 Analysis to be Done	79

References	81
A Experiment Forcing Functions	85
B Padé Approximation in Time-Varying Simulation	89
III Scientific Article Appendices	93
A Full Simulation Analysis Results	95
B Participant 8 Data	99
C Performance Metrics Sorted by Forcing Function and Order in Experiment	101
D Details of Experiment FPs and FN	105
E Including u in Model Threshold	107
F DYN21 Transition Time Trace Analysis	111
G Experiment Latin Squares	115
H Documents for Human Research Ethics Committee	117

List of Figures

2.1	Diagram of inverse and forward models.	39
2.2	Block diagram for the compensatory tracking system.	41
3.1	Block diagram for Young and Stark's model [22].	44
3.2	Experimental results from Miller and Elkind [23] showing the strong correlation between ΔC and $\Delta \epsilon$ at the moment participants detected the change in CE dynamics.	45
3.3	Relationship between when participants detected the change in CE dynamics and $\Delta \epsilon$ at that moment [23].	45
3.4	Setup of Young's experiment [24].	46
3.5	Block diagram for the compensatory tracking task of Phatak and Bekey [25].	46
3.6	Supervisory control algorithm as defined by Phatak and Bekey [25].	48
3.7	Phase plane showing the moment the HO detected a change in the CE dynamics with DR-1 superimposed [26].	49
3.8	Block diagram for Niemela and Krendel's compensatory tracking experiment [28].	50
3.9	Location of the detection boundary resulting from Niemela and Krendel's second experiment [28].	51
3.10	Block of diagram of Hess' adaptive model for pursuit tracking tasks [30].	52
3.11	Block diagram for Plaetinck's compensatory tracking experiment [35].	54
3.12	Results of the ANN-based identification procedure [36].	56
3.13	Block diagram of Terenzi's MRAC model [37].	56
4.1	Block diagram for the compensatory tracking system in the simulation.	60
4.2	Hypothetical SAS implementation.	61
4.3	Bode plots for the two CE dynamics conditions.	62
4.4	Implementation of the compensatory tracking simulation in Simulink. The top half is the simulation for the adaptive pilot, while the bottom half is the constant pilot.	65
4.5	Schematic showing how the second step of the analysis was done for the DYN12 condition.	68
4.6	Schematic showing how the second step of the analysis was done for the DYN21 condition.	69
4.7	Comparison of the rms of each of the signals in the steady-state pre-transition phase with the adaptive pilot in the expected detection period, and the adaptive pilot in the steady-state post-transition phase for the DYN12 condition.	70
4.8	Number of TPs, FPs, and FNs for the constant pilot with various thresholds in the DYN12 condition.	71
4.9	Comparison of the rms of each of the signals in the steady-state pre-transition phase with the constant pilot in the expected detection period, the adaptive pilot in the expected detection period, and the adaptive pilot in the steady-state post-transition phase for the DYN21 condition.	72
4.10	Spread in rms for the constant pilot in the detection region over all 100 forcing function realizations.	73
4.11	Accuracy plotted as a function of the threshold (multiple of steady-state standard deviation) and window length for the three signals that were identified to be most relevant for the detection phase.	74
4.12	Number of TPs, FPs, and FNs over varying window sizes for the constant pilot in the DYN21 condition and a threshold of 0.5σ	75
5.1	Sketch of the pursuit display to be used in the experiment.	78
A.1	Distribution of power in the forcing function over a two-second window.	86
A.2	Q-Q plot of a forcing function that will be used in the experiment.	86

B.1	Comparison of the output of the MATLAB only model and the model using Simulink in the DYN1 condition.	90
B.2	Difference between the outputs of the two models for the DYN1 condition.	90
B.3	Transient in the control input and output as a result of using a third order Padé approximation in the DYN12 condition.	91
B.4	Transient in the control input and output when the time delay is kept constant in the DYN12 condition.	91
B.5	Transient in the control input and output in the DYN21 condition.	92
B.6	Control input and output in the transition region for the final model used in this report. .	92
A.1	Comparison of the standard deviations of each of the signals in the steady-state pre-transition phase with the adaptive HO in the detection period and the adaptive HO in the steady-state post-transition phase for the DYN12 condition. $N = 135$ in all of the box plots.	96
A.2	Accuracy of a model based on each of the candidate signals in the DYN12 condition. .	97
A.3	Comparison of the standard deviations of each of the signals in the steady-state pre-transition phase with the constant HO in the detection period, the adaptive HO in the detection period, and the adaptive HO in the steady-state post-transition phase for the DYN21 condition. $N = 135$ in all of the box plots.	98
B.1	RMSE over the six time-invariant trials in DYN1 and DYN2.	99
B.2	Crossover frequencies calculated using the first 30 s of the measurement time for the six time-invariant trials in DYN1 and DYN2.	100
B.3	Detection lags for the six DYN12 and DYN21 trials that did not use the validation forcing functions.	100
C.1	RMSE sorted by forcing function in the DYN1 and DYN2 trials.	101
C.2	RMSE sorted by the order in which the DYN1 and DYN2 trials were performed.	102
C.3	Crossover frequencies sorted by forcing function in the DYN1 and DYN2 trials.	102
C.4	Crossover frequencies sorted by the order in which the DYN1 and DYN2 trials were performed.	102
C.5	Detection lags sorted by the order in which the DYN12 and DYN21 trials were performed.	103
E.1	Comparison of the model's TPs with the participant's TPs in the original model without u in the threshold ($w = 0$).	108
E.2	Comparison of the model's TPs with the participant's TPs over various weights w when considering a 0.2 s window for calculating the power in u	108
E.3	Comparison of the model's TPs with the participant's TPs over various weights w when considering a 0.5 s window for calculating the power in u	109
E.4	Comparison of the model's TPs with the participant's TPs over various weights w when considering a 1 s window for calculating the power in u	110
F.1	Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 1. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 1.	112
F.2	Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 3. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 3.	112
F.3	Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 4. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 4.	112

F.4	Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 6. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 6.	113
F.5	Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 7. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 7.	113
F.6	Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 9. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 9.	113

List of Tables

3.1	Time taken (in seconds) for the three groups of participants in Young's experiment to detect that the CE dynamics had changed for a transition from single to double integrator and vice versa [24].	44
3.2	Pre-failure tracking standard deviation compared to DR-1 values for all participants in van Ham's experiment [26].	49
3.3	Accuracy and detection lags for the methods analysed by Plaetinck [35], compared with the HO performance.	54
4.1	Parameters of the ten sine waves used to construct the forcing function.	60
4.2	Parameters of the CE dynamics for the two dynamics cases.	61
4.3	Parameters of the HO.	63
4.4	Phase margin and crossover frequency for the constant pilot model post-transition in the time-varying conditions.	63
4.5	Values of the noise filter gain k_n	64
4.6	Comparison of the rms values in steady-state conditions for the two simulations.	66
4.7	Comparison of the rms values in the time-varying conditions with both inputs.	67
A.1	Phase shifts and power in the two-second window centered around the transition for the nine forcing functions to be used in the experiment.	87
D.1	Summary of the details surrounding the FPs and the FN in the experiment.	105

Nomenclature

Abbreviations

Abbreviation	Definition
AMT	Adaptive Model Theory
ANN	Artificial Neural Network
ARX	Autoregressive Exogeneous
CE	Controlled Element
FN	False Negative
FP	False Positive
HO	Human Operator
MA	Moving Average
MIMO	Multiple-input Multiple-output
MRAC	Model Reference Adaptive Control
rms	Root Mean Square
SA	Sensory Analysis
SAS	Stability Augmentation System
SLLANN	Single Layer Linear Artificial Neural Network
TICA	Time-Invariant Condition Average
TN	True Negative
TP	True Positive
RE	Response Execution
RP	Response Planning
VAF	Variance Accounted For

Symbols

Symbol	Definition	Unit
A_t	Amplitude	[-]
e	Error	[rad]
\dot{e}	Error rate	[rad/sec]
f_t	Forcing function	[rad]
G	Maximum rate of change	[1/sec]
k_c	Controlled element gain	[-]
k_n	Noise gain	[-]
k_p	Human operator gain	[-]
M	Moment of maximum rate of change	[secs]
n_t	Integer multiple of fundamental frequency	[-]
T_L	Lead time constant	[secs]
T_n	Noise time constant	[secs]
u	Control input	[rad]
\dot{u}	Control input rate	[rad/sec]
y	Output	[rad]
\dot{y}	Output rate	[rad/sec]
ω_b	Controlled element break frequency	[rad/sec]
ω_{nm}	Neuromuscular system natural frequency	[rad/sec]
φ_m	Phase margin	[deg]

Symbol	Definition	Unit
ϕ_t	Phase shift	[rad]
σ	Standard deviation	[-]
σ^2	Variance	[-]
τ_e	Human operator effective time delay	[sec]
ζ_{nm}	Neuromuscular system damping ratio	[-]

Part I

Scientific Article

Modeling the Human Operator's Detection of a Change in Controlled Element Dynamics

Author: M. Barragan; Supervisors: M. Mulder, M. M. van Paassen, D. M. Pool,
Control & Simulation Section, Faculty of Aerospace Engineering
Delft University of Technology, Delft, The Netherlands

While human control behavior is well-understood in continuous control tasks, little is known about how human operators detect sudden changes in the controlled element dynamics. This paper focuses on modeling this detection phase for pursuit tracking tasks. Potential triggers for the human operator to detect changes in the controlled element dynamics were investigated via a time-varying computer simulation. Based on the results, hypotheses were generated and later tested in a single-axis pursuit tracking experiment with fifteen participants, where good-quality data were collected. Transitions from approximate single to approximate double integrator dynamics and vice versa were investigated, and participants indicated if they detected the transition by pressing a button. Using the button push data, a model for each transition was developed and validated. The models work under the assumption that human operators use a threshold, a multiple of the steady-state standard deviation, on certain signals to detect transitions. The models developed for the transition from single to double integrator dynamics and vice versa are based on the tracking error and system output acceleration, respectively. They have an accuracy of 88.9% and 99.4%, respectively. However, the estimation of the detection lags remains a limitation of both models. Nonetheless, this research helped confirm the tracking error can be used in a model for the transition from single to double integrator dynamics, proposed a model for the opposite transition, and identified that the relationship between control inputs and the system's response is an important factor in the detection phase.

Nomenclature

$A_i [i]$	=	Amplitude of i^{th} sine wave in forcing function (rad)
e	=	Error (rad)
\dot{e}	=	Error rate (rad/s)
\ddot{e}	=	Error acceleration (rad/s ²)
f_i	=	Forcing function (rad)
FN	=	False negative
FP	=	False positive
G	=	Maximum rate of change (s ⁻¹)
$H_c(s, t)$	=	Controlled element transfer function (-)
$H_{del}(s, t)$	=	Human operator delay transfer function (-)
$H_n(s, t)$	=	Remnant filter transfer function (-)
$H_p(s, t)$	=	Human operator transfer function (-)
i	=	Sine wave index (-)
IQR	=	Interquartile range (-)
k_c	=	Controlled element gain (-)
k_n	=	Remnant filter gain (rad)
k_p	=	Human operator gain (-)
k	=	Multiple of base frequency (-)
M	=	Centered on time (s)
N_t	=	Number of sine waves used to construct forcing functions (-)
N	=	Number of samples (-)
n	=	Human operator remnant (rad)
P_1	=	Initial parameter value (-)

P_2	=	Final parameter value (-)
RMSE	=	Root-mean-square error (deg)
T_L	=	Human operator lead time constant (s)
T_m	=	Measurement time (s)
T_n	=	Remnant filter time constant (s)
t	=	Time (s)
TN	=	True negative
TP	=	True positive
u	=	Control input (rad)
\dot{u}	=	Control input rate (rad/s)
w	=	Gaussian white noise input (-)
y	=	System output (rad)
\dot{y}	=	System output rate (rad/s)
\ddot{y}	=	System output acceleration (rad/s ²)
ω_b	=	Controlled element break frequency (rad/s)
ω_m	=	Fundamental frequency (rad/s)
ω_{nm}	=	Neuromuscular system natural frequency (rad/s)
$\omega_t[i]$	=	Frequency of i^{th} sine wave in forcing function (rad/s)
$\phi_t[i]$	=	Phase shift of i^{th} sine wave in forcing function (rad)
ρ	=	Pearson correlation coefficient (-)
σ	=	Standard deviation (-)
τ_e	=	Effective time delay (s)
ζ_{nm}	=	Neuromuscular system damping ratio (-)

I. Introduction

Over the past decades, the role of automation has increased significantly in society. In the context of control tasks, this has led to human operators taking on more of a supervisory role. A situation where interventions made by the human operator are often in moments where adaptability, accuracy, and quick-thinking are required. Adaptability, in particular, is currently one of the main strengths humans have over automation. However, with the advancement of technology and development in how human operators interact with systems, the models used to describe their behavior have failed to keep up [1]. It is beneficial to understand and model the adaptive behavior of human operators for several reasons. First, it would help improve the design of current systems, since different design choices can be tested through simulation to identify potential problems. Additionally, adaptive models can be used to create better support systems for human operators. In the case of failures, for example, this can help reduce the number of loss of control incidents, which are still one of the largest worldwide contributors to fatal accidents [2]. Finally, it would allow for the development of safer and improved vehicle designs [3].

Young et al. [4] identified that the adaptive process of a human operator consists of three phases, namely detection, identification, and modification. In the detection phase, the human operator realizes that something has changed and that adaptation is required. This triggers the identification phase to begin, in which the human operator identifies the nature of the change. Finally, in the modification phase, the human operator modifies their control behavior to suit the new situation. It can be argued that the identification and modification phases happen simultaneously, though the detection phase is certainly distinct. A good adaptive model would be capable of approximating the human operator's behavior throughout all three phases. However, the scope of this paper is limited to the detection phase, specifically the case where there is a change in the controlled element dynamics.

Significant research efforts have been dedicated to studying the adaptive nature of human operators and how they are able to detect changes to the controlled element dynamics [5][6][7]. Several models have been proposed, especially for compensatory tracking tasks in which the controlled element dynamics transition from more stable to less stable dynamics, that use a threshold on either the tracking error or error rate [8][9]. The motivation behind these models is that well-trained human operators keep track of the statistical properties of those signals and are thus able to identify when either becomes abnormally large. Furthermore, human operators are able to predict the value of a particular signal in the (near) future via an internal model [10][11] of the controlled element dynamics and knowledge of the control inputs. A significant mismatch between the observed and predicted value can also lead to the human operator detecting a change in controlled element dynamics. Building on this, Hess [12][13] proposed to model the detection phase in

pursuit tracking tasks based on a threshold of a signal proportional to the error between the observed system output rate and the (internally generated) desired system output rate.

Thus far the majority of research efforts have focused on the transition from more stable to less stable dynamics (e.g., from approximate single to approximate double integrator dynamics). Modeling the human operator's detection of this transition based on outliers (i.e., extreme peaks) in the error or error rate has been successful [8][14]. However, the same cannot be said for a transition the other way around, because the system does not become unstable if the current control strategy is kept, and thus the error does not grow abnormally large [15][16]. This suggests that only part of the detection phase is properly understood. Furthermore, the majority of research efforts thus far have studied the detection phase for tracking using compensatory displays.

In real-life applications, compensatory displays are the exception rather than the norm. Studying the detection phase with a pursuit display, where more information is available to the human operator, can increase the applicability of the research. The presence of additional information also significantly impacts the control strategies available to the human operator [17] and can therefore also impact the way changes to the controlled element dynamics are detected. For example, the explicit knowledge of the relationship between the control inputs given by the human operator and the system's response can aid with the detection phase when using a pursuit display. The difference between compensatory tracking tasks and pursuit tracking tasks is illustrated in Figure 1. The same signals appear in both Figure 1a and Figure 1b, namely the forcing function $f_t(t)$, error $e(t)$, human operator remnant $n(t)$, control input $u(t)$, and the system output $y(t)$. However, the human operator is modeled differently in both as a result of only the error being displayed with a compensatory display, whereas the forcing function and system output are both displayed with a pursuit display. For any controlled element dynamics $H_c(s, t)$, the human operator can only respond to the error (modeled as $H_{p,e}(s, t)$) when using a compensatory display. With a pursuit display, however, the human operator can also respond to the forcing function (modeled as $H_{p,ft}(s, t)$), and the system output (modeled as $H_{p,y}(s, t)$) in addition to the error [17]. Thus, it is clear that the control behavior and the way changes to the controlled element dynamics are detected can be very different with a pursuit display compared to a compensatory display.

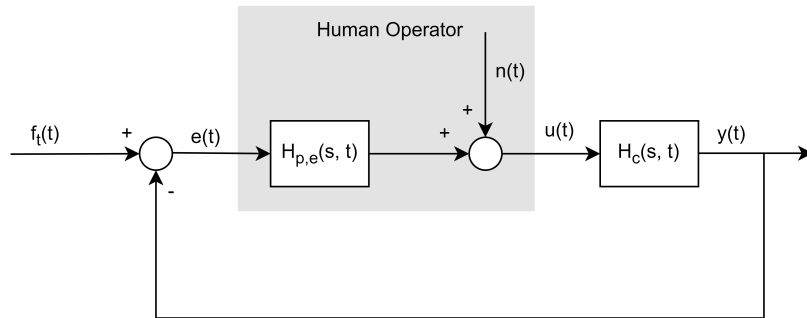

The goal of this paper is to investigate and provide insight into how human operators are able to detect changes in the controlled element dynamics and consequently adapt their control behavior when using pursuit displays. Particular attention is paid to the transition from approximate double to approximate single integrator dynamics since this is less well understood. First, a simulation was developed to investigate what the trigger for a human operator to detect a change in controlled element dynamics might be. This involved analyzing six candidate signals (the system output y , the control input u , the tracking error e , and all of their derivatives), and evaluating potential models that are based on when the respective signal exceeds a certain threshold (a multiple of the steady-state standard deviation). Afterwards, a human-in-the-loop pursuit tracking task experiment was performed at the Faculty of Aerospace Engineering, TU Delft in which several transitions from approximate single integrator dynamics to approximate double integrator dynamics and vice versa were conducted. Participants indicated if and when they detected the change by pressing a button on a joystick. The goal was then to find a connection between the subjective button push data and some property in one or more of the signals that may have triggered participants to detect that the controlled element dynamics changed. Ultimately, two models are proposed and validated, one for each transition.

The structure of the paper is as follows. First, the design and results from a time-varying computer simulation will be presented in Section II. Next, the experiment design and data analysis methodology will be presented in Section III. The results of the experiment data are presented in Section IV, followed by a discussion of the results in Section V. Finally, the conclusions are presented in Section VI.

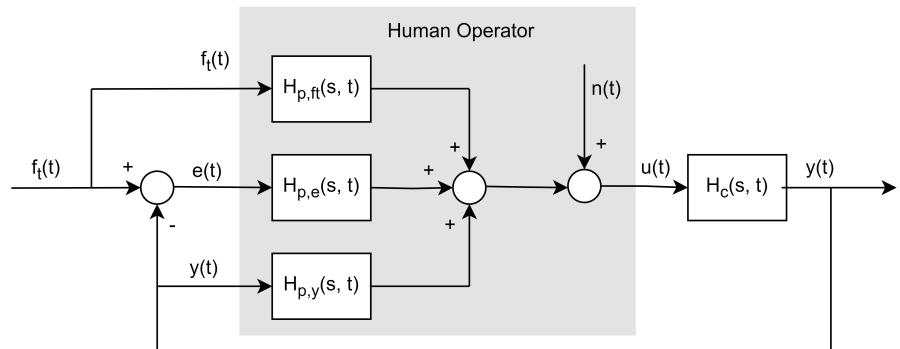

II. Computer Simulation

A. Simulation Design

Prior to performing the experiment, a time-varying simulation was implemented using MATLAB and Simulink to generate hypotheses for the experiment and gain a better understanding of what might trigger a human operator (HO) to detect a change in controlled element (CE) dynamics. To do this, two HO models were compared with each other following a sudden change in CE dynamics. One of the HO models is of an adaptive HO that immediately adapts their own dynamics to fit the new CE dynamics after a transition, while the other is of a constant HO that never adapts their own dynamics and continues controlling the new CE in the same way as before the transition. The latter is representative of a HO who has not (yet) detected that there has been a change in CE dynamics. Comparison between the signals of the adaptive HO and constant HO can then be used to generate hypotheses on what triggers a HO to detect a transition.



(a) Compensatory display (left) and control diagram (right).



(b) Pursuit display (left) and control diagram (right). The control diagram is adapted from Mulder et al. [17].

Fig. 1 Difference between compensatory and pursuit displays, and the accompanying control diagrams.

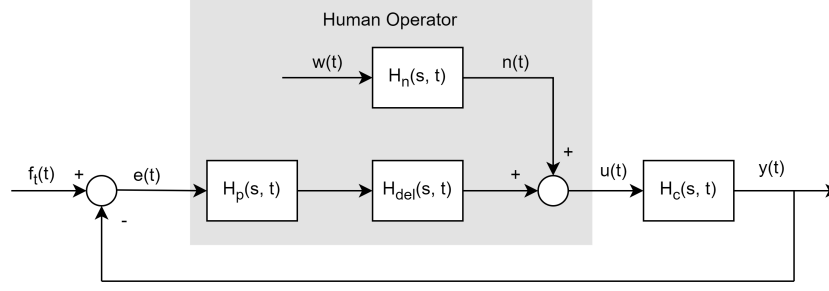


Fig. 2 Block diagram for the compensatory tracking system in the simulation.

The naming convention for the CE dynamics is as follows: the numbers are the order of the system and the sequence is the order in which the conditions occur [15]. For example, DYN12 is a trial where the CE dynamics started off as an approximate single integrator (first order) and transitioned to an approximate double integrator (second order) during the trial, while DYN21 has the same conditions but with the order reversed. A time-invariant trial only has one number. For example, DYN1 is a trial where the CE dynamics remained an approximate single integrator throughout the entire trial. The simulation was primarily used to investigate the time-varying conditions, DYN12 and DYN21. The time-invariant trials, DYN1 and DYN2, were only used to calculate steady-state reference values to be used in the analysis, as will be described in Subsection II.B.

Despite the fact that a pursuit display was used in the experiment, the simulation is made for compensatory tracking because there is a universally accepted model for it, namely the simplified precision model [18]. The block diagram for the compensatory tracking simulation is given in Figure 2. $H_p(s, t)$ is the linear part of the HO dynamics (including the neuromuscular system dynamics but excluding the time delay), $H_{del}(s, t)$ is the effective HO time delay, $H_n(s, t)$ is the remnant filter dynamics used to model the HO remnant, and $H_c(s, t)$ is the CE dynamics.

The structure of $H_c(s, t)$ is given in Eq. (1), where k_c is the CE gain and ω_b is the CE break frequency, both of which are time-varying. The break frequency is used to change the order of the CE dynamics in the crossover region. For the DYN1 condition, ω_b is set to 20 rad/s, whereas for the DYN2 condition, it is set to 0.2 rad/s. The gain k_c is varied to keep the level of control activity approximately constant for both conditions and has a value of 15 in DYN1 and 2 in DYN2.

$$H_c(s, t) = \frac{k_c(t)}{s(s + \omega_b(t))} \quad (1)$$

The HO dynamics are given in Eq. (2) and have the form dictated by the simplified precision model plus a term for the neuromuscular system dynamics [18]. For the results presented here, the HO gain k_p was tuned after performing the experiment to match the mean DYN1 and mean DYN2 crossover frequencies from the experiment. This results in a k_p of 0.233 for DYN1 and 3.73 for DYN2. The lead time constant T_L is 0 s for DYN1 (since no lead generation is required) and 5 s for DYN2, to perfectly compensate for the second integrator that has a break frequency of 0.2 rad/s. The neuromuscular system parameters were kept constant for both conditions with a natural frequency ω_{nm} of 15 rad/s and a damping ratio ζ_{nm} of 0.7. $H_{del}(s, t)$ accounts for the HO effective time delay, $e^{-\tau_e(t)s}$, and was modeled using the 'variable time delay' block in Simulink, with a delay of 0.09 s in DYN1 and 0.23 s in DYN2.

$$H_p(s, t) = k_p(t)[1 + T_L(t)s] \cdot \frac{\omega_{nm}^2}{s^2 + 2\zeta_{nm}\omega_{nm}s + \omega_{nm}^2} \quad (2)$$

Finally, the remnant filter dynamics have the form given in Eq. (3), which Van Grootheest et al. [19] found to be the best structure for fitting the remnant. The time constant T_n was kept constant at 0.06 s, and the gain k_n was chosen such that the ratio of the power in the control input due to the remnant to the total power in the control input was 0.2. This resulted in a value of k_n of 0.00787 for DYN1 and 0.00813 for DYN2. The signal $w(t)$ in Figure 2, which is Gaussian white noise with zero mean and unit variance, is the input to the remnant filter, and the HO remnant $n(t)$ is the output. All of the parameter values are summarized in Table 1.

$$H_n(s, t) = \frac{k_n(t)}{(T_n s + 1)^2} \quad (3)$$

For the time-varying conditions, the transition from one parameter value to the other is done according to the sigmoid function defined in Eq. (4), where P_1 is the initial parameter value, P_2 is the final parameter value, G is the maximum rate of change (equal to 100 s^{-1}), and M is the centered on time (equal to 75 s). These values of G and M apply to all of the time-varying parameters in both the CE dynamics and the adaptive HO model.

$$P(t) = P_1 + \frac{P_2 - P_1}{1 + e^{-G(t-M)}} \quad (4)$$

Table 1 Values for all of the parameters used in the simulation.

CE Dynamics	k_c (-)	ω_b (rad/s)	k_p (-)	T_L (s)	τ_e (s)	ω_{nm} (rad/s)	ζ_{nm} (-)	k_n (-)	T_n (s)	G (s^{-1})	M (s)
DYN1	15	20	0.233	0	0.09	15	0.7	0.00787	0.06	100	75
DYN2	2	0.2	3.73	5	0.23	15	0.7	0.00813	0.06	100	75

The resulting crossover frequencies and phase margins are given in Table 2. In addition to the standard DYN1 and DYN2 conditions, two further conditions arise after the change in CE dynamics as a result of the constant HO keeping the same control strategy. These are shown in the bottom two rows of Table 2. 'DYN1 CE with DYN2 HO' occurs post-transition in the DYN21 trials while 'DYN2 CE with DYN1 HO' occurs post-transition in the DYN12 trials. As can be seen, the system with the constant HO becomes unstable post-transition in DYN12, while the system with the constant HO remains stable post-transition in DYN21 but with a very low crossover frequency and large phase margin. The crossover frequencies in the standard DYN1 and DYN2 conditions are 2.77 rad/s and 2.33 rad/s, respectively, and the phase margins are 52.8 deg and 46.7 deg, respectively.

Table 2 Crossover frequencies and phase margins for the different conditions in the simulation.

Condition	ω_c (rad/s)	Phase Margin (deg)
DYN1	2.77	52.8
DYN2	2.33	46.7
DYN1 CE with DYN2 HO	0.359	143
DYN2 CE with DYN1 HO	2.73	-24.6

B. Simulation Analysis Methodology

The goal of the analysis of the simulation data is to identify differences in the properties of different signals throughout separate phases of the simulation, as this could be something that a HO can use to detect that there has been a change in CE dynamics. The candidate signals that will be analyzed are the system output y , the system output rate \dot{y} , the control input u , the control input rate \dot{u} , the error e , and the error rate \dot{e} . These were chosen because they are the signals that can either be explicitly seen or can be visually estimated by the HO during a trial.

Each trial in the simulation lasts 120 s. However, all of the forcing functions were designed to have a period of 30 s, meaning there are exactly four periods in each trial [14][15]. This is useful to compare different phases of the simulation because, if a whole period is considered, there is no dependence on the local forcing function properties. As such, each trial is divided into four 30 s phases, the first of which ($0 \text{ s} < t < 30 \text{ s}$) is used as run-in time, the second ($30 \text{ s} < t < 60 \text{ s}$) is the pre-transition steady-state phase, the third ($60 \text{ s} < t < 90 \text{ s}$) is the transition region, and the fourth ($90 \text{ s} < t < 120 \text{ s}$) is the post-transition steady-state phase [14]. The transition region is further narrowed down to a phase called the "detection period", which is 2.2 s - 7.4 s after the transition, and represents the period of time where a detection is expected [14]. Since the transition always occurs at $t = 75 \text{ s}$ (hence why $M = 75 \text{ s}$ in Table 1), the detection period is $77.2 \text{ s} < t < 82.4 \text{ s}$.

A total of nine forcing functions were analyzed with the simulation to ensure the conclusions are not specific to local forcing function properties. For each forcing function, fifteen remnant realizations were simulated, resulting in 135 trials in each condition (DYN1, DYN2, DYN12, and DYN21). The reasoning behind using specifically nine forcing functions and fifteen remnant realizations for each will become clear in subsubsection III.A.3 and subsubsection III.A.5.




Fig. 3 Illustration showing how the second step of the simulation data analysis is done.

The analysis of the results took place in two steps. In the first step, the standard deviation (σ) of the six candidate signals was compared over the different phases of the simulation. Of particular interest is the comparison of the pre-transition steady-state standard deviations with the standard deviations of the constant HO in the detection period, since this is representative of a HO that has not yet detected the change in CE dynamics. Thus, the goal in this step is to identify the signals that have the largest differences in their properties before and immediately after the transition, because that makes it likely that they are important in the process of detecting a change in CE dynamics.


In the second step of the analysis, a model is made using the signals identified to be the most relevant in the first step, and the accuracy of the model is evaluated. An illustration of the method used to evaluate the accuracy of the model is given in Figure 3. The model triggers (i.e., predicts the HO would detect a change in CE dynamics) whenever the instantaneous value of a particular signal exceeds a particular threshold, of which a whole range is evaluated [14]. The threshold is a multiple of the steady-state standard deviation of the respective signal, which is calculated from the time-invariant trials in the simulation (this is the only point in the simulation analysis where the time-invariant trials are used) [14]. The 135 trials in DYN12 and 135 trials in DYN21 are then evaluated. If, for a particular trial, the threshold is exceeded before the transition, as would be the case at time t_1 for a threshold of 2σ in Figure 3, it counts as a false positive (FP, "too early detection"). If the threshold is *not* exceeded before the transition and *is* exceeded in the detection period, as would be the case at time t_2 for a threshold of 3σ in Figure 3, it counts as a true positive (TP). Finally, if the threshold is never exceeded, as would be the case for a threshold of 4σ in Figure 3, it counts as a false negative (FN) [14]. The accuracy of the model can then be calculated using the equation below. Note that it is impossible to have a true negative (TN) in the simulation since only the time-varying trials were analyzed with the model, but it is included in Eq. (5) for completeness and because it will be needed when evaluating the experimental results in Section IV.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (5)$$


C. Simulation Results

1. Sample Time Traces


To get a better feel for how the simulation results look in the time domain, sample time traces for each of the six candidate signals in each of the four CE dynamics conditions are presented in Figure 4-6 (for one forcing function and remnant realization). The different phases of the simulation introduced in Subsection II.B can be seen by the shaded regions in the plots. Furthermore, a sample sigmoid function (Eq. (4)) is also shown in Figure 4a, where it can be seen that the transition is almost instantaneous. Note that there is only one line plotted in each of the figures before the transition because the adaptive HO and the constant HO are the same during that phase of the simulation. After the transition, the signals are shown for both HO models. In DYN12, the constant HO signals are only shown for 5 s




(a) y DYN1 and DYN12.



(b) y DYN2 and DYN21.



(c) \dot{y} DYN1 and DYN12.



(d) \dot{y} DYN2 and DYN21.

Fig. 4 Comparison of sample time traces for y and \dot{y} in the four CE dynamics conditions, and between the constant HO and the adaptive HO post-transition.


following the transition because they become unstable and start congesting the figures.

Several main observations can be made. First, it can be seen that the tracking performance is better in DYN1 than in DYN2, as expected due to the more difficult nature of the DYN2 condition. Second, the adaptive HO is able to quickly adapt to the CE dynamics post-transition and resume steady-state tracking. Finally, the constant HO post-transition in DYN21 gives very aggressive control inputs (Figure 5b), resulting in large (and quickly oscillating) values for \dot{y} (Figure 4d) that lead to oscillations in y (Figure 4b). Thus, despite the the system with the constant HO not becoming unstable post-transition, the tracking performance is still significantly worse than that of the adaptive HO.


2. DYN12 Results

Next, the results for DYN12 will be presented. In the first step of the analysis (as described in Subsection II.B), a significant difference between the standard deviations in the pre-transition steady-state phase and the standard deviations of the constant HO in the detection period was found for all of the candidate signals since the system with the constant HO becomes unstable after the transition. For this reason and due to space constraints, the figures from the first step of the analysis are not presented here. The full results can be found in Appendix A.


Since it is impossible to tell which of the six candidate signals is the first to diverge, e and \dot{e} will be used in the second step of the analysis. This is because previous research efforts (analyzing compensatory displays) have been successful in modeling the DYN12 transition using e and \dot{e} [8][9], and large increases to either is something that a HO




(a) u DYN1 and DYN12.



(b) u DYN2 and DYN21.




(c) \dot{u} DYN1 and DYN12.




(d) \dot{u} DYN2 and DYN21.


Fig. 5 Comparison of sample time traces for u and \dot{u} in the four CE dynamics conditions, and between the constant HO and the adaptive HO post-transition.




(a) e DYN1 and DYN12.



(b) e DYN2 and DYN21.



(c) \dot{e} DYN1 and DYN12.



(d) \dot{e} DYN2 and DYN21.

Fig. 6 Comparison of sample time traces for e and \dot{e} in the four CE dynamics conditions, and between the constant HO and the adaptive HO post-transition

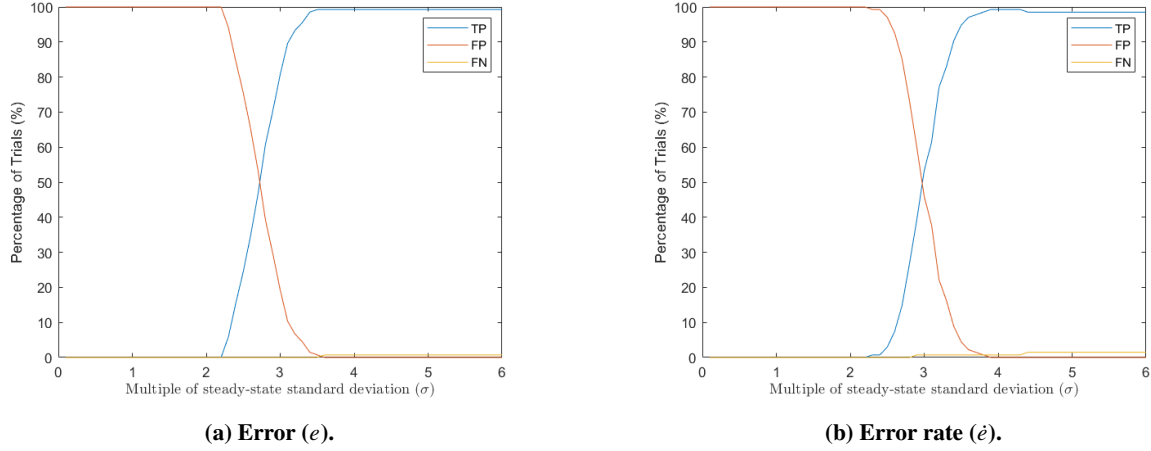


Fig. 7 Number of TPs, FPs, and FNs for the constant HO over various thresholds in the DYN12 condition.

can notice and use as a trigger to detect a change in CE dynamics. This is particularly true when using a compensatory display since they are the *only* two signals the HO can see (or estimate in the case of \dot{e}). The number of TPs, FPs, and FNs over different thresholds (i.e., different multiples of the steady-state standard deviation) are given in Figure 7. As can be seen, the number of FPs drops to zero for thresholds larger than 3.6σ for e and 3.9σ for \dot{e} (the same threshold on \dot{e} proposed by Van Ham [20]), while there is a TP in more than 99% of trials for a threshold larger than 3.5σ for e and more than 98.5% of trials for a threshold larger than 3.9σ for \dot{e} . Thus, above the threshold of 3.5σ and 3.9σ for e and \dot{e} , respectively, an accuracy greater than 98.5% can be achieved. This suggests that e and \dot{e} could be good signals off which to base a model for the DYN12 transition.

3. DYN21 Results

Finally, the results for DYN21 will be presented. In the first step of the analysis it was found that the biggest difference in standard deviations between the steady-state pre-transition phase and the constant HO in the detection period occurs for y and \dot{u} , as shown in Figure 8. The number of samples $N = 135$ in each box plot. As can be seen, there is a 56.2% decrease in the median standard deviation of y (Figure 8a) and a 30.4% increase in the median standard deviation of \dot{u} (Figure 8b). However, there is still overlap with the standard deviations of the pre-transition steady-state phase in the case of \dot{u} , which makes developing a model more difficult. It can also be seen that there is a large difference between the constant HO in the detection period, and the adaptive HO in the detection period and post-transition steady-state phase, suggesting that adaptation is still beneficial in the DYN21 transition, despite the system with the constant HO not becoming unstable post-transition.

In the second step of the analysis, the accuracy of a model over a range of thresholds was evaluated (according to Eq. (5)), with the results shown in Figure 9. As can be seen, there is no threshold for any of the six candidate signals that results in an accuracy greater than 31%. The highest accuracy is achieved for \dot{u} , as predicted from the first step of the analysis, but it is very low in comparison to the model for DYN12 based on e or \dot{e} . As alluded to earlier, the big overlap between the pre-transition steady-state standard deviations and the standard deviations in the detection period for the constant HO contributes significantly to the low accuracy. Regarding y , since the magnitude of the oscillations for the constant HO becomes smaller post-transition (Figure 4b), it is impossible to achieve an accuracy greater than 0% as there are no TPs for any threshold. Either the threshold is too low and results in a FP, or the threshold is too high and results in a FN. However, by introducing the error acceleration \ddot{e} (which is related to \dot{u} in DYN2) into the analysis, a model with an accuracy of 73.0% can be achieved with a threshold of 3.3σ . Visually estimating accelerations is difficult for humans [21], however, especially of e when using a pursuit display since the error itself has to be estimated. Thus, these results suggest that detecting the DYN21 transition will be significantly harder than detecting the DYN12 transition, assuming HOs detect the transition based on a threshold for a particular signal, which is simply an assumption at this stage.

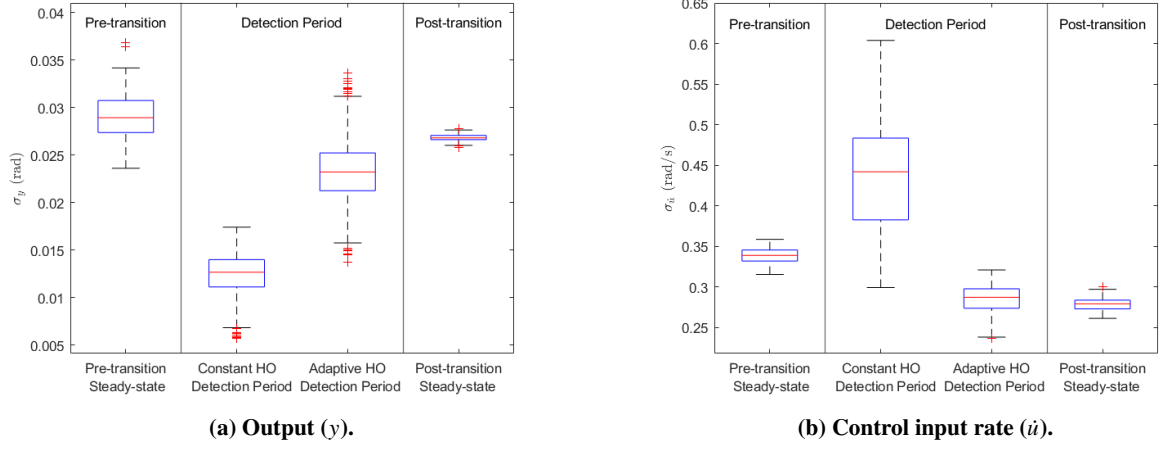


Fig. 8 Comparison of the standard deviations of y and \dot{u} in the steady-state pre-transition phase with the constant HO in the detection period, the adaptive HO in the detection period, and the adaptive HO in the steady-state post-transition phase for the DYN21 condition.




Fig. 9 Accuracy of a model based on each of the six candidate signals and the error acceleration \ddot{e} over various thresholds for the DYN21 transition.




Fig. 10 Diagram of the pursuit display used in the experiment.




Fig. 11 Set-up during the experiment, with the exception of the joystick used for the button push.

4. Simulation Analysis Conclusions

Overall, several conclusions can be made from the simulation results. First, it can be seen that in both conditions (DYN12 and DYN21) there are significant differences between the constant HO and the adaptive HO's control behavior. In DYN12, the system with the constant HO becomes unstable, while in DYN21 both HO models remain stable but the tracking performance of the adaptive HO is significantly better. Thus, detecting the transition and adapting to the new CE dynamics is crucial to maintaining good tracking performance. Second, for the DYN12 transition, a model based on e or \dot{e} was able to achieve a high accuracy ($>98.5\%$) suggesting these are good signals off which to base a model, which is consistent with previous research. [9][14]. On the other hand, detecting the DYN21 transition may be more difficult because there are no significant differences in any of the six candidate signals. Only for a model based on \ddot{e} could a high accuracy be achieved, but visually perceiving accelerations of a signal that has to be estimated itself is difficult for humans [21].

III. Method

A. Experiment Data

For developing and validating the model presented in this paper, a human-in-the-loop experiment was conducted at the Faculty of Aerospace Engineering, TU Delft. Details of the experiment will be provided here.

1. Control Task

The experiment consisted of 30 runs, excluding training runs, of a single-axis (horizontal) pursuit tracking task. A diagram of the display can be seen in Figure 10. The forcing function value was shown by the white square and the system output was shown by the green square, meaning participants controlled the green square and had to keep it as close to the white square as possible. The error $e(t)$ could be visually estimated by participants as the distance between the two squares. Participants were told that their primary goal was to keep the error as low as possible for the duration of the trial. If, at any moment, they believed the CE dynamics had changed, they were instructed to immediately press the trigger button on a Logitech joystick, separate from the one used to control the system. The decision to use a separate joystick was based on a recommendation from a previous experiment [14] in which it was found that reaction times for detecting a transition could increase if both the primary task of keeping the error as low as possible and the secondary task of pressing a button were done with the same hand on the same side-stick.

2. Apparatus

The experiment was conducted in a fixed-base simulator in the Human-Machine-Interaction Laboratory at the Faculty of Aerospace Engineering, TU Delft. The set-up during the experiment can be seen in Figure 11. Participants sat in the right seat and gave control inputs using their right hand. The side-stick could only be moved left and right during the experiment, which consequently made the CE output move to the left or to the right, respectively. The

maximum stick deflection was ± 15 deg. The center of the stick is 9 cm above the axis of rotation. It had a stiffness of $3.54 \text{ N m rad}^{-1}$, a damping coefficient of $0.22 \text{ N m s rad}^{-1}$, an inertia of 0.01 kg m^2 , and a breakout moment of 0 N m . The joystick used to indicate they noticed a change in CE dynamics (not pictured) was positioned on their left leg and help with their left hand.

3. Independent Variables

Two independent variables characterized each trial, namely the CE dynamics and the forcing function realization. The CE dynamics were varied to introduce time-varying behavior and allow for the investigation into the detection lags, and several forcing function realizations were used to investigate the effect of the forcing function properties around the transition moment, prevent predictability in the forcing function, and improve the generalization of the model [14].

Like in the simulation, there were a total of four CE dynamics conditions tested in the experiment, namely DYN1, DYN2, DYN12, and DYN21. The structure of the CE dynamics is the same as in the simulation and given in Eq. (1), and the beginning and end values of the gain k_c and break frequency ω_b are given in Table 3. Transitions between the beginning and end values of the two time-varying parameters were done according to the same sigmoid function as used in the simulation and given in Eq. (4). The value of G is 100 s^{-1} and M is 45 s (the middle of the measurement time).

Table 3 Parameters of the CE dynamics and sigmoid function for the four dynamics cases.


CE Dynamics	k_{c1} (-)	k_{c2} (-)	ω_{b1} (rad/s)	ω_{b2} (rad/s)	G (s^{-1})	M (s)
DYN1	15	15	20	20	-	-
DYN2	2	2	0.2	0.2	-	-
DYN12	15	2	20	0.2	100	45
DYN21	2	15	0.2	20	100	45

The forcing functions used in the experiment all follow the structure given in Eq. (6), where N_t is the number of sine waves used to construct the forcing function (ten in this experiment), and $A_t[i]$, $\omega_t[i]$, and $\phi_t[i]$ are the amplitude, frequency, and phase shift of the i^{th} sine wave, respectively. The amplitudes and frequencies used for each of the ten sine waves were kept the same for all forcing function realizations, with the phase shifts being the only difference. The amplitudes and frequencies are given in Table 4. It should be noted that, in order to avoid spectral leakage, the frequency of all of the sine waves must be an integer multiple k of the fundamental frequency, defined as $\omega_m = \frac{2\pi}{T_m}$, where T_m is the measurement time. For this purpose, the measurement time is taken to be 30 s.

$$f_t(t) = \sum_{i=1}^{N_t} A_t[i] \sin(\omega_t[i]t + \phi_t[i]) \quad (6)$$

Table 4 Parameters (excluding the phase shifts) of the ten sine waves used to construct the forcing functions.

n (-)	k (-)	ω_t (rad/s)	A_t (rad)
1	2	0.419	$2.905 \cdot 10^{-2}$
2	5	1.047	$1.916 \cdot 10^{-2}$
3	9	1.885	$1.020 \cdot 10^{-2}$
4	13	2.723	$6.032 \cdot 10^{-3}$
5	19	3.979	$3.356 \cdot 10^{-3}$
6	27	5.655	$1.983 \cdot 10^{-3}$
7	39	8.168	$1.230 \cdot 10^{-3}$
8	51	10.681	$9.331 \cdot 10^{-4}$
9	67	14.032	$7.541 \cdot 10^{-4}$
10	83	17.383	$6.674 \cdot 10^{-4}$



(a) Low power forcing functions (numbers 1, 2, and 3). (b) Medium power forcing functions (numbers 4, 5, and 6). (c) High power forcing functions (numbers 7, 8, and 9).

Fig. 12 All nine forcing functions used in the experiment, grouped by power level, in the 10 s window centered around the transition.

One of the reasons different forcing function realizations were used in the experiment is that previous experiments have found that the gradient of the forcing function around the transition moment can have a large impact on the detection lags [14][16]. Specifically, larger gradients are expected to lead to smaller detection lags. To quantify how large the gradients were around the transition moment, the power of the forcing functions was calculated for the two-second window centered around the transition. Consequently, three forcing functions were chosen that had low power, three that had medium power, and three that had high power, for a total of nine forcing function realizations used in the experiment (the same nine forcing functions were used in the simulation, hence the choice of specifically nine forcing functions in Subsection II.B). What counts as "low", "medium", and "high" power was determined by looking at the distribution of power over an arbitrary two-second window in 1,000 forcing functions that had an average crest factor (calculated from 100,000 randomly generated forcing functions). From this, it was concluded that a power between 0 rad^2 and $2.5 \cdot 10^{-5} \text{ rad}^2$ can be considered low power, between $2.5 \cdot 10^{-5} \text{ rad}^2$ and $7.5 \cdot 10^{-5} \text{ rad}^2$ can be considered medium power, and above $7.5 \cdot 10^{-5} \text{ rad}^2$ can be considered high power. The forcing functions, numbered 1-9, are ordered in terms of power in the two-second window centered around the transition, with one being the lowest power and nine being the highest power. Time traces of the forcing functions, grouped by power level (low, medium, or high), in the transition region can be seen in Figure 12. The shades of the colors in each subfigure correspond to the amount of power in the two-second window centered around the transition, with the lightest shade being the least power and the darkest shade being the most power. Three of the nine forcing functions were used to collect data exclusively for validation, namely the middle one in each power level (forcing function numbers 2, 5, and 8). Thus, the validation forcing function for each power level is the line that is neither the lightest nor the darkest shade in each subfigure. The phase shifts and power in the two-second window centered around the transition for each forcing function are given in Table 5.

4. Experiment Procedures

Each participant performed a total of 30 measurement trials (i.e., excluding training trials) in the experiment, each with a unique combination of CE dynamics and forcing function. For each of the three validation forcing functions, one trial was performed in each of the two time-varying conditions (DYN12 and DYN21), for a total of six trials. For the other six forcing functions, one trial was performed in each of the four conditions (DYN1, DYN2, DYN12, and DYN21), for a total of 24 trials. The experiment was split into two halves by combining the fifteen DYN1 and DYN12 trials into one half and the fifteen DYN2 and DYN21 trials into the other half. Within each half, the order of the conditions was determined via a random incomplete Latin square, so that participants could not predict whether a transition would occur in a particular trial.

The experiment started with several training runs in which participants could get familiar with the four CE dynamics conditions. Depending on the participant's performance, 1-5 training runs were done in each condition, usually two or three. If participants showed good tracking and detection performance on the first trial, only one additional training trial would be done in that condition, provided the performance remained stable. If the performance improved or was not at the level expected, an additional training run would be done in that condition until performance stabilized at a good level. All training runs used forcing function number 1.

Each trial lasted between 95 s and 105 s. The measurement time was always 90 s, but a random run-in time of either

Table 5 Phase shifts and power in the two-second window centered around the transition for the nine forcing functions used in the experiment. The forcing function numbers with "(V)" after the number are the ones used to collect validation data.

	Forcing Function Number								
	1	2 (V)	3	4	5 (V)	6	7	8 (V)	9
ϕ_1 , rad	6.065	3.454	0.885	4.315	5.608	3.671	1.040	3.219	6.172
ϕ_2 , rad	0.223	3.073	1.120	3.308	6.027	2.757	1.958	5.615	2.662
ϕ_3 , rad	4.232	5.440	4.806	0.988	3.106	5.220	2.037	4.049	0.903
ϕ_4 , rad	1.880	1.296	4.251	3.925	3.109	4.162	4.768	4.820	2.290
ϕ_5 , rad	4.757	4.062	1.260	6.259	5.506	3.685	1.522	4.160	6.112
ϕ_6 , rad	1.072	3.464	4.764	6.236	5.560	5.188	2.820	0.440	2.162
ϕ_7 , rad	2.754	0.153	5.692	0.609	5.668	5.261	1.520	2.525	4.821
ϕ_8 , rad	1.593	1.678	3.003	4.782	2.511	5.699	1.577	3.368	0.101
ϕ_9 , rad	5.404	4.585	0.116	2.776	1.259	5.566	0.114	0.213	6.090
ϕ_{10} , rad	4.101	0.583	1.338	4.126	2.843	0.776	0.755	1.461	1.853
Power Level	Low	Low	Low	Medium	Medium	Medium	High	High	High
Power, rad ²	$6.2 \cdot 10^{-6}$	$1.2 \cdot 10^{-5}$	$1.6 \cdot 10^{-5}$	$4.4 \cdot 10^{-5}$	$5.8 \cdot 10^{-5}$	$6.3 \cdot 10^{-5}$	$1.3 \cdot 10^{-4}$	$2.2 \cdot 10^{-4}$	$2.3 \cdot 10^{-4}$

5 s, 10 s, or 15 s was used to avoid participants anticipating the moment a transition would occur (since, if it was a time-varying condition, the transition would always occur 45 s into the measurement time) and to minimize the chance a participant could identify which forcing function was being used based on the starting position of the target. Despite the measurement time being 90 s, the period of all of the forcing functions was 30 s (hence why T_m in subsection III.A.3 is 30 s), which means there are exactly three periods of the forcing function in the measurement time. The first of those could be used to look at steady-state tracking in the first condition (e.g., DYN1 in DYN12), the second could be used to consider the transition period, and the third could be used to look at steady-state tracking in the second condition (e.g., DYN2 in DYN12).

If a participant made a large control error in any of the trials (due to a lack of concentration, for example) that trial would be re-done at the end of that half of the experiment. However, this only happened with one participant, where they had a FP in the first trial and two FPs in the second trial. For the remainder of that half of the experiment, the participant did not have any FPs or FNs and they stated the reason for the three FPs was a lack of concentration and sufficient knowledge of the CE dynamics (only one training run per time-varying condition was done with that participant).

5. Participants and Instructions

A total of sixteen participants took part in the experiment. One participant (Participant 8) had particularly poor tracking and detection performance and was thus excluded from the data presented in this paper, though that participant's data are given in Appendix B. Thus, the data from fifteen participants are used to generate the results presented here, hence why fifteen remnant realizations were simulated per forcing function in Subsection II.B. Each of the fifteen participants performed the experiment according to the order given in a different row of the Latin square to balance out order effects. As a result of the odd number of participants, and the inability to increase the number of participants to 30 due to time constraints, eight participants performed the DYN1 and DYN12 conditions first, while seven performed the DYN2 and DYN21 conditions first. All fifteen participants are university students or staff. They were chosen either based on the fact that they had experience with tracking tasks and had shown good performance previously, or had hobbies that generally translate to good performance in tracking tasks, such as playing video games that require quick reaction times (e.g., first-person shooting games) or simulator race car drivers.

Before starting the experiment, participants were briefed on the experiment procedure. Special attention was given to highlighting that their primary task throughout the experiment should be to keep the tracking error as low as possible, while the button push was a secondary task. However, in the event that they detected a change in CE dynamics, the button push could momentarily be elevated to the primary task. Participants were also told to push the button when they "think" the CE dynamics changed. After the training runs, no feedback was given to participants on their performance

other than the root-mean-square error (RMSE) score that was shown on the display.

This experiment was reviewed and approved by the Human Research Ethics Committee at TU Delft under application number 3307.

6. Hypotheses

A total of four hypotheses were developed for the experiment, as listed below.

- **H1:** In both time-varying conditions all participants will adapt their control strategy to avoid the degradation in tracking performance observed in the computer simulations (i.e., their behavior will be similar to that of the adaptive HO).
- **H2:** In both time-varying conditions participants will be quicker to detect transitions in the CE dynamics when the gradients in the forcing function around the moment of transition are larger (quantified by the amount of power in the two-second window centered around the transition).
- **H3:** In DYN12, a potential model with a high accuracy for predicting the HO's detection of a change in CE dynamics can be based on significant increases in the error and error rate signals.
- **H4:** In DYN21, a potential model with a high accuracy for predicting the HO's detection of a change in CE dynamics cannot be based on any of the six candidate signals ($y, \dot{y}, u, \dot{u}, e, \dot{e}$).

H1, **H3**, and **H4** were formulated based on the conclusions from the simulation analysis (subsubsection II.C.4), while **H2** is based on Van Ham's findings [14].

B. Data Analysis

1. Performance Metrics

There are three primary performance metrics analyzed from the experiment data, namely the RMSE, crossover frequencies, and detection lags. The first two relate purely to tracking performance, while the latter has to do with detection performance.

The RMSE is useful for two reasons, one of which is to verify the quality of the experiment data and ensure it is representative of skilled HO behavior. For this purpose, the RMSE was calculated over the entire measurement time for all of the trials in DYN1 and DYN2. The other reason is that the RMSE in the final 30 s of the measurement time in the time-varying trials can be used to accept or reject **H1**. If participants always adapt their control behavior following a transition, the distribution of RMSE values when looking at the final 30 s of a time-varying trial (such that it covers one period of the forcing function) should be similar to the distribution of the RMSE over the six time-invariant trials in the same condition.

The crossover frequencies are also useful to verify the quality of the experiment data. For each trial, they were calculated using either the first or last 30 s of the measurement time, again to cover one period of the forcing function. Therefore, DYN12 trials were considered DYN1 trials when using the first 30 s of the measurement time, and DYN2 trials when considering the last 30 s of the measurement. The same can be said for the DYN21 trials but reversed. The crossover frequencies were then found by calculating the crossover frequency of the HO and CE frequency response. The HO frequency response was found by dividing the frequency response of the control input by the frequency response of the error *at the frequencies* of the ten sine waves which make up the forcing function v_j , as given by Eq. (7). The corresponding phase margins were also calculated along with the crossover frequencies.

$$\hat{Y}_p(j\omega; v_j) = \frac{U(j\omega; v_j)}{E(j\omega; v_j)} \quad (7)$$

Finally, the detection lags were calculated by looking at how long it took between the transition (45 s into the measurement time) and the moment the button push occurred. This information is useful both to have a reference as to where to look in the time traces to determine what may have triggered the detection, as well as to be able to compare with the model predictions to determine the quality of the model.

2. Model with Multiple of Steady-State Standard Deviation as Threshold

The primary goal of the data analysis is to find a relationship between a certain property in one of the six candidate signals and the moment participants pressed the button to indicate they detected the change in CE dynamics. To do this a very similar method to what was used to develop the models for the simulation data (outlined in Subsection II.B) will

be followed. The model for the experiment data also uses a threshold based on the steady-state standard deviation of the respective signal and the moment it is exceeded is considered the moment the model predicts the HO would detect the change in CE dynamics.

Referring to Figure 3 and the accompanying explanation of it in Subsection II.B, there are only a few differences with the method that will be followed here. First, for the analysis of the experiment data, the "detection period" is no longer considered, meaning a TP can occur at *any* moment after a transition and before the end of the trial. Second, it is possible to have TNs here, since some of the experiment trials were time-invariant. For these cases, a TN occurs if the threshold is *never* exceeded during the trial, and a FP occurs if the threshold is exceeded at any moment during the trial. In the time-varying trials, TPs, FPs, and FNs are calculated in the same way as was explained in Subsection II.B (with the exception of TPs only being able to occur in the detection period, as just described). Finally, the only other difference with respect to the simulation data analysis is that for the experiment data, the reference steady-state standard deviation that is used to determine the threshold is calculated *per participant*, as opposed to using every participant's data. The accuracy of the model is then calculated using Eq. (5). As a second measure of the quality of the model, the (time) difference between the model detections (when the threshold is exceeded) and the participant detections (when participants pressed the button) was analyzed.

IV. Results

A. RMSE


To verify the quality of the experiment data, the RMSE sorted by participant over the six time-invariant trials of DYN1 is shown in Figure 13a, alongside the RMSE for the six time-invariant trials of DYN2 in Figure 13b. As a reference, the right-most box plot in each subfigure is the data for *all* participants combined. Therefore $N = 6$ for each box plot except the right-most in each subfigure, where $N = 90$. It is clear that there is a difference in tracking performance between participants. However, different skill levels are to be expected between different people in any task, and the differences in our dataset are not large enough to consider a specific (group of) participant(s) as outliers or unskilled. This is supported by the fact that there are no data points considered outliers in the box plot for all of the participant's data (right-most box plot) in Figure 13a and only three outliers in Figure 13b. Thus, these results suggest that the quality of the experiment data is good, and all participants can be considered skilled HOs.

With regards to using the RMSE as a metric to determine if participants always adapted their control behavior to the new CE dynamics (as hypothesized by **H1**), the RMSE is sorted by participant for the last 30 s of the DYN21 trials in Figure 13c and the last 30 s of the DYN12 trials in Figure 13d. The former can be compared to the DYN1 trials (Figure 13a), while the latter can be compared to the DYN2 trials (Figure 13b). As can be seen, the distribution in RMSE for each participant is very similar between the two pairs, indicating that the control behavior was similar and thus that adaptation took place. One noticeable difference is that the spread in RMSE, quantified by the interquartile range (IQR), is somewhat larger in Figure 13c and Figure 13d than in Figure 13a and Figure 13b, respectively. However, this is to be expected given that the former two are calculated from only one period of the forcing function (30 s) while the latter two are calculated from three periods of the forcing function (90 s). Thus, mistakes, or lack thereof, have a bigger impact on the RMSE when only looking at a 30 s period.


As a final point, sorting the RMSE by forcing function and the order in which the conditions were performed was found not to impact the RMSE, as can be seen in Appendix C.

B. Crossover Frequencies


The crossover frequencies and phase margins sorted by participant, calculated using the first 30 s of the measurement time according to the method described in subsection III.B.1, are presented in Figure 14. The crossover frequencies and phase margins from the six DYN1 trials and six DYN12 trials that did not use the validation forcing functions are shown in Figure 14a and Figure 14c, respectively, while the equivalent results for the DYN2 and DYN21 trials are shown in Figure 14b and Figure 14d, respectively. Therefore, $N = 12$ in all of the box plots, except the right-most in each subfigure where $N = 180$ since it shows all of the participant's data combined. The scattered points overlayed are the results of the three DYN12 trials (or DYN21 trials for Figure 14b and Figure 14d) that used the validation forcing functions. The mean crossover frequency in DYN1 is 2.77 rad/s, while in DYN2 it is 2.33 rad/s. The phase margins are also good, with a mean of 55.6 deg in DYN1 and 41.1 deg in DYN2. In the context of verifying the quality of the experiment data, these results show that all participants had good crossover frequencies and phase margins, and no




(a) Time-invariant DYN1 trials.



(b) Time-invariant DYN2 trials.




(c) Last 30 s of the DYN21 trials, so representative of DYN1.




(d) Last 30 s of the DYN12 trials, so representative of DYN2.


Fig. 13 RMSE by participant in different conditions and phases of the trials.




(a) DYN1 crossover frequencies.



(b) DYN2 crossover frequencies.



(c) DYN1 phase margins.




(d) DYN2 phase margins.

Fig. 14 Crossover frequencies and phase margins by participant calculated using the first 30 s of the measurement time in the six time-invariant trials and the six time-varying trials that did not use the validation forcing functions. The scattered points are the results of the three time-varying trials that used the validation forcing functions.


(group of) participant(s) performed significantly worse than the others. The latter can be seen by the fact that there is only one outlier in the box plot with all of the participant's data in Figure 14a, two in the box plot with all of the participant's data in Figure 14c and Figure 14d, and none in Figure 14b. Thus, like with the RMSE in Subsection IV.A, these results suggest the experiment data is of good quality. Given the relationship between crossover frequency and RMSE, the results from Figure 14 are consistent with the results from Figure 13 in that participants who had low RMSE scores generally had high crossover frequencies. For the time-invariant DYN1 trials, the correlation coefficient ρ between the RMSE and crossover frequency is -0.81, and for the time-invariant DYN2 trials, ρ is -0.52, both of which can be considered high [22].

Figure 15 shows the crossover frequencies and phase margins calculated using the final 30 s of the measurement time in the six time-varying trials that did not use the validation forcing functions. Thus, $N = 6$ in each of the box plots, except the right-most one in each subfigure where $N = 90$ since it shows all of the participant's data combined. The scattered points overlaid are the results of the three time-varying trials that used the validation forcing functions. It should be noted that only the time-varying trials are included in Figure 15 because the goal of this figure is to compare it with Figure 14 to see whether adaptation to the new CE dynamics took place. As can be seen, the means and distributions are similar. Thus it can be concluded, like with the RMSE, that participants adapted their control behavior.


In comparison to similar previous experiments [14][15], the crossover frequencies found here are significantly higher. Possible reasons for this are that the display used in this experiment is different (not only the fact that it is a




(a) DYN1 crossover frequencies.



(b) DYN2 crossover frequencies.



(c) DYN1 phase margins.



(d) DYN2 phase margins.

Fig. 15 Crossover frequencies and phase margins by participant calculated using the last 30 s of the measurement time in the six time-varying trials that did not use the validation forcing functions. The scattered points are the results of the three time-varying trials that used the validation forcing function.




Fig. 16 Detection lags by participant in the time-varying trials.

pursuit display but also the symbols on the display are different), the CE dynamics were modified, and the horizontal axis was used in this experiment as opposed to the vertical (pitch) axis.

Like with the RMSE, the forcing function and order of conditions were found to have minimal impact on the crossover frequencies, as can be seen in Appendix C.

C. Detection Lags and Detection Performance

The detection lags sorted by participant for the DYN12 transition are shown in Figure 16a, and the detection lags for the DYN21 transition are shown in Figure 16b. Again, the right-most box plot shows the data for all participants combined. Participant 2 had a FP in one of the DYN21 trials, thus $N = 5$ for that box plot, and Participant 17 had one FP and one FN in the DYN21 trials so $N = 4$ for that box plot. For the remaining box plots $N = 6$, except the right-most in each subfigure where $N = 90$ and $N = 87$ in Figure 16a and Figure 16b, respectively. The scattered points overlaid are the detection lags for the trials that used the validation forcing functions. Note that only TPs are shown and some (outliers in the) detection lags in Figure 16b are outside of the y-axis range so are indicated by an upwards arrow along with the corresponding detection lag.

Despite the results from Section II suggesting the DYN21 transition would be much harder to detect, it is interesting to note that the detection lags for the two conditions are not vastly different. The mean for the DYN12 transition is 1.82 s compared to 3.03 s for the DYN21 transition. However, the medians are very similar, with a median detection lag of 1.73 s in DYN12 and 1.83 s in DYN21. The large difference between the mean and median in DYN21 is evident from the box plot with all of the participant's data (right-most box plot) in Figure 16b, where it can be seen that the median is much lower than the mean, and it is the third quartile line that is approximately equal to the mean. This indicates that the mean is being dragged up by a few slow detections. It can also be seen in Figure 16 that the spread per participant is generally larger for the DYN21 transition than for the same participant in the DYN12 transition. The larger spread can be seen by the bigger IQRs in Figure 16b compared to Figure 16a. For DYN12, the mean IQR over all participants is 0.93 s, while for DYN21 it is 1.35 s. The larger variability suggests the DYN21 transition was more difficult to detect.

In comparison to similar previous experiments, the detection lags in this experiment were much lower. Van Ham [14] and Plaetinck [23] both performed compensatory tracking experiments investigating the detection of the DYN12 transition where participants were also instructed to press a button when they detected a transition. The median detection lag for Van Ham's experiment [14] was 6.6 s and for Plaetinck's experiment [23] it was 5.2 s (both significantly higher than the median for DYN12 and DYN21 in this experiment). Two factors likely contributed to this. The first is that the break frequency ω_b of the CE dynamics in DYN1 was increased from 6 rad/s to 20 rad/s for this experiment, so the system responded more like pure a single integrator in the crossover region and made the two conditions more distinct. The second is the fact that a pursuit display was used in this experiment, which will be further discussed in Section V.

The detection lags sorted by (non-validation) forcing function can be seen in Figure 17a for DYN12 and in Figure 17b for DYN21. Again, the outliers that exceed the y-axis range in Figure 17b are shown by an upwards arrow along with the corresponding detection lag. As can be seen, there is no clear downward trend in either condition, so the impact of the

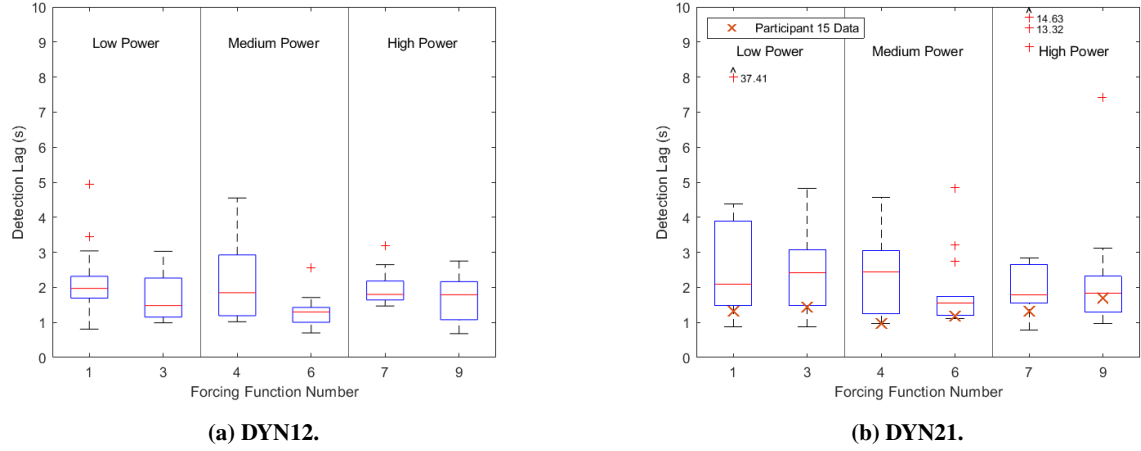


Fig. 17 Detection lags sorted by forcing function.

forcing function was not as large as expected by **H2**. For DYN21, it can be seen that the spread decreases for the latter forcing functions, suggesting the transition is slightly easier to detect when the gradients in the forcing function are large. However, the fact that the lower bound is similar across all forcing functions suggests that the forcing function does not have a very big influence for a well-trained HO. For example, looking at the data for Participant 15 (the scattered crosses in Figure 17b), who arguably had the best detection performance in DYN21 (and good tracking performance in both DYN1 and DYN2), it is clear that increasing gradients in the forcing function did not lead to smaller detection lags.

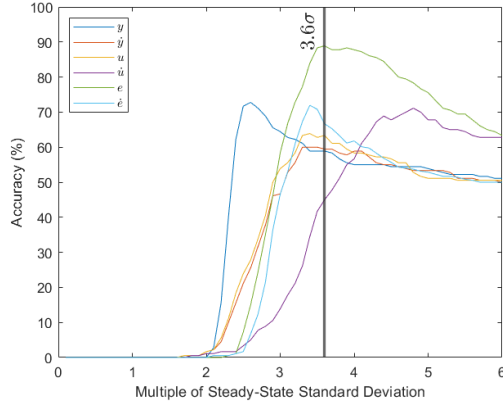
For the DYN12 detection lags, a repeated-measures ANOVA test revealed no significant differences in the means depending on the power in the forcing function in the transition region $F(2, 58) = 1.71$, $p = 0.19$. For the DYN21 detection lags, Mauchly's test indicated that the sphericity assumption had been violated $\chi^2(2) = 26.2$, $p = <0.001$. As a result, the degrees of freedom were corrected using Greenhouse-Geisser estimates of sphericity ($\epsilon = 0.61$). Again, the results revealed no significant difference in the means depending on the power in the forcing function in the transition region $F(1.22, 33.0)$, $p = 0.353$. Note that the two (non-validation) forcing functions in each power level were combined for the ANOVA test, such that the comparison was between low power, medium power, and high power.

Overall, there were a total of twelve FPs, split over six participants, and one FN. Thus, the accuracy over all of the participants, calculated according to Eq. (5), is 97.1%. Of the twelve FPs, seven were in DYN2, two were in each of DYN21 ("too early detections") and DYN1, and one was in DYN12 (also a "too early detection"). The FN was in DYN21. Therefore, the participant accuracy in DYN1 and DYN12 was 98.7%, and in DYN2 and DYN21 it was 95.6%, again suggesting the DYN21 transition was (slightly) more difficult to detect. Six of the twelve FPs occurred during one of the first three trials in the experiment (excluding training trials), and a further three occurred in one of the first three trials of the second half of the experiment, where the set of conditions being tested was different from the first half. This suggests that either there was insufficient training, or the level of confidence participants used to press the button was too low, and after realizing it resulted in too many false button pushes, they adjusted and waited a bit longer before pressing the button. All of this information is summarized in a table in Appendix D.


D. Model for DYN12 Transition Detection

Using the data reserved for developing the model and the methodology given in subsubsection III.B.2, a model based on each of the six candidate signals analyzed can be made and evaluated. In Figure 18a, the accuracy of a model based on the different signals over various thresholds can be seen. As expected based on **H3**, e is the signal that results in the highest accuracy, with a maximum accuracy of 88.9% at a threshold of 3.6σ . For the validation data set, the accuracy at this threshold is 84.4%. The proposed 3.6σ threshold is very similar to Van Ham [14], who proposed a threshold of 3.9σ on the error. In Figure 18b, the breakdown of the number of TPs, FPs, TNs, and FNs over the same range of thresholds can be seen. At the threshold of 3.6σ , there are 80 TPs (out of 90 possible), 13 FPs, 80 TNs, and 7 FNs. Since 50% of the trials considered here (i.e., excluding the trials used to collect validation data) were time-varying, a 100% accuracy would be achieved if there was a TP in 50% of the trials and a TN in 50% of the trials.

In Figure 19, a comparison can be made between the model detections (i.e., when the threshold is exceeded), and



(a) Accuracy of a model based on each of the signals over a range of thresholds.



(b) Breakdown of the number of TPs, FPs, TNs, and FNs for a model based on e .

Fig. 18 Analysis of which signal results in the best accuracy for modeling the DYN12 transition.

participant detections (i.e., when participants pressed the button). The ordering of the trials on the x-axis is arbitrary, as the intention is merely to show a comparison of the model detections with the participant detections for each trial. However, the trials are numbered according to the following (arbitrary) structure. First, the DYN1 trials are grouped (trials 1-90), and the DYN12 trials are grouped (trials 91-180). Within each half, the trials are ordered first by participant and then by forcing function number (in ascending order). Therefore, trials 1-6 are the six DYN1 trials for Participant 2 in order of ascending forcing function number, trials 7-12 are the six DYN1 trials for Participant 3 in order of ascending forcing function number, etc. The same applies to trials 91-180 but with the DYN12 trials instead of DYN1. Therefore, if a detection occurs in one of the areas shaded in red, it is a FP (the darker red is a FP in DYN1 and light red is a "too early detection" in DYN12), while if a detection occurs in the area shaded in green, it is a TP.

The first observation that can be made is that the model has more FPs than the participants, which is the source of the lower model accuracy compared to the accuracy of the participants in the DYN1 and DYN12 trials. Focusing on the TPs, it can also be seen that there are a considerable amount of trials where the model detections are slow. The slowest participant detection is 4.95 s, whereas there are thirteen trials where the model detections are slower than 4.95 s.

Overall, there are 80 trials (out of 90 DYN12 trials) where both the model and the participant have a TP. For these 80 trials, the mean difference between the model's detection lags and the participant's detection lags is 3.46 s. However, it is interesting to consider the cases where the model detections are quicker and the cases where the participant detections are quicker separately. Out of the 80 trials where both the model and the participant have a TP, the model detections were quicker in 42 trials and the participant detections were quicker in the remaining 38 trials. In the 42 trials where the model detections were quicker, the mean difference between the model detections and the participant detections is 0.61 s, which is close to a human reaction time [24]. In the 38 trials where the participant detections were quicker, the mean difference between the model detections and participant detections is 6.62 s. This suggests that in approximately half of the trials (the 42 trials where the model detections were quicker), the model detections are a good estimate of the actual participant detection lags. However, in the other half of the trials (the 38 trials where the participant detections were quicker), the model detections are not representative of the actual participant detection lags. A possible reason for this is that there is a different, or multiple, mechanisms at play in the detection phase. For example, it could be that the participant detections were triggered by a filtered version of e . Another possibility is that the detection was based on the proposed threshold for e in some cases (e.g., the 42 trials where the model detections were quicker), but on a threshold for \dot{e} or the observed relationship between the control inputs and the system's response in other cases.

All in all, it can be concluded that the model proposed based on a threshold for e is capable of achieving a high accuracy in detecting *if* a transition occurred, but not as good at estimating the participants' detection lags.

E. Model for DYN21 Transition Detection

Following the same procedure that was just described for the DYN12 transition, a model can be made for the DYN21 transition. The accuracy for all of the signals over a range of thresholds can be seen in Figure 20a. Note that an accuracy




Fig. 19 Comparison of the moment the threshold is exceeded in the model and when participants pressed the button. FPs appear on the left (DYN1) and bottom half of the right (DYN12 but before the transition) while TPs appear on the top half of the right side.

greater than 0.5 cannot be achieved using the six candidate signals, as hypothesized by **H4**. As a result, it was decided to include the output acceleration \ddot{y} in the analysis as well. Further discussion on the implications of this choice of signal will be given in Section V. However, with a model based on \ddot{y} it is possible to achieve an accuracy as high as 99.4% at a threshold of 4σ . This results in a total of 90 TPs, 1 FP, 89 TNs, and 0 FNs. The accuracy of the validation data set at a threshold of 4σ is 100%.


In Figure 21, the same comparison that was shown in Figure 19 for the DYN12 model is presented for the DYN21 model. Again, the ordering of the trials on the x-axis is arbitrary but is done in the same way as in Figure 19. Contrary to the model for DYN12, there are more participant FPs in this case, with only one model FP.

Out of the 90 DYN21 trials, there are 87 trials in which both the model and the participant had a TP. For these 87 trials, the mean difference between the model's detection lags and the participant's detection lags is 3.12 s. Again, it is interesting to consider the cases where the model detections were quicker and the cases where the participant detections were quicker separately. Here, there is a big difference in the number of cases of each, with the model detections being quicker in 72 of the 87 trials, and the participant detections being quicker in the remaining 15 trials. When the model detections are quicker, the mean difference between the model's detection lags and the participant's detection lags is 2.59 s. Over the same trials, the mean participant detection lag is 3.08 s, which indicates that the model detections are extremely quick in these cases. In fact, in 29 of the 72 trials where the model detections are quicker, the model detection occurs before $t = 45.03$ s, so before the transition is even complete. A discussion on the implications of this will be provided in Section V. In the 15 trials where the participant detections were quicker, the mean difference in detection lags between the model and the participant is 5.65 s. Thus, the model's detection lags are not representative of the actual participant detection lags in neither the case where the model detections were quicker nor the case where the participant detections were quicker.


Overall, it can be concluded that the DYN21 model is able to achieve a near-perfect accuracy in detecting the transitions, but often the model's detection lags are extremely small and are therefore not representative of the actual participant detection lags.

F. Relationship Between $u-\dot{y}$ and $u-\ddot{y}$ in DYN21

To illustrate one of the benefits of using a pursuit display over a compensatory display, the relationship between u and \dot{y} , as well as between u and \ddot{y} , is presented here. This represents the relationship between what participants do (i.e., give control inputs) and what they see on the screen in response when using a pursuit display. The relationship between u and \dot{y} for the ten-second window centered around the transition in every DYN21 trial is shown in Figure 22a. The relationship between u and \ddot{y} for all of the DYN21 trials over the same period is given in Figure 22b. The red crosses are the moment participants pressed the button. Prior to the transition, the DYN2 system is being controlled which means control inputs are proportional to \ddot{y} . Thus, it is logical that the relationship is linear, as in Figure 22b. After the



(a) Accuracy of a model based on each of the signals over a range of thresholds.



(b) Breakdown of the number of TPs, FPs, TNs, and FNs for a model based on \ddot{y} .

Fig. 20 Analysis of which signal results in the best accuracy for modeling the DYN21 transition.




Fig. 21 Comparison of the moment the threshold is exceeded in the model and when participants pressed the button. FPs appear on the left (DYN2) and bottom half of the right (DYN21 but before the transition) while TPs appear on the top half of the right side.




Fig. 22 Relationship between u - \dot{y} and u - \ddot{y} in the ten-second window centered around the DYN21 transition for all participants. The button pushes are indicated by the red crosses.

transition, the control inputs are proportional to \dot{y} , and thus a linear relationship emerges in Figure 22a. As a result of this, all of the button pushes occur in the linear portion of Figure 22a (some of the button pushes towards the bottom left of Figure 22a appear to be outside the linear portion but this is not the case as the linear portion is relatively wide in DYN1). In Figure 22b, the button pushes do not exclusively occur in the linear portion, and when they do, the trajectory taken to get there always goes outside of the linear portion first. Therefore, in addition to a threshold on a particular signal (e.g., \ddot{y} as proposed in Subsection IV.E), deviations from the expected (linear) relationship between u and \dot{y} or u and \ddot{y} can also act as a trigger for participants to detect a transition.

V. Discussion

The goal of this research was to gain a better understanding of what triggers a HO to detect a change in CE dynamics. A human-in-the-loop single-axis pursuit tracking experiment was performed in which the CE dynamics sometimes transitioned from approximate single to approximate double integrator dynamics or vice versa. If and when participants detected the change in CE dynamics, they pressed a button on a joystick to indicate that they detected it. A total of fifteen participants took part, each performing a total of six trials in DYN1, nine trials in DYN12, six trials in DYN2, and nine trials in DYN21. The data from all of the DYN1 and DYN2 trials, as well as six of the trials from each of DYN12 and DYN21 were used to develop a model that predicts when participants would press the button. It was found that the error signal e was best for modeling the DYN12 transition, where an accuracy of 88.9% could be achieved for a threshold of 3.6σ . For DYN21, none of the six candidate signals (y , \dot{y} , u , \dot{u} , e , or \dot{e}) resulted in an accuracy greater than 50%, so instead the model is based on the output acceleration \ddot{y} , which resulted in an accuracy of 99.4% for a threshold of 4σ .

The experiment was successful in collecting good-quality data for developing the models. Throughout the entire experiment, the participant detection accuracy was 97.1%, meaning in the large majority of cases they were able to detect when the CE dynamics changed. The median detection lags in this experiment were also small compared to previous experiments [14][23], with a median of 1.73 s in the DYN12 transition and 1.83 s in the DYN21 transition, again indicating that the quality of the participants was good. The mean crossover frequencies were high as well, with a mean of 2.77 rad/s in DYN1 and 2.33 rad/s in DYN2. A similar experiment was previously conducted [15] where the mean participant crossover frequency was 0.97 rad/s in DYN1 and 1.67 rad/s in DYN2. Such crossover frequencies make it difficult to develop a good, representative model of HO control behavior because they do not really represent skilled HO control behavior. Thus, it is likely that the results presented here are more generalizable to skilled HO behavior, which is ultimately where the applications of this research lie.

The first hypothesis investigated, **H1**, predicted that all participants would adapt their control behavior after a transition to approximate the behavior of the adaptive HO in the simulation (Subsection II.C). In every DYN12 trial of the experiment, participants were able to stabilize the system after the transition and continue tracking the forcing function with the new CE dynamics while keeping the system output symbol on the display until the end of the trial

(even if it momentarily left the display immediately after the transition). On the other hand, it was found that the system with the constant HO in the simulation became unstable after the transition, so this shows that the participants all adapted their control behavior in the DYN12 trials. In DYN21, the system with the constant HO did not become unstable in the simulation, but did suffer from worse tracking performance than the adaptive HO after the transition. It can be seen in Figure 13c and Figure 15 that this was not the case for participants in the experiment, as the distribution of RMSE and crossover frequencies calculated using the last 30 s of the DYN21 trials is very similar to that of steady-state tracking in DYN1. Thus, **H1 is accepted**.

The second hypothesis, **H2**, predicted that larger gradients in the forcing function in the transition region would lead to smaller detection lags. To quantify how large the gradients in the transition region were, the power in the forcing functions in the two-second window centered around the transition was used. Thus, the higher the forcing function number, the larger the gradients in the transition region. From Figure 17 and the repeated-measures ANOVA tests (Subsection IV.C) it can be concluded that larger gradients in the transition region did *not* lead to smaller detection lags, and thus **H2 is rejected**. Despite the fact that Van Ham's experiment [14] had previously found that large gradients in the forcing function *did* generally lead to smaller detection lags, there is one crucial difference with this experiment that likely explains it, namely the type of display. Van Ham [14] used a compensatory display, whereas a pursuit display was used in this experiment. With the pursuit display, participants can see the relationship between their control inputs and the system's response, whereas this relationship is impacted in an unknown way (to the participant) by the forcing function when using a compensatory display.

Knowledge of this relationship is less important when the gradients in the forcing function are large because large control inputs are required to track the large gradients. This causes the magnitude of the majority of the signals in the control loop to increase. Therefore, if exceeding a threshold (e.g., on e for DYN12 and on \ddot{y} for DYN21) is a trigger for the human operator to detect a change in CE dynamics, the large control inputs will make it more likely that this threshold is exceeded.

On the other hand, knowledge of the relationship between the control inputs and the system's response can be more useful when the gradients in the forcing function are small. This is because the magnitude of the relevant signal may not be large enough to exceed the threshold. In this case, participants can use the fact that, in DYN1, u is proportional to \dot{y} , while in DYN2, u is proportional to \ddot{y} . Thus, in DYN1, moving the side-stick to the left or right causes the system output to *immediately* move in that direction. On the other hand, in DYN2, it is necessary to give a control input in the opposite direction to the direction of \dot{y} to *first* slow it down, and *only then* it will start moving in the direction of the control input. Therefore, if participants notice that the relationship between the control inputs and the output has changed, namely that the sign of \dot{y} does not immediately change when the sign of u changes for the DYN12 transition, or that the sign of \dot{y} *does* immediately change when the sign of u changes for the DYN21 transition, that can also be used as a trigger to detect the transition. This relationship is something that cannot be seen by the HO with a compensatory display, so it is hypothesized here that this is why the gradients in the forcing function did not have a big impact on the detection lags.

The final two hypotheses, **H3** and **H4**, have to do with which signal is best to base a model on. For the DYN12 transition, e is the only one of the six candidate signals that results in a high accuracy. **H3** predicted the model could be based on e or \dot{e} so is therefore **rejected**, but the first half of the hypothesis is accepted. This is consistent with previous research that found that a model with a threshold based on e performs well in detecting (specifically) the DYN12 transition [8][9][14]. The accuracy of the model with a threshold of 3.6σ on e is high, though the estimation of the detection lags can still be improved. In approximately half of the DYN12 trials, the model's detection lags were representative of the actual participant detection lags, but in the other half, there was a big difference in detection lags.

For the DYN21 transition, it can be seen in Figure 20a that none of the six candidate signals could be used to make a model with an accuracy greater than 50%. Thus, **H4 is accepted**. Instead, a model based on \ddot{y} is proposed. Despite the fact that humans are not good at visually perceiving accelerations [21], there is still merit to the model because it is not argued that participants need to be able to estimate the value of \ddot{y} . Rather, they only need to be aware of significant changes to its properties, which, post-transition in DYN21, can be perceived as a system that is more aggressive, more responsive to control inputs, and less sluggish, or in the words of some of the participants, less "oversteery". \ddot{y} also captures the relationship between u and \dot{y} , and between u and \ddot{y} described just above. This is because, in order for \dot{y} to be proportional to u (as is the case in DYN1), it is necessary for \ddot{y} to be larger in DYN1 than DYN2 so that the sign of \dot{y} changes quicker after a change in the sign of u , and this can be used by participants to detect the transition.

A limitation of this model, however, is that the model detections (i.e., the moment the threshold is exceeded) are extremely quick, often occurring before $t = 45.03$ s for a transition at $t = 45$ s. This suggests the model could be fit to a property of the CE dynamics rather than something participants use to detect the transition, and thus may not be generalizable to other CE dynamics. On the other hand, it can be argued that this information *is* what participants use to

detect the transition, but the fact that it is an acceleration and requires more precise knowledge of the expected response means it takes longer for the participant to notice.

The next step towards using this research in real-life applications is adjusting the models to more closely estimate the real HO detection lags. One complication in estimating the detection lags, however, is that it is not a binary decision to press the button, such as if someone was instructed to press the button when a picture pops up on the screen. Thus, there is a statistical element to the data [14]. Furthermore, the button push data are subjective, and different people have different thresholds for when they think they should press the button, in addition to cognitive decision time [24]. For this reason, it is recommended to introduce lags into the detection process. This can be done, for example, by passing the relevant signal through a low pass filter as done by Hess [13] or adding an explicit delay to model reaction times.

Another recommendation is to include u in the model. For this research, an attempt was made to include u in the threshold though it was unsuccessful in improving the estimation of the detection lags. Nonetheless, the results presented in this paper suggest that the relationship between u and the system's response is important, so additional investigation into how to include this in the model is recommended. One possibility is to develop a model that generates expectations of the behavior of u given the current y , \dot{y} , and f_i . Significant mismatch between the expectation and reality can then also be used as a trigger, or can be incorporated into the threshold (e.g., the higher the mismatch, the lower the threshold on \ddot{y} for the DYN21 model). This is similar to Phatak and Bekey [9] and Van Ham [14] who proposed independent thresholds on e and \dot{e} but with a (possible) difference. The recommendation here is to try and include u in a *single* model such that it uses (at least) u and \ddot{y} to determine a threshold, as opposed to a model that can trigger based on individually looking at different signals. To further understand the role that the relationship between the control inputs and system response plays, a similar experiment to the one performed for this research could be conducted but with several disturbance signals, each with differing amounts of power. If the relationship between the control inputs and the system's response is indeed critical, it should be found that increasing the power in the disturbance signal leads to larger detection lags because the relationship is lost.

Overall, however, this research provides the first steps to understanding how HOs detect a change in CE dynamics, particularly the DYN21 transition, using pursuit displays. A model for the DYN21 transition based on \ddot{y} is proposed and recommendations for future improvements to the model as well as future investigations are made.

VI. Conclusion

This paper aimed to gain an understanding of what triggers a HO to detect a change in CE dynamics and see if this can be predicted with a model. A human-in-the-loop pursuit tracking task experiment was performed to collect data to develop and validate the models. Fifteen participants took part and had a combined accuracy of 97.1% in detecting transitions. The mean detection lag was 1.82 s for the DYN12 transition and 3.03 s for the DYN21 transition. The observed crossover frequencies were high, with a mean of 2.77 rad/s for DYN1 and 2.33 rad/s for DYN2, indicating that skilled participants were selected. Two models were developed, one for DYN12 and one for DYN21. The model for DYN12 is based on e and has an accuracy of 88.9% at a threshold of 3.6σ . The model for DYN21 is based on \ddot{y} and has an accuracy of 99.4% at a threshold of 4σ . The estimation of detection lags in both models remains one of their limitations, particularly the model for DYN21 where the model's detection lags are often extremely small, and, in the majority of cases, well before the participants pushed the button. Overall, however, this research helped confirm previous findings that a high-accuracy model can be developed based on e for the DYN12 transition, proposed a model with a high accuracy for the DYN21 transition, and identified that the relationship between the control inputs and system response is likely key in the detection process, particularly for the DYN21 transition when using a pursuit display.

References

- [1] Mulder, M., Pool, D. M., Abbink, D. A., Boer, E. R., Zaal, P. M., Drop, F. M., Van Der El, K., and Van Paassen, M. M., "Manual Control Cybernetics: State-of-the-Art and Current Trends," , 10 2018. <https://doi.org/10.1109/THMS.2017.2761342>.
- [2] "Loss of Control In-flight (LOC-I)," , 2023. URL <https://www.iata.org/en/programs/safety/operational-safety/loss-of-control-inflight/#:~:text=LOC%2DI%20refers%20to%20accidents,failures%2C%20icing%2C%20or%20stalls>.
- [3] Drop, F., "Control-Theoretic Models of Feedforward in Manual Control," 2016. <https://doi.org/10.4233/uuid:7c1f62db-9a5a-4e02-8f11-488d6a299500>, URL <https://doi.org/10.4233/uuid:7c1f62db-9a5a-4e02-8f11-488d6a299500>.
- [4] Young, L. R., Green, D. M., Elkind, J. I., and Kelly, J. A., "Adaptive Dynamic Response Characteristics of the Human

- Operator in Simple Manual Control,” *IEEE Transactions on Human Factors in Electronics*, Vol. HFE-5, No. 1, 1964, pp. 6–13. <https://doi.org/10.1109/THFE.1964.231648>.
- [5] Young, L. R., and Stark, L., “Biological control system - a critical review and evaluation, developments in manual control,” Tech. rep., In collab. with NASA, 3 1965. URL https://archive.org/details/nasa_techdoc_19650009660/page/n107/mode/2up.
 - [6] Young, L. R., “On Adaptive Manual Control,” *Ergonomics*, Vol. 12, No. 4, 1969, pp. 635–674. <https://doi.org/10.1080/00140136908931083>.
 - [7] Niemela, R. J., and Krendel, E. S., “Detection of a Change in Plant Dynamics in a Man-Machine System,” *Correspondence*, 1975.
 - [8] Elkind, J. I., and Miller, D. C., “Process of Adaptation by the Human Controller,” *Second Annual NASA University Conference on Manual Control*, 1966, pp. 47–63.
 - [9] Phatak, A. V., and Bekey, G. A., “Model of the Adaptive Behavior of the Human Operator in Response to a Sudden Change in the Control Situation,” *IEEE Transactions on Man-Machine Systems*, Vol. 10, No. 3, 1969, pp. 72–80. <https://doi.org/10.1109/TMMS.1969.299886>.
 - [10] Bastian, A. J., “Moving, sensing and learning with cerebellar damage,” *Current Opinion in Neurobiology*, Vol. 21, No. 4, 2011, pp. 596–601. <https://doi.org/10.1016/j.conb.2011.06.007>.
 - [11] Imamizu, H., Miyauchi, S., Tamada, T., Sasaki, Y., Takino, R., Pütz, B., Yoshioka, T., and Kawato, M., “Human cerebellar activity reflecting an acquired internal model of new tool,” *Nature*, Vol. 403, 2000, pp. 192–195.
 - [12] Hess, R. A., “Modeling pilot control behavior with sudden changes in vehicle dynamics,” *Journal of Aircraft*, Vol. 46, No. 5, 2009, pp. 1584–1592. <https://doi.org/10.2514/1.41215>.
 - [13] Hess, R. A., “Modeling human pilot adaptation to flight control anomalies and changing task demands,” *Journal of Guidance, Control, and Dynamics*, Vol. 39, No. 3, 2016, pp. 655–666. <https://doi.org/10.2514/1.G001303>.
 - [14] van Ham, J. M., Pool, D. M., and Mulder, M., “Predicting Human Control Adaptation from Statistical Variations in Tracking Error and Error Rate,” *IFAC-PapersOnLine*, Vol. 55, No. 29, 2022, pp. 166–171. <https://doi.org/10.1016/j.ifacol.2022.10.250>.
 - [15] Terenzi, L., Zaai, P. M., Pool, D. M., and Mulder, M., “Adaptive Manual Control: a Predictive Coding Approach,” *AIAA Science and Technology Forum and Exposition, AIAA SciTech Forum 2022*, American Institute of Aeronautics and Astronautics Inc, AIAA, 2022. <https://doi.org/10.2514/6.2022-2448>.
 - [16] Jakimovska, N., Pool, D. M., van Paassen, M. M., and Mulder, M., “Using the Hess Adaptive Pilot Model for Modeling Human Operator’s Control Adaptations in Pursuit Tracking,” *American Institute of Aeronautics and Astronautics (AIAA)*, 2023. <https://doi.org/10.2514/6.2023-0541>.
 - [17] Mulder, M., Pool, D. M., van der El, K., Drop, F. M., and Van Paassen, M. M., “Manual Control with Pursuit Displays: New Insights, New Models, New Issues,” *IFAC-PapersOnLine*, Vol. 52, No. 19, 2019, pp. 139–144. <https://doi.org/10.1016/j.ifacol.2019.12.125>.
 - [18] McRuer, D. T., and Jex, H. R., “A Review of Quasi-Linear Pilot Models,” *IEEE Transactions on Human Factors in Electronics*, Vol. HFE-3, No. 3, 1967, pp. 231–249. <https://doi.org/10.1109/THFE.1967.234304>.
 - [19] Van Grootheest, H. A., Pool, D. M., Van Paassen, M. M., and Mulder, M., “Identification of time-varying manual-control adaptations with recursive ARX models,” *AIAA Modeling and Simulation Technologies Conference*, American Institute of Aeronautics and Astronautics Inc, AIAA, 2018. <https://doi.org/10.2514/6.2018-0118>.
 - [20] Van Ham, J. M., *Adaptive Manual Control - The Human Response to Sudden Changes in Controlled Element Dynamics (MSc thesis)*, 2021. URL <https://surfdrive.surf.nl/files/index.php/s/p3sY6Ur0IFjsjIX>.
 - [21] Brenner, E., Rodriguez, I. A., Muñoz, V. E., Schootemeijer, S., Mahieu, Y., Veerkamp, K., Zandbergen, M., Van der Zee, T., and Smeets, J. B., “How can people be so good at intercepting accelerating objects if they are so poor at visually judging acceleration?” *i-Perception*, Vol. 7, No. 1, 2016. <https://doi.org/10.1177/2041669515624317>.
 - [22] Gignac, G. E., and Szodorai, E. T., “Effect size guidelines for individual differences researchers,” *Personality and Individual Differences*, Vol. 102, 2016, pp. 74–78. <https://doi.org/10.1016/j.paid.2016.06.069>.
 - [23] Plaetinck, W., Pool, D. M., van Paassen, M. M., and Mulder, M., “Online Identification of Pilot Adaptation to Sudden Degradations in Vehicle Stability,” *IFAC-PapersOnLine*, Vol. 51, No. 34, 2019, pp. 347–352. <https://doi.org/10.1016/j.ifacol.2019.01.020>.
 - [24] Thorpe, S., Fize, D., and Marlot, C., “Speed of processing in the human visual system,” *Nature*, Vol. 381, 1996, pp. 520–522.

N.B.: The following part has been assessed
for the course AE4020 Literature Study.

Part II

Preliminary Thesis Report

1

Introduction

Over the past decades, the role of automation in society has increased significantly. However, in the context of control tasks, this does not mean the role of human operators has become any less important. On the contrary, it can be argued that the role of human operators has become even more important as it has shifted to a more supervisory role, in which interventions made by the human operator are in situations where caution, precision, and adaptability may be required. The adaptability in particular is currently one of the main strengths human operators have over automation. However, with the advancement of technology and development in how human operators interact with systems, the models used to describe their behaviour have failed to keep up [1].

It is beneficial to understand the adaptive behaviour of human operators for several reasons. First, it would help improve the design of current systems, since different design choices can be tested through simulation to identify potential problems. Additionally, adaptive models can be used to create better support systems for human operators. In the case of failures, for example, this can help reduce the number of loss of control incidents, which are still one of the largest worldwide contributors to fatal accidents [2]. Finally, it would allow for the development of safer and better vehicle designs [3]. For example, the introduction of human-like automation in self-driving cars [1].

Young et al [4] identified that the adaptive process of a human operator consists of three phases, namely detection, identification, and modification. In the detection phase, the human operator realizes that something has changed and that adaptation is required. This triggers the identification phase to begin, in which the human operator identifies the nature of the change. Finally, in the modification phase, the human operator modifies their control behaviour to suit the new situation. It can be argued that the identification and modification phases happen simultaneously, though the detection phase is certainly distinct. A good adaptive model would be capable of approximating the human operator's behaviour throughout all three phases. This thesis, however, will look only at the detection phase, specifically the case where there is a change in the controlled element dynamics. Future research can later extend the results from this research to include the identification and modification phases.

The goal of this research is thus to develop a model that is able to predict when a human operator will detect a change in the controlled element dynamics, and consequently help improve understanding of what causes the detection. The main research question that will be answered is,

How can a human operator's detection of a change in controlled element dynamics be modeled?

To help answer the main research question, the following sub-questions have been formulated.

1. What are the signals that human operators use to detect that the controlled element dynamics have changed, and why are they relevant?
2. What are the threshold values for the relevant signals?
3. How does the nature of the original dynamics affect the process of detecting a change?

4. How can the experiment best be designed to help develop and validate the model?

To get a preliminary answer for the first and third sub-questions, research will be done on concepts and theories relevant to human manual control, and on the strengths and weaknesses of existing models. The former will be presented in Chapter 2 while the latter will be presented in Chapter 3. In Chapter 4, results from a simulation will be presented to help answer the second sub-question, and finally, in Chapter 5, future work (including the experiment design) will be presented to help answer the fourth sub-question.

2

Relevant Concepts for Manual Control

Before starting with the literature review, some concepts and theories that are relevant to human manual control will be introduced. Their relevance comes from the fact that they are an underlying assumption in several models to be introduced in Chapter 3 and they help explain observations that have been made. First, the concept of internal models will be explained in Section 2.1, followed by an introduction to adaptive model theory in Section 2.2, and finally a brief summary of McRuer's crossover model in Section 2.3.

2.1. Internal Models

Internal models are believed to be internal representations stored in the cerebellum that contain information about our own movement dynamics, and the dynamics of the systems we interact with around us. Their role is to generate predictions about future sensory information, in turn minimizing surprise [5]. Significant evidence, some of which will be presented here, has been found to support the existence of internal models, and their existence is commonly accepted [6] not only in humans but also in other animals [7] [8]. With respect to control tasks, two types of internal models are hypothesized to exist, namely forward models and inverse models [9]. A diagram showing how inverse and forward models work is given in Figure 2.1. Inverse models invert the controlled element (CE) dynamics and estimate the control input required to achieve a certain (desired) output. On the other hand, forward models use the control input to estimate the output given the current states. As such, they can be used to perform mental simulations of what would happen if various control inputs were to be given.

With regards to internal models encoding information about our own movement dynamics, the first pieces of evidence presented here are point-to-point (arm) reaching experiments performed with external force fields. The force fields were generated with either a robot manipulandum [10] or a spinning room [11] [12], and they change the effective arm dynamics. In participants' first few trials, their arm movements were distorted (instead of straight as is usually the case) and the end point errors were large. In subsequent trials, the movements started becoming straighter and the end point errors de-

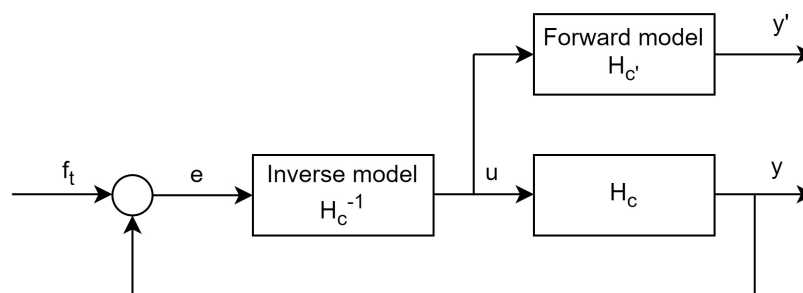


Figure 2.1: Diagram of inverse and forward models.

creased. When the external force field was removed, the opposite effect to what was seen in the first few trials was observed, namely that the arm movements became distorted but in the opposite direction to the external force fields. This strongly suggests the existence of an adaptive inverse model that quickly adapted to compensate for the external force field. Further evidence includes that patients with cerebellar damage display inferior performance in motor tasks and increased difficulty in learning new tasks compared to a person with a healthy cerebellum [13]. Both of these effects can be explained by the lack of internal models. The delay in processing sensory feedback, such as visual inputs, is large, making it inadequate for rapid movements. People with healthy cerebellums are able to overcome this time delay through the use of internal models, but since that is not possible for people with cerebellar damage, they have poorer performance in motor tasks and increased jittering in their movements [14]. The greater difficulty in learning new tasks is due to the belief that internal models play a more vital role during learning, rather than in retaining learned tasks [13].

Evidence has also been found for the development of internal models of systems we interact with (i.e. the CE). Davidson et al [15] performed an experiment to detect the presence of an adaptive inverse internal model of the CE dynamics. Participants had to track a target signal composed of a sum of sine waves, but at certain points during the run, the system output was blanked and they could only see the target signal. The intention was to see whether the temporal improvements in blanked and unblanked runs (compared to the first trial) were equal, as this would show whether an inverse model is present at all, as well as how much it contributed to learning compared to a possible forward model. If the improvement in blanked runs was equal to the improvement in unblanked runs, that would suggest the existence of only an inverse model since only the inverse model can be used during blanked runs as a result of there not being feedback of the system output. During unblanked runs, on the other hand, both internal models could be tuned to improve performance. The results showed an improvement of 33% in unblanked trials, but only 18% in blanked trials. Since the improvement in the blanked trials can be attributed to an improvement in an adaptive inverse internal model, the remaining improvement is likely caused by an improvement in an adaptive forward model. Thus, this experiment provided evidence for the existence of both an inverse model and a forward model. Further evidence was presented by Imamizu et al [16] as they found increased cerebellar activity during and after visuomotor learning. As learning plateaued, cerebellar activity decreased once again, suggesting the internal model was already developed.

Rather than having an internal model for each activity, it is hypothesized that we store multiple internal models that we can quickly switch between [17]. This is supported by Imamizu et al's [16] observation that several regions of the cerebellum showed increased activity during visuomotor learning, which suggests different internal models use the new information for tuning. Further evidence includes the fact that tasks requiring similar skills, such as playing tennis and ping-pong, are easier to learn when one already has experience (and therefore a developed internal model) in one of the tasks [17]. Finally, learning a new condition in a certain task takes considerably longer than re-adapting to a previously learned condition, likely due to the fact that the internal model for the previously learned condition is already in memory. For tracking tasks, this means it is likely that (for a well-trained participant) an internal model exists for the different CE dynamics, so if and when the human operator (HO) detects that the CE dynamics have changed, it is 'simply' a matter of switching from one internal model to the other.

2.2. Adaptive Model Theory

Adaptive model theory (AMT) is a theory developed by Neilson along with a team of researchers which they describe as, "a computational theory of the brain processes that control purposive coordinated human movement." [18] It is applicable to pursuit tracking tasks and outlines a biologically realistic way in which information is processed by the brain and ultimately leads to (control) actions taken. Based on the theory, Neilson et al [19] made a simulation to demonstrate the capabilities of AMT. The resulting simulation is adaptive, changing its behaviour in response to changes in the statistical properties of the forcing function or disturbance signal, or changes in the CE dynamics which, as outlined by Mulder et al [20], is precisely what is required in future cybernetics models.




Figure 2.2: Block diagram for the compensatory tracking system.

AMT is made up of three stages, namely sensory analysis (SA), response planning (RP), and response execution (RE) [19]. The first stage is the SA stage, which operates continuously, and has three inputs as well as three outputs. The inputs are the forcing function, disturbance signal, and system output, while the outputs are estimates for the next value of the forcing function, disturbance signal, and system output. For predicting the system output, it uses the current system output along with the planned system output trajectory (to be discussed in the RP stage) that has not yet been executed. The second stage is the RP stage, which introduces intermittency as it has a fixed duration of 150 ms. The inputs to this stage are the outputs of the SA stage, and it outputs the planned system output trajectory. The planned system output trajectory is a trajectory taking the current system output to the predicted (future) forcing function value. Finally, the RE stage, which also operates continuously, takes the planned system output trajectory as input and determines the required control input via an inverse internal model. The inverse internal model is a series of three inverse models, one for the muscle control system, one for the biomechanical system, and one for the external system (i.e. the CE), so it goes down to very low-level motor commands.

Particularly with the first two of the inverse models, there are significant technical challenges in explaining how the nervous system is able to develop those models [18]. Amongst the challenges include the fact that the system is high dimensional, multiple-input multiple-output (MIMO), redundant, time-varying, and nonlinear. Formulating good models for MIMO systems is not possible when the inputs are highly correlated. Neither is it possible to formulate good inverse models when there is a lot of noise or when there are more inputs than outputs. Somehow, however, the nervous system is able to remove redundancy and still get high-accuracy inverse models. This is where the power of AMT lies, the ability to propose a biologically realistic way in which this is done. Despite insufficient neurobiological data to verify it, the ability of the simulation to accurately capture measured HO adaptive behaviour during a change in CE dynamics is sufficient for AMT to, at a minimum, be used to suggest and check theories on HO adaptive behaviour.

2.3. McRuer's Crossover Model

The previous two sections focused on concepts and theories more applicable to time-varying conditions. In this section, a brief summary of McRuer's crossover model [21], applicable to compensatory tracking in steady-state conditions, will be provided. Developed in the early 1960s, it remains one of the leading universal models for modeling HO behaviour in compensatory tracking due to its accuracy and simplicity. Unfortunately, compensatory displays are the exception in day-to-day applications, but since several of the models to be discussed in Chapter 3, as well as the simulation developed in Chapter 4, use the crossover model, a brief summary will be presented here.

The block diagram for the system is given in Figure 2.2. $Y_p(j\omega)$ is the pilot/HO dynamics, $Y_c(j\omega)$ the CE dynamics, f_t the forcing function, e the error, u the control input, and y the system output. McRuer's primary finding was that the open loop dynamics, $Y_{OL}(j\omega)$, approximate an integrator with a time delay in the crossover region for a range of CE dynamics, as given by Equation (2.1).

$$Y_{OL}(j\omega) = Y_p(j\omega)Y_c(j\omega) = \frac{\omega_c}{j\omega} e^{-j\omega\tau_e} \quad (2.1)$$

The HO is able to achieve this by altering the structure of their own dynamics depending on the CE dynamics. The simplified precision model, Equation (2.2), is the simplest way in which the HO dynam-

ics can be represented. k_p is the gain, T_L the lead time constant, T_I the lag time constant, and τ_e the effective time delay, which accounts for all sources of delay. The gain is used to place the crossover frequency where desired, while the lead and lag time constants are used to equalize the dynamics and ensure the open loop dynamics are integrator-like in the crossover region. It should be noted that one or both of the time constants may not be required for certain CE dynamics. For example, if the CE dynamics are already integrator-like, neither is required, and if the CE dynamics are double integrator-like, only the lead time constant is required. The values for all of the parameters in Equation (2.2) can be found using the verbal adjustment rules [21]. It should be noted that McRuer's model is a quasi-linear model, and Equation (2.2) only accounts for the linear portion of the HO. The nonlinear part, called the remnant, is usually added to the output of the HO (u in Figure 2.2) but its structure is not specified by McRuer.

$$Y_p(j\omega) = k_p \frac{1 + T_L j\omega}{1 + T_I j\omega} e^{-j\omega\tau_e} \quad (2.2)$$

It is interesting to note that when considering the crossover model in terms of an internal model, Equation (2.1) can be rewritten to Equation (2.3) and the inverse internal model of the CE dynamics can be explicitly seen as the term $\frac{1}{Y_c(j\omega)}$. In this formulation, the integrator can then be used to reduce any steady-state errors. As an example of how this ties back to the simplified precision model, if the CE dynamics are second order, one of the poles cancels with the integrator thus leaving a lead term multiplied with a gain (ω_c), which is exactly what is dictated by the simplified precision model in Equation (2.2).

$$Y_p(j\omega) = \frac{1}{Y_c(j\omega)} \cdot \frac{\omega_c}{j\omega} e^{-j\omega\tau_e} \quad (2.3)$$

3

Review of Existing Models

Significant research efforts have been dedicated to the development and investigation of adaptive models. Different types of such models exist, including some that focus exclusively on the detection phase, some that focus exclusively on modeling the adaptive response without investigating what triggered the HO to detect the transition and start adapting, and some that are a mixture of the two. In this chapter, the results from a literature review of these models will be presented. This will help answer the first sub-research question ("What are the signals that human operators use to detect that the controlled element dynamics have changed, and why are they relevant?") and the third sub-research question ("How does the nature of the original dynamics affect the process of detecting a change?"). Models that investigate the detection phase will be presented in Section 3.1, while models that consider the adaptation phase but do not investigate the detection phase will be presented in Section 3.2. Finally, conclusions drawn from the literature review will be presented in Section 3.3.

3.1. Models Investigating the Detection Phase

This section will discuss models that investigate the detection phase, including some that propose explicit boundaries for certain signals which, when exceeded, are hypothesized to trigger the HO to detect a transition.

3.1.1. Young and Stark Model

Young and Stark [22] developed a model reference approach for compensatory tracking in 1965 to describe the potential process through which HOs detect a change in the CE dynamics and consequently adapt their control behaviour. They distinguished between two types of models, namely model reference models, such as the one they developed, and error pattern recognition models, which will be discussed in the coming subsections. The block diagram of their model is given in Figure 3.1. The dotted lines represent hypothetical connections in the adaptive process, while the solid lines are a conventional closed-loop control diagram. The principle on which the model is based is the idea that HOs keep an internal model of the CE dynamics, given by the "model of controlled element" block. With knowledge of the current error and the current control input, a mental simulation occurs in which the future error **rate** is predicted. If there is a significant mismatch, then adaptation begins. The deviation filter determines what a "significant" mismatch is and consequently filters out insignificant errors while passing on important ones to the "adaptive control operator" block to adjust the current control law until the deviations return to a normal level. However, the model has not been developed further and is thus only a qualitative model with no details on how to actually implement the internal model, deviation filter, or adaptation logic.

Building on the model reference approach, Miller and Elkind [23] developed a model, also for compensatory tracking, for the detection phase based on the error rate and, crucially, added the side-stick movement (i.e. control input) as part of the detection criteria. Miller and Elkind [23] suggest that a HO will detect a change to the CE dynamics when the difference between the predicted error rate (based on the current error rate and control input) and the observed error rate, $\Delta\epsilon$, exceeds a certain thresh-

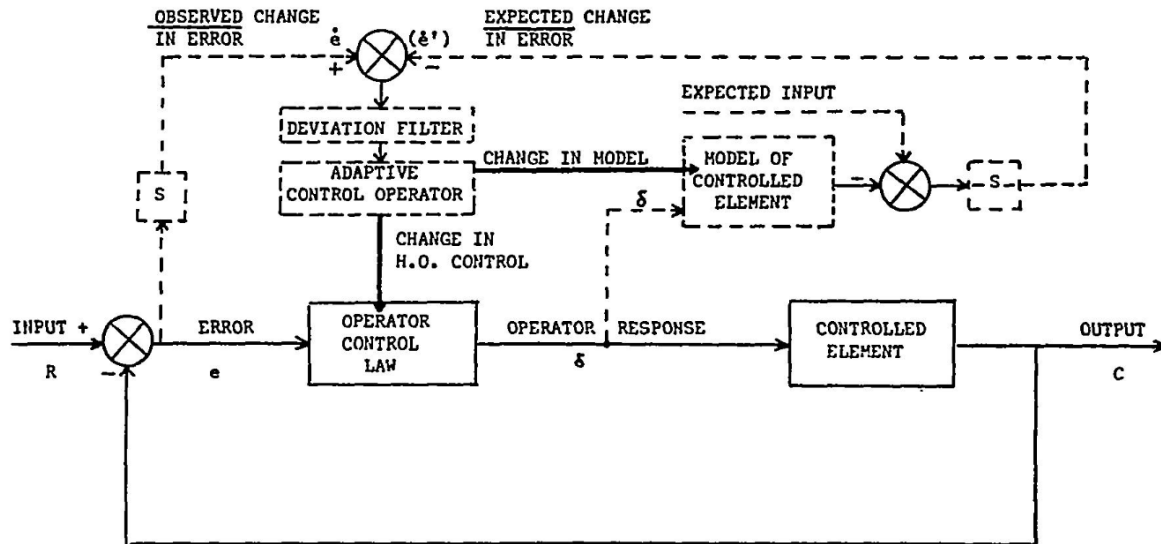


Figure 3.1: Block diagram for Young and Stark's model [22].

old. The threshold is dependent not only on the standard deviation of $\Delta\epsilon$ but also on the change in the control input, ΔC . While the exact way in which the control input is incorporated into the threshold is not specified, some results showing the success of adding it to the detection criteria are. In Figure 3.2, the strong correlation between ΔC and $\Delta\epsilon$ at the moment of detection can be seen. This suggests that when control inputs are small, a small mismatch in predicted and observed error rate can be enough to trigger the HO to detect a transition, but when the control inputs are larger, the mismatch also needs to be larger.

The success of including the control inputs in the detection criteria can be seen in Figure 3.3. A strong relationship between when participants detected the transition (button release time) and the time at which $\Delta\epsilon$ exceeds the threshold (dependent on the control inputs) can be seen. This suggests that knowledge of the control inputs is important to be able to detect a change in CE dynamics.

To further investigate this, Young [24] performed another compensatory tracking experiment. A schematic of the experiment setup is given in Figure 3.4. Participants were divided into three groups, namely active controllers, inactive controllers, and passive monitors. The active controllers were in control of the system the entire run, while the inactive controllers were initially in control, but after some time their control inputs were disconnected without their knowledge and they saw the same error as was being shown to the active controller. The passive monitors simply watched the active controller during a run. Thus, the active controllers had knowledge of the system (so were able to generate an internal model) and of the control inputs, while the inactive controllers had knowledge of the system but incorrect knowledge of the control inputs, and the passive monitors had neither. During the run, all three groups had to indicate via a button press when they detected a change in the CE dynamics. The time taken to detect the change for transitions from single to double integrator and vice versa is given in Table 3.1.

Table 3.1: Time taken (in seconds) for the three groups of participants in Young's experiment to detect that the CE dynamics had changed for a transition from single to double integrator and vice versa [24].

Participant group	Single to Double Integrator Transition (secs)	Double to Single Integrator Transition (secs)
Active controller	1.36	1.23
Inactive controller	1.56	1.45
Passive observer	5	2.48

As can be seen in Table 3.1, the active controllers took the least amount of time to detect a change.




Figure 3.2: Experimental results from Miller and Elkind [23] showing the strong correlation between ΔC and $\Delta \epsilon$ at the moment participants detected the change in CE dynamics.




Figure 3.3: Relationship between when participants detected the change in CE dynamics and $\Delta \epsilon$ at that moment [23].

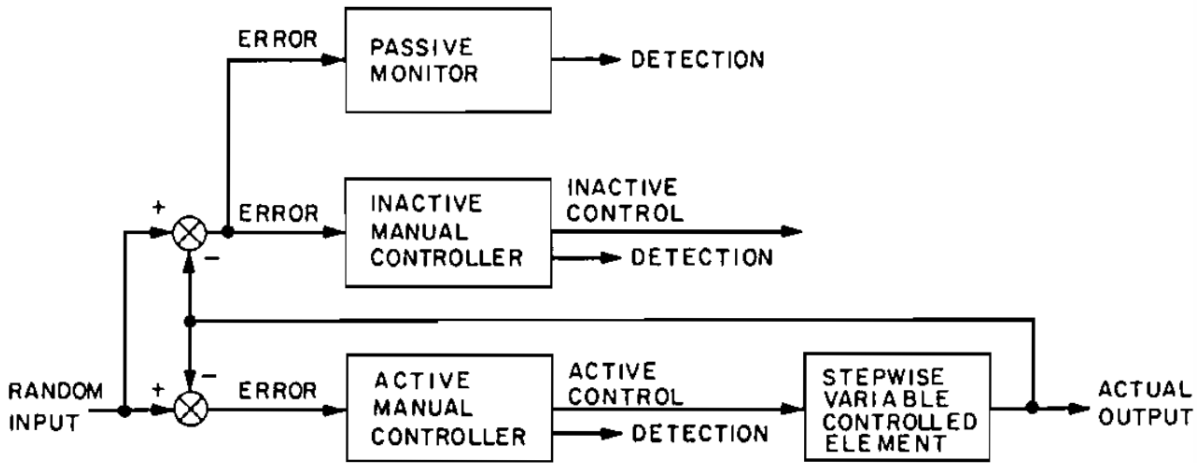


Figure 3.4: Setup of Young's experiment [24].

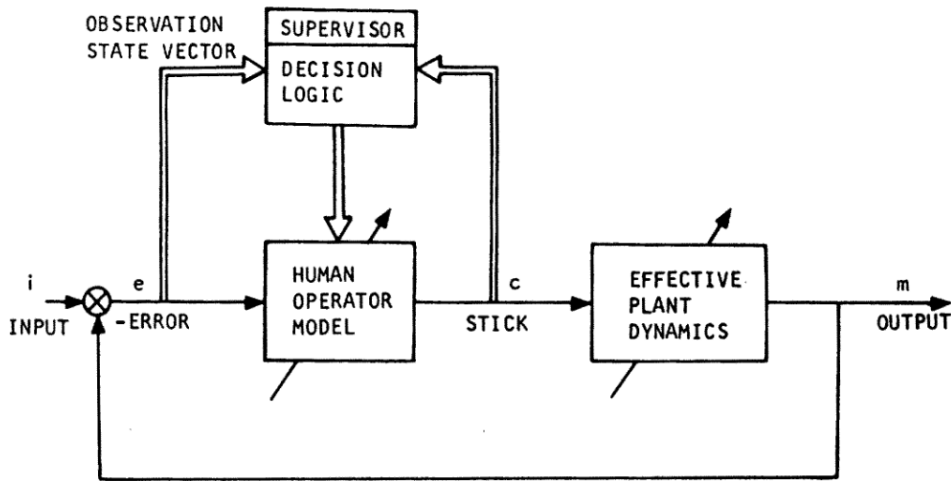


Figure 3.5: Block diagram for the compensatory tracking task of Phatak and Bekey [25].

The inactive controllers had a similar, though slower time and the passive observers took significantly longer. This suggests the existence of an internal model for participants that had experience controlling the system, as well as it being an important factor in determining the transition moment. The quicker detection times for the active controllers compared to the inactive controllers also suggests that knowledge of the stick inputs does indeed contribute to quicker detection times, though to a lesser extent than the internal model of the system dynamics.

3.1.2. Phatak and Bekey Model

Phatak and Bekey [25] developed an adaptive model for compensatory control tasks featuring a multistage discrete decision making process called the supervisory control algorithm. The block diagram is given in Figure 3.5. The HO dynamics are dependent on the CE dynamics and are calculated according to McRuer's crossover model. For the CE dynamics, an approximate second-order system is chosen, but in the nominal condition, it has both rate and attitude feedback to improve the stability and decrease the order of the effective dynamics. When modeling failure, one or both of the feedback loops fail to varying degrees. The supervisor block represents the internal model developed by the HO that keeps track of the (maximum) error and error rate throughout the run.

It is hypothesized by Phatak and Bekey that the magnitude of the maximum value of the error or error rate during nominal conditions is the critical piece of information that is stored by the HO which triggers

them to detect a change in the CE dynamics. If the error or error rate exceeds twice the maximum of the respective signal during the run, designated as decision region 1 (DR-1), the adaptation process begins. For this, the supervisory control algorithm is introduced, as presented in Figure 3.6. To interpret Figure 3.6 it should be noted that this algorithm relies on there being a finite set of post-transition CE dynamics that the (experienced) HO can choose from once a failure is detected, which are called augmentation level A-D. Furthermore, all of the possible post-transition dynamics were designed to be unstable if the HO did not adapt their control strategy, and thus, the error or error rate would always exceed the bounds of DR-1 and trigger the start of adaptation.

Referring to Figure 3.6, adaptation begins if and when the error or error rate exceeds the DR-1 boundary. If this is the case, then the HO adopts the control strategy required to stabilize one of the other three possible CE dynamics (augmentation level B in this case). If that succeeds in changing the sign of the error rate and the error or error rate do not grow larger to the point of exceeding the bounds of the next decision region (DR-2, which has higher bounds), then this is the control strategy adopted by the HO. If the error or error rate do exceed the bounds of DR-2, then the process is repeated with the HO adopting the control strategy required to stabilize one of the other two remaining possible CE dynamics (augmentation level C in this case), and the process is repeated. As such, it is clear that the HO needs to be familiar with all of the possible failure cases as well as which control strategy is required for each.

The decision regions (DR-1, DR-2, and DR-3) were determined based on data from a roll tracking experiment that had one well-trained participant. Later, a simulation was made to check the performance of the supervisory control algorithm developed. Different failure moments and inputs were tested in the simulation, where it was found that the characteristics of the adaptive response, including retention time (the length of time between when the failure occurs and when it is detected by the HO), maximum error and error rate, and settling time, were similar to those found with the human-in-the-loop experiment. Thus, it was concluded that the proposed model was feasible for modeling adaptive behaviour. As a final check, the same input and failure times as in the experiment were used in the simulation, and it was found that the model output closely matched the HO data for the first three seconds after failure. Afterwards, the model predicted a better performance of the HO than what was observed in the experiment which the authors attributed to the HO potentially reacting to disturbances or an incorrect estimation of the error or error rate.

More recently, van Ham [26] performed a compensatory pitch tracking experiment with six participants in an attempt to experimentally validate the model and test the robustness when exposed to different CE dynamics. In the experiment, the CE dynamics transitioned from approximate single to approximate double integrator, and the participants could indicate they noticed a change in the CE dynamics by pressing a button on the side-stick. To estimate which signal (the error or error rate) triggered the detection, and to know what the value of that signal was at the time, van Ham looked for a local maximum in one of the signals 0.3 seconds to 0.6 seconds before the button was pressed. This time window was taken as it was found in previous research that it is the typical amount of time required for a HO to press a button in response to a visual input [27]. Using this method, van Ham found that in 92.7% of cases, the boundaries of DR-1 were not exceeded so they would be classified as "no detection". In the remaining 7.3% of cases, the boundaries of DR-1 were exceeded at some point after 7.4 seconds following the CE dynamics transition, which is considered a late detection [26]. In no cases was the boundary of DR-1 exceeded in the expected interval of 2.2-7.4 seconds following the transition. Figure 3.7 shows the phase plane along with DR-1 superimposed. As can be seen, the majority of detections came well within its bounds, suggesting the limits proposed by Phatak and Bekey [25] are too high.

However, van Ham noticed the standard deviation of the error and error rate signals pre-transition was extremely consistent amongst all participants, as can be seen in Table 3.2. P1-P6 are the six participants in the experiment, e is the error, and \dot{e} is the error rate. As a result of this observation, alternate limits were proposed that depend on the standard deviation of the error and error rate instead of the maximum. The best compromise between detections that were too early (within 2.2 seconds following the transition) and too late (after 7.7 seconds following the transition) was a value of 3.9σ for both the error and error rate. This would result in 46.9% of detections coming within the expected time interval (2.2-7.4 seconds), 9.4% would be too early, 32.3% would be too late, and 11.5% of cases

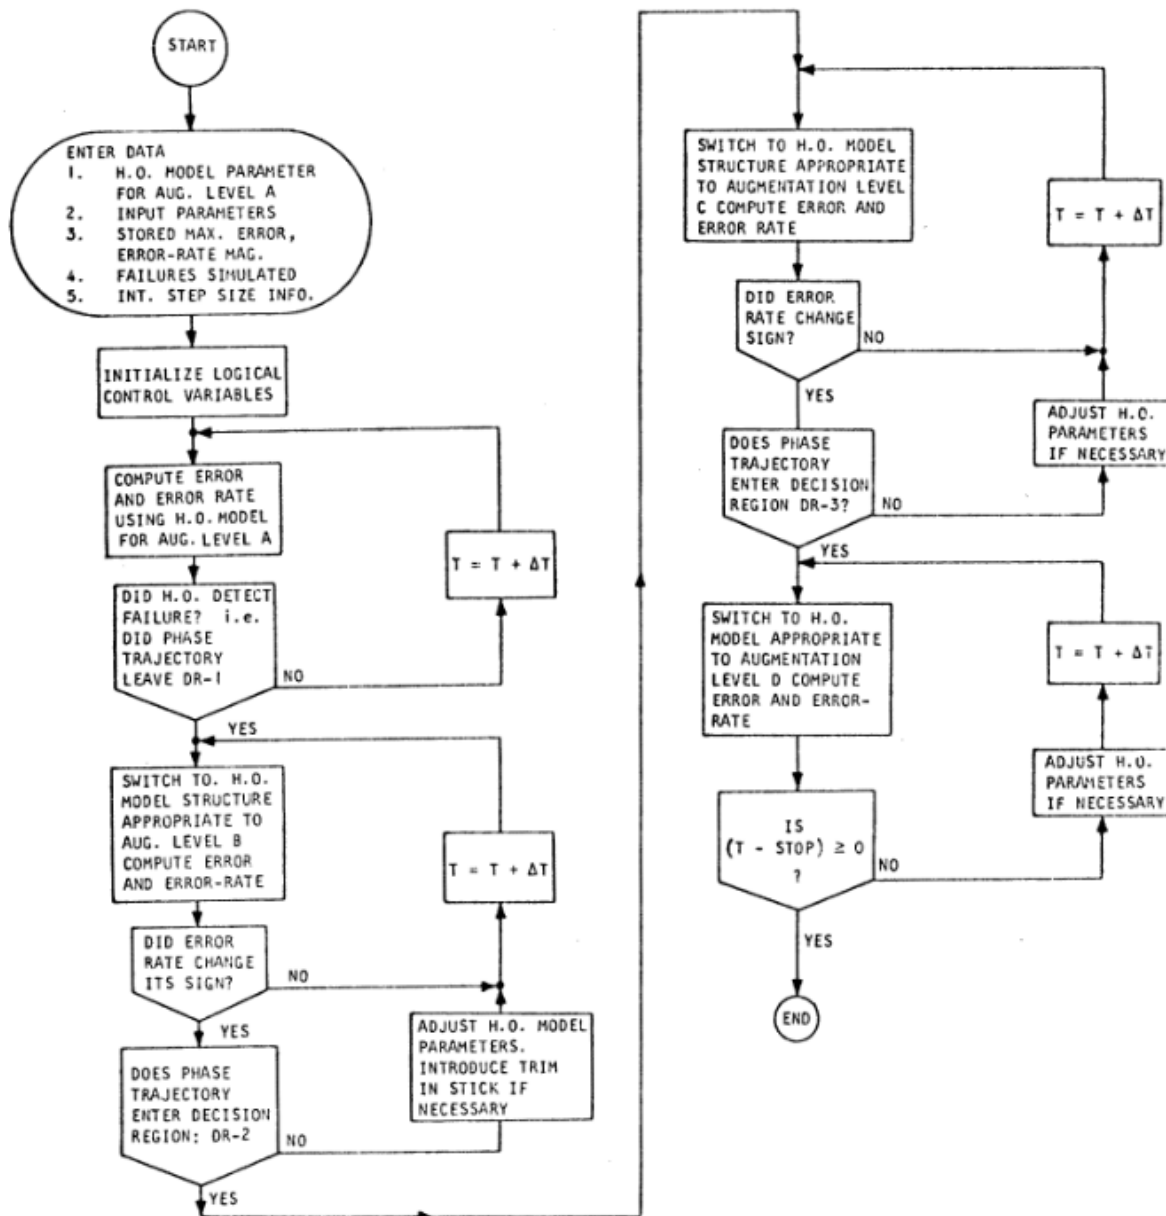


Figure 3.6: Supervisory control algorithm as defined by Phatak and Bekey [25].




Figure 3.7: Phase plane showing the moment the HO detected a change in the CE dynamics with DR-1 superimposed [26].

would result in "no detection", which is similar to the experimental results. As a point of reference, the boundaries of DR-1 are 12σ for the error and 6σ for the error rate.

Table 3.2: Pre-failure tracking standard deviation compared to DR-1 values for all participants in van Ham's experiment [26].

	P1	P2	P3	P4	P5	P6
σ_e [-]	0.023	0.022	0.023	0.023	0.024	0.022
$\sigma_{\dot{e}}$ [-/s]	0.047	0.043	0.046	0.042	0.046	0.043
$2 \cdot \dot{e}_{max} $ [-/s]	0.312	0.291	0.276	0.255	0.299	0.253

It is also interesting to note that the median detection lag (time taken for the participants to press the button on the side-stick following the change in CE dynamics) was 6.6 seconds. In comparison to the values found by Young [24] (Table 3.1), this is significantly longer. Several possibilities exist for explaining this difference. First, Young used pure integrator and double integrator CE dynamics whereas van Ham [26] used approximate single and approximate double integrator CE dynamics, which could make it more difficult for participants to detect a transition. Additionally, it is possible that the participants in Young's experiment were better trained. Finally, since Young's experiment was performed in the late 1960s, it used an analog computer so there may have been some artefacts modeling time-varying systems that the HO could pick up on to aid with the detection of a change in CE dynamics.

Overall, it can be concluded that the model developed by Phatak and Bekey [25] cannot be generalized to different CE dynamics, at least not without some modifications, as evidenced by van Ham's experiment. Another critical limitation is that, since the model relies on large errors and error rates post-transition, its performance would likely be far inferior when considering a transition from approximate double integrator to approximate single integrator dynamics, though this was not tested in van Ham's experiment.

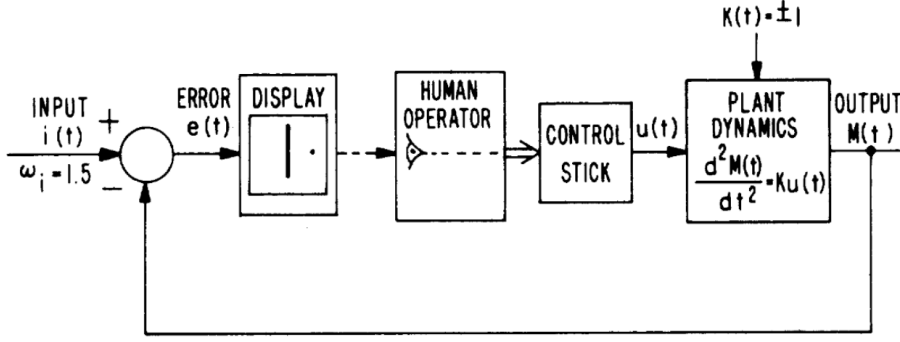


Figure 3.8: Block diagram for Niemela and Krendel's compensatory tracking experiment [28].

3.1.3. Niemela and Krendel Model

Niemela and Krendel [28] designed an experiment to determine the region of error vs error rate wherein the HO's boundary for detecting a change in the CE dynamics lies. The experiment took part in two steps with two distinct (but similar) experiments. In the first one, five participants controlled a double integrator CE using a compensatory display. The block diagram for this setup is given in Figure 3.8. Partway through the run, the polarity of the CE changed and the participants had to adapt their control strategy. From this experiment, it was concluded that the detection moment was characterized by a peak in error rate when the error and error rate had the same sign. Thus, they postulated that the trigger for a HO was some multiple of the steady state root mean square (rms) of the error signal and some (possibly different) multiple of the steady state rms of the error rate signal, which led them to their next experiment.

In the second experiment, one participant took part and controlled the same system but was given an audio warning after the transition when certain boundaries of the error and error rate were exceeded. If the peak error decreased significantly in magnitude, it meant the audio warning helped the participant and thus they had not yet realized the transition had occurred. If it did not help, then it was considered redundant information meaning the participant had already noticed that the transition occurred and was in the process of adapting his behaviour. Consequently, the limits for detection without external (audio) aid must lie somewhere between the points where the audio warning helped improve the peak error and where it did not. This led to the generation of Figure 3.9. The three points labeled "error trajectory features: means and $\pm 1\sigma$ " show the statistical properties of what the authors argue are the three dominant features of the error trajectory immediately post-transition. From top to bottom, these are the maximum positive error rate along with the error at that point, the error when the error rate is 0, and the maximum negative error rate along with the error at that point. For the purpose of identifying the detection boundaries, the primary area to consider is the cross-hatched region labeled "detection boundary location". In the area below and to the left of the cross-hatched region (i.e. below $1\sigma_e$ and to the left of $1\sigma_{\dot{e}}$), the audio warning helped significantly improve the peak error, while in the region above and to the right of the cross-hatched region, it did not. Thus, it led the authors to conclude that the detection boundary used by the participant must lie somewhere in the cross-hatched region. The fact that the maximum positive error rate and error at that point lie to the right of the cross-hatched region supports the conclusion that the detection boundary must lie inside the cross-hatched region.

3.1.4. Hess Model

Most recently, Hess developed a model for adaptive control behaviour with a pursuit display [29] [30]. The model structure is given in Figure 3.10. It is composed of an inner rate loop, with gain K_r , and an outer position loop, with gain K_p , which together make up the model of the HO. The two gains are the only two variables that can be tuned in the model. Hess argues that the former is used as the primary means of adapting to changes in the CE dynamics, while the latter is used to make minor adjustments to improve the tracking performance. The block G_{nm} represents the neuromuscular system dynamics, which is given by Equation (3.1). The limiter block represents the physical limitations of the device being used and is included to make the simulation more realistic. Finally, the vehicle block contains the




Figure 3.9: Location of the detection boundary resulting from Niemela and Krendel's second experiment [28].

CE dynamics, which vary between several cases of approximate gain, approximate single integrator, and approximate double integrator dynamics in Hess' computer simulation.

$$G_{nm}(s) = \frac{10^2}{s^2 + 2(0.707)(10)s + 10^2} \quad (3.1)$$

The adaptive logic is what drives the changes in the gains when there is a change in the CE dynamics. There are two inputs to this block, namely the output rate, \dot{M} , and the signal R , which when subtracted from each other give the inner loop error. The trigger for beginning to adapt is determined based on a signal x defined by Hess as given in Equation (3.2). Thus, the x signal is the sign and (squared) magnitude of the inner loop error multiplied with a low pass filter for smoothing and to introduce lags into the adaptation process. The adaptation process then begins if the value of $K_{trigger}$, defined in Equation (3.3), is equal to 1. Note that the $\text{rms}[\sqrt{|x|}]$ is calculated based on a steady state condition and the time at which the CE dynamics change, t_c , is sufficiently long into the simulation that it has reached an asymptotic value and the HO would have enough time to make an internal model for the x signal. Also note that the square root of x is used in Equation (3.3) since the square of the error is used to calculate the x signal.

$$x = \text{sgn}\{|R| - |\dot{M}|\} \cdot [|R| - |\dot{M}|]^2 \cdot \frac{1.5^2}{s^2 + 2(1.5)s + 1.5^2} \quad (3.2)$$

$$K_{trigger} = \begin{cases} 0 & \text{if instantaneous value of } \sqrt{|x|} < 3 \cdot \text{rms}[\sqrt{|x|}] \text{ or } t < t_c \\ 1 & \text{if instantaneous value of } \sqrt{|x|} \geq 3 \cdot \text{rms}[\sqrt{|x|}] \text{ and } t > t_c \end{cases} \quad (3.3)$$

The motivation for choosing this as the trigger for adaptation is as follows. Hess argues that the primary signal from a HO's internal model that triggers the detection of a change in CE dynamics is the relationship between proprioceptive feedback and the system output rate [31]. The proprioceptive signal allows the HO to have an approximate idea of the position of their limbs and thus an expectation of what the system output rate should be. When this relationship is no longer proportional as would be expected for a pitch or roll tracking task, it leads the HO to begin adapting. Furthermore, Hess made this model around four laws, two of which dictate that triggering changes to the gains K_p and K_r must be based on signals available to the HO, and the adaptation logic must also be based on these signals. In a pursuit tracking task, the output rate can be visually estimated by the HO, and evidence for the

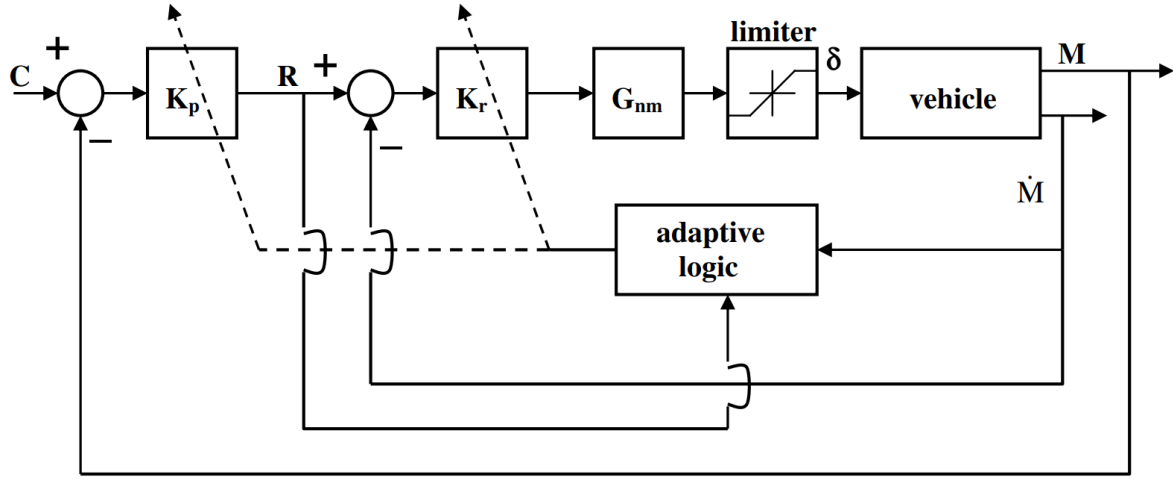


Figure 3.10: Block of diagram of Hess' adaptive model for pursuit tracking tasks [30].

existence of the proprioceptive signal has also been found [32], thus satisfying the two laws.

Once the adaptive logic is triggered, the adjustment of the two gains are given by Equation (3.4)-(3.6) [29]. In Equation (3.5), the $\text{rms}[R^2]$ is also calculated in a steady state condition, and once again the low pass filters introduce smoothing and lag. N is the number of response variables being controlled by the HO, so the factor $\frac{1}{N}$ is present due to the assumption that the HO will be less aggressive in their adaptive response when controlling more than one axis.

$$\Delta K_r = x_n K_{trigger} \cdot \frac{1}{s^2 + 2s + 1} \quad (3.4)$$

$$x_n = \frac{x}{\text{rms}[R^2]} \cdot \frac{1}{s^2 + 2s + 1} \cdot \frac{1}{N} \quad (3.5)$$

$$\Delta K_p = \begin{cases} 0.35 \Delta K_r & \text{if } \Delta K_r > 0 \\ 0 & \text{if } \Delta K_r \leq 0 \end{cases} \quad (3.6)$$

Limits on the maximum change to the gains are also introduced, as defined by Equation (3.7) and Equation (3.8).

$$|K_p|_{max} = 2|K_{p-nom}| \quad (3.7)$$

$$|K_r|_{max} = 10|K_{p-nom}| \quad (3.8)$$

In his original model, Hess demonstrated the capabilities of the adaptive model via computer simulation [30]. Several CE dynamics including approximate gain, approximate single integrator, and approximate double integrator dynamics were tested. These were primarily done for single-axis tracking tasks, but the model's capabilities were also demonstrated for a two-axis task. In general, the model showed good tracking performance both before and after the CE dynamics changed. However, there were significant oscillations when transitioning to dynamics that are more difficult to control.

Later, Hess performed a human-in-the-loop pitch tracking experiment to try and validate the model [29]. The experiment was done using a compensatory display, which differs from the pursuit display for which the model was developed, and one well-trained participant took part. The CE dynamics transitioned to more difficult dynamics during the run to imitate a failure. Hess quantified the quality of fit using the variance accounted for (VAF), as defined in Equation (3.9). The VAF was found to be 0.73 which Hess described as "quite high" [29]. Nonetheless, the tracking performance of the participant was worse than the model prediction. Possible explanations for this are the exclusion of a HO remnant in the model

and pulsive control behaviour which occurred after the CE dynamics transition that can be attributed to the lead generation required [33].

$$\text{VAF} = 1 - \frac{\text{var}(\delta(t_i)_{\text{human}}) - \text{var}(\delta(t_i)_{\text{model}})}{\text{var}(\delta(t_i)_{\text{human}})} \quad (3.9)$$

Jakimovska et al [34] looked further into the model and attempted to validate it with a pursuit tracking experiment using more participants (ten). In the experiment, the CE dynamics transitioned from single to double integrator and vice versa. Several changes were made compared to the original Hess model, but the triggering mechanism remained the same bar one change. That change is that the $\text{rms}[\sqrt{|x|}]$ is calculated recursively (meaning that at the current time t , the values of the x signal from time 0 until time t are used to calculate the rms) which is argued to be more realistic since it represents the HO's internal model that is being continuously tuned during the run. With this change and an increase in the break frequency of the low pass filter in Equation (3.2) from 1.5 rad/sec to 3 rad/sec, the triggering mechanism was appropriately triggered for eight of the ten participants in the transition from approximate single to approximate double integrator (DYN12) CE dynamics. For these cases, the model was then able to appropriately capture the adaptive behaviour of the HO. However, the triggering mechanism was not triggered for the other two participants due to a relatively large tracking error pre-transition as a result of them adopting a low-gain strategy and thus having a low crossover frequency, which highlights the model's sensitivity to the pre-transition tracking performance.

For the transition from approximate double to approximate single integrator (DYN21) CE dynamics, however, the triggering mechanism could not be activated for any of the participants. In some cases, the x signal was even seen to decrease in magnitude after the transition when keeping the same gains used to control the double integrator dynamics. The VAF of the output also increased after the transition (without adaptation of the gains) due to the model's ability to model the single integrator condition better than the double integrator condition. Thus, given that the post-transition tracking error never became abnormally large due to the easier dynamics, it proved impossible to trigger the triggering mechanism.

Overall, Hess' model has several advantages and also significant disadvantages. The advantages include that it is a simple rule-based model based on signals available to the HO, and the pre- and post-transition open loop dynamics of the pilot-vehicle system are integrator-like in the crossover region, as dictated by the crossover model. However, the biggest disadvantage is that the triggering mechanism cannot be activated when transitioning from more difficult dynamics to simpler dynamics. Furthermore, the adaptive rules are based on the author's own experience and are thus not theoretically grounded. Finally, parameters must be tuned for specific participants in order to activate the triggering logic, which is not known apriori.

3.2. Models Without Explicit Boundaries for Detection

Attempts have also been made to develop adaptive models in which the HO's parameters are continuously estimated (either online or offline). Thus, they can identify that the HO detected a change in the CE dynamics and consequently changed their control behaviour by looking at the values of the estimated parameters, but the underlying reason for the change is not investigated.

3.2.1. Recursive Autoregressive-Exogeneous Model

Plaetinck [35] used a low-order recursive autoregressive-exogeneous (ARX) model to identify online when the HO detects a change in the CE dynamics in compensatory tracking. It is based on an experiment in which participants could indicate that they noticed a change in the CE dynamics by pressing a button on the side-stick. Data from that experiment was used by Plaetinck to test two methods of determining the moment the HO detected the transition. A block diagram of the setup is given in Figure 3.11. $H_p(s, t)$ is the linear part of the HO, $H_n(s, t)$ is the noise dynamics used to model the HO remnant, and $H_c(s, t)$ is the CE dynamics, all of which are time-varying.

The HO dynamics, $H_p(s, t)$, are given in Equation (3.10). K_e is the error response gain, T_L is the lead time constant, τ is the HO delay, and $H_{nm}(s)$ is the neuromuscular system transfer function. Additionally, Plaetinck introduces the error rate response gain, $K_{\dot{e}}(t)$, defined as $K_e(t) \cdot T_L(t)$. The error




Figure 3.11: Block diagram for Pläetinck's compensatory tracking experiment [35].

response gain and error rate response gain, as well as the neuromuscular system parameters, were calculated based on the ARX coefficients estimated via recursive least squares. It was found that the neuromuscular system parameters do not vary significantly in the DYN21 transition so the focus is on the error response gain and the error rate response gain.

The two methods for determining the moment the HO noticed the change in CE dynamics alluded to earlier are time-invariant condition average (TICA) and moving average (MA), both of which were tested on K_e and $K_{\dot{e}}$. The TICA method involves averaging the parameter values for a steady-state condition of the pre-transition dynamics, and the MA method looks at the average over a window of a certain number of samples up to the current time. In both methods, if the current value of K_e or $K_{\dot{e}}$ exceeds the bounds of the average plus or minus a certain margin (unique to each method) for a period of 3 seconds, that is the moment it is hypothesized the HO detected the change in CE dynamics.

$$H_p(s, t) = K_e(1 + T_L s)e^{-s\tau} H_{nm}(s) \quad (3.10)$$

Pläetinck defined the accuracy of a method based on Equation (3.11). 'TP' stands for true positive, 'TN' for true negative, 'FP' for false positive and 'FN' for false negative. By this metric, the best performing method was using TICA on the error rate response gain $K_{\dot{e}}$, which resulted in an accuracy of 57%. The accuracy of the other options evaluated is given in Table 3.3.

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{TN} + \text{FN}} \quad (3.11)$$

Table 3.3: Accuracy and detection lags for the methods analysed by Pläetinck [35], compared with the HO performance.

Method	Accuracy	Median Detection Lag (secs)
TICA on K_e	22%	-
TICA on $K_{\dot{e}}$	57%	6.5
MA on K_e	12%	-
MA on $K_{\dot{e}}$	43%	4.6
Participants	88%	5.2

However, in the experiment, the accuracy of the participants was found to be 88%, so even the TICA method on $K_{\dot{e}}$ does not perform as well as the HO, due to its high number of false positives. On the other hand, when looking only at true positives, there is no statistically significant difference in the time taken to detect a change following the CE transition between the HO and TICA method on $K_{\dot{e}}$. As can be seen in Table 3.3, the TICA method on $K_{\dot{e}}$ was 1.3 seconds slower in detecting the transition than the HO. In comparison to the detection lags found by van Ham [26] (median of 6.6 seconds), both the HO and, especially the TICA method on $K_{\dot{e}}$, were very close. Given that the dynamics used in this experiment were the same as those used in van Ham's experiment, it further suggests that the choice of CE dynamics could be one of the reasons for the significantly longer detection lags when compared to Young [24].

Overall, the similarity in detection lags between the HO and TICA method on K_e , as well as the reasonable accuracy of the TICA method on K_e suggests that the model, at a minimum, captures some elements of what the HO uses to detect a CE transition. However, only the DYN12 transition was considered in this study, which limits the applicability to a transition the other way around.

3.2.2. Artificial Neural Networks

Jiao et al [36] extended the Hess model [29] [30] by adding a lead time constant and time delay for the HO, and attempted to identify the three parameters (the two gains from Figure 3.10 and the lead time constant) using an artificial neural network (ANN). The triggering mechanism and adaptive logic from Hess' model were therefore eliminated and replaced with the ANN's estimates of the three parameters identified from experimental data. For the identification procedure, the authors use a single-layer linear ANN (SLLANN). The inputs to the SLLANN are the current error, error rate, control input, and output rate, and the output of the SLLANN is the system output. Four experimental conditions were carried out in which there was a control system failure, aircraft failure, or actuator failure. The experiment used a compensatory display despite the model being developed for pursuit displays, though Jiao et al [36] argue that well-trained participants are able to adopt behaviour typical of a pursuit display using a compensatory display. In the ten second period following the failure, the ANN-based identification procedure was used to estimate the values of the three parameters, and the system output was compared to the HO performance observed in the experiment. The results for the identification can be seen in Figure 3.12. One big limitation of the method in the context of identifying when the HO detected the failure can be observed here, namely that the identification procedure is only done for the 10 second period following the failure. As a result, there is a discontinuity in the parameter estimates at $t = 70$ secs and $t = 80$ secs which hides the moment the HO detected the change and began adapting their behaviour.

Overall, the adaptations observed in Figure 3.12, are consistent with the expected results based on the specific CE transitions. The VAF of the model when compared to the experimental results for the ten second period following the failure is 0.73, 0.89, 0.86, 0.93 for the four experimental conditions, respectively. The accuracy is therefore concluded to be "satisfactory" by Jiao et al [36].

3.2.3. Model Reference Adaptive Control

Terenzi [37] and Tohidi [38][39] have developed adaptive pilot models based on model reference adaptive control (MRAC) for adaptive control behaviour with pursuit displays. MRAC models have several attractive features for modeling the time-varying HO, namely that they are adaptable, include an internal model, and there is a guarantee that the model will follow the chosen internal model. The block diagram from Terenzi's research, which has a very similar structure to Tohidi's work, is given in Figure 3.13. The HO is modeled using the gains k_r and K_x , which are constantly being adapted based on the mismatch between the internal model and the actual output. The adaptive law, derived based on the Lyapunov stability theorem, defines how the gains change throughout the run. The block H_{mCL} represents the (closed loop) internal model of the HO. Terenzi used the closed loop form (Equation (3.13)) of McRuer's open loop model (Equation (3.12)).

Using this model, Terenzi [37] found that the steady-state performance of the MRAC model was inferior to a state-of-the-art time-varying parametric model [40]. However, the model was able to approximate the DYN12 transition well, as evidenced by an increase in VAF following the transition. On the other hand, the DYN21 transition was not well modeled, with a significant decrease in the model's VAF following the CE transition. Terenzi [37] argues that this could be because the gains do not change quickly enough, and increasing their rate of change can make the controller unstable. Thus, it was concluded that transitions to increasing orders of CE dynamics could be modeled using the proposed MRAC model, but not transitions where the order of the CE dynamics decreased.

$$H_{mOL} = \frac{\omega_c}{s} \cdot e^{-\tau s} \quad (3.12)$$

$$H_{mOL} = \frac{H_{mOL}}{1 + H_{mOL}} \quad (3.13)$$

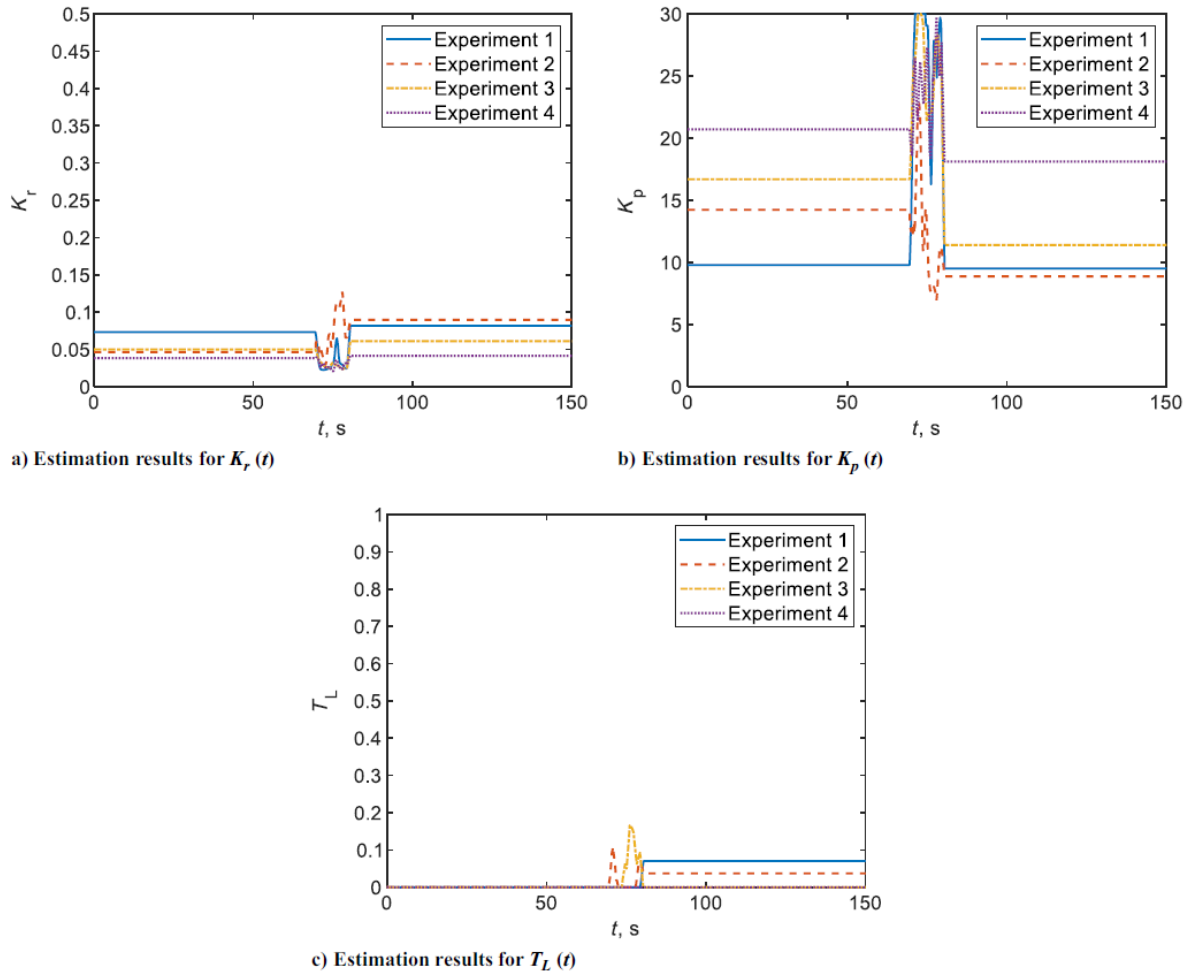


Figure 3.12: Results of the ANN-based identification procedure [36].

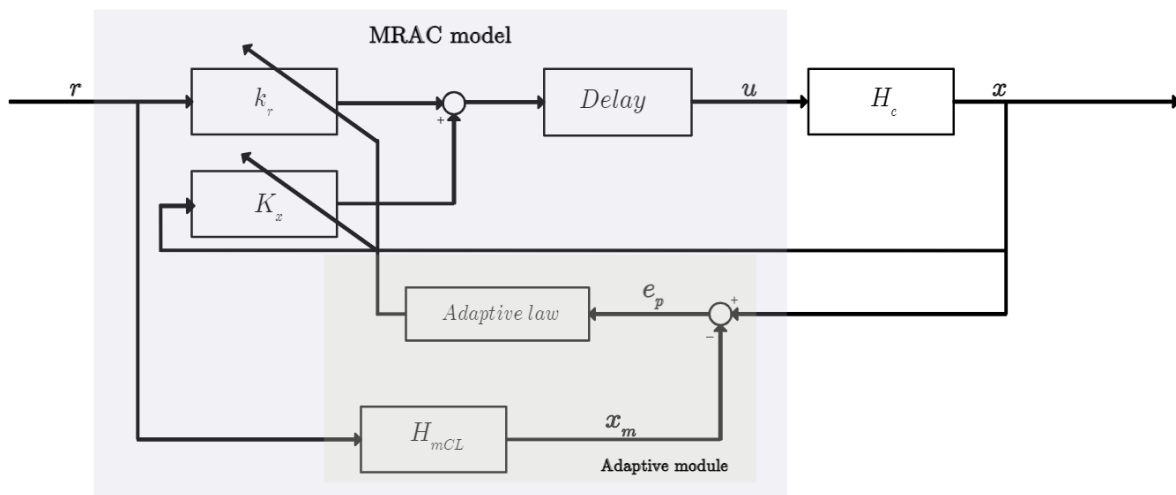


Figure 3.13: Block diagram of Terenzi's MRAC model [37].

One limitation of Terenzi's [37] research is that the internal model must be time-invariant, meaning the crossover frequency is the same for both CE dynamics conditions, but a change in crossover frequency is possible after both the DYN12 and DYN21 transitions. To bypass this, Tohidi [39] kept approximate single integrator dynamics for both conditions in his experiment and instead modified the CE dynamics by changing the location of the pole and gain by 50%. The internal model used is the same as Terenzi [37] and the adaptive law was also derived based on the Lyapunov stability criteria. With Tohidi's [39] experimental setup, the model results were significantly better with the average deviation between the model and observed behaviour in the experiment being very close to 0, and the standard deviation of the difference also being small. The estimated frequency response for the participants also closely matched the frequency response of the adaptive model before and after the change in CE dynamics.

3.3. Conclusions

Overall, several important conclusions can be drawn from the literature review. First, the majority of research efforts have focused on compensatory displays. Out of all of the models described in this chapter, only Hess' model [29] [30] and the MRAC models just described were developed to capture adaptive control behaviour with pursuit displays. Since compensatory style displays are an exception in real-life applications, it would be beneficial to dedicate further efforts on analysing adaptive behaviour with pursuit displays. Secondly, and as a consequence of the first point, the majority of research efforts have focused on analysing the characteristics of the error and/or error rate signals, with limited inclusion of the properties from other signals such as the control input, system output, or the forcing function. When using pursuit displays, more information is available to the HO that can aid them in detecting changes in the CE dynamics, so the role this additional information plays should be investigated. Finally, using the error and error rate has shown success in modeling the DYN12 transition as a result of the system becoming unstable if no action is taken by the HO, but the same cannot be said for the DYN21 transition. Since the system becomes more stable, there are no large peaks in the error or error rate which can be used by the HO to detect a change in CE dynamics, or in a model to trigger the adaptive logic. Thus, it remains beneficial to study the DYN21 transition using a pursuit display and the additional information (such as explicit knowledge of the system output and forcing function) that comes with it to gain a better understanding of the process that a HO goes through to detect a change in the CE dynamics. Such understanding would be helpful in developing better adaptive models that trigger their adaptive logic at an appropriate moment and can be used to develop better human-machine interfaces, support systems, and human-like automation.

4

Simulation Results

In order to get an idea of what the most important aspects to investigate during the experiment are, as well as to get a preview of what the experiment results may look like and formulate a preliminary answer to the second sub-research question ("What are the threshold values for the relevant signals?"), a simulation was made in which an adaptive pilot model is compared with a constant (non-adaptive) pilot model. This allows for the comparison of different signals in the loop, such as the error, control input, and system output (along with their rates of change), which leads to hypotheses that can be tested in the experiment. The design of the simulation will be discussed in Section 4.1 followed by the results and a discussion of the results in Section 4.2. Finally, conclusions from the analysis will be presented in Section 4.3.

4.1. Simulation Design

The simulation was made in MATLAB and Simulink for a compensatory tracking task. Despite the fact that the experiment to be performed for this research will use a pursuit display, a compensatory model was chosen for the simulation because there is a universally accepted model for it (McRuer's crossover model [21]), and it still allows for the analysis of the majority of signals present with a pursuit display. The block diagram for the compensatory tracking system is shown in Figure 4.1. As can be seen, there are four blocks containing time-varying dynamics. They are the HO dynamics (including the neuromuscular system), $H_p(t)$, the HO delay, $H_{del}(t)$, the CE dynamics, $H_c(t)$, and the noise dynamics, $H_n(t)$. All of the time-varying parameters can take one of two values, one of which corresponds to the approximate single integrator condition (DYN1) and the other to the approximate double integrator condition (DYN2). The variation from one value to the other is done according to the sigmoid function defined by Equation (4.1).

$$P(t) = P_1 + \frac{P_2 - P_1}{1 + e^{-G(t-M)}} \quad (4.1)$$

Where P_1 and P_2 are the beginning and end values, respectively, G defines how quickly the transition from one value to the other takes place, and M defines the moment of the maximum rate of change. For the simulation, G was set to 100 sec^{-1} , which corresponds to a near instantaneous change in the parameter value, and M was set to 60 sec.

In the following subsections, the forcing function will be described, along with the exact dynamics in each of the four blocks from Figure 4.1.

4.1.1. Forcing Function

As is commonly used for manual tracking tasks, the forcing function, f_t , is a sum of sine waves, as defined by Equation (4.2).

$$f_t = \sum_{n=1}^{10} A_t[n] \sin(\omega_t[n]t + \phi_t[n]) \quad (4.2)$$

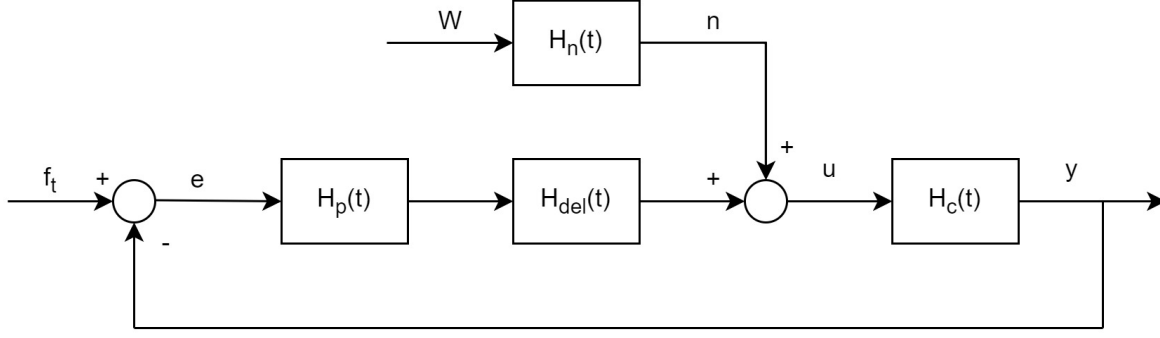


Figure 4.1: Block diagram for the compensatory tracking system in the simulation.

Where $A_t[n]$ is the amplitude of the n th sine wave, $\omega_t[n]$ is the frequency, and $\phi_t[n]$ is the phase shift. Ten sine waves are used to construct the forcing function, which is why the summation is until ten. In total, 100 different forcing function realizations were used in the simulation. However, for verifying the simulation, the values for the amplitude, frequency, and phase shift given in Table 4.1 are used since those are the ones used in Terenzi's experiment [37]. For generating the other 99 forcing functions, everything except the phase shift ϕ_t was kept constant. For the phase shifts, 100,000 random sets of ten phases were generated, and the first 99 forcing functions that resulted in an average crest factor, as defined by Equation (4.3), were chosen to be used in the simulation in addition to the one specified in Table 4.1.

Table 4.1: Parameters of the ten sine waves used to construct the forcing function.

k [-]	n_t [-]	ω_t [rad/sec]	A_t [rad]	ϕ_t [rad]
1	2	0.419	$2.905 \cdot 10^{-2}$	2.841
2	5	1.047	$1.916 \cdot 10^{-2}$	3.319
3	9	1.885	$1.020 \cdot 10^{-2}$	0.718
4	13	2.723	$6.032 \cdot 10^{-3}$	0.768
5	19	3.979	$3.356 \cdot 10^{-3}$	2.925
6	27	5.655	$1.983 \cdot 10^{-3}$	5.145
7	39	8.168	$1.230 \cdot 10^{-3}$	2.085
8	51	10.681	$9.331 \cdot 10^{-4}$	0.383
9	67	14.032	$7.541 \cdot 10^{-4}$	0.763
10	83	17.383	$6.674 \cdot 10^{-4}$	3.247

$$\text{Crest Factor} = \frac{\max(|f_t(t)|)}{\text{rms}(f_t(t))} \quad (4.3)$$

It should be noted that, in order to avoid spectral leakage, the frequency of all of the sine waves must be an integer multiple n_t of the fundamental frequency defined as $\omega_m = \frac{2\pi}{T_m}$, where T_m is the measurement time. While the simulation lasts 90 sec, like in Terenzi's experiment, the measurement time was taken to be 30 sec, so the fundamental frequency is 0.209 rad/sec [37]. Thus, there are three periods of the signal in the simulation, of which the first is used as run-in time, the second to analyze the first condition (for example the double integrator condition in DYN21), and the third to analyze the second condition (for example the single integrator condition in DYN21). This is why M in Equation (4.1) is set to 60 sec.

4.1.2. Controlled Element Dynamics

The first of the dynamics blocks to be discussed is the CE dynamics. Again, to be consistent with Terenzi's [37] research, similar dynamics to what was used in his experiment will be implemented in the simulation. The structure of the CE dynamics is defined by Equation (4.4).

$$H_c(t) = \frac{k_c}{s(s + \omega_b(t))} \quad (4.4)$$

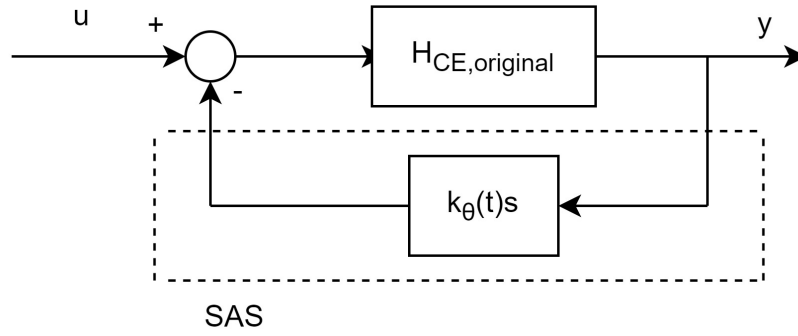


Figure 4.2: Hypothetical SAS implementation.

In Terenzi's experiment, both the gain, k_c , and the break frequency, ω_b , are time-varying. For this application, however, only the break frequency will be time-varying, as will be discussed in this subsection. The break frequency can be varied from a low value to approximate a double integrator, to a high value to approximate a single integrator, while the gain (if it were time-varying) could be varied to keep the required control activity approximately the same for the two conditions. The values of k_c and ω_b for the single and double integrator-like conditions are shown in Table 4.2.

Compared to Terenzi's experiment [37], there are two differences. The first is that the break frequency in the single integrator condition is increased from 6 rad/sec to 20 rad/sec. The reason for this change is that having a break frequency at 6 rad/sec means the effect of the second integrator will already be felt in the crossover region, so to make the condition more closely approximate a pure single integrator, the break frequency for the second integrator was put further away from the crossover region. The second difference is that the gain is kept constant at 90, instead of being 90 for the DYN1 condition and 30 for the DYN2 condition. This is to make the failure case more realistic. If the gain were to change from 90 to 30, it would not be possible to model the DYN12 transition as a single failure of the stability augmentation system (SAS), for example, but rather it would have to be two simultaneous failures, which is less realistic. Consider a situation where the DYN2 dynamics represent the CE dynamics without a SAS. If a SAS is added, as shown in Figure 4.2, the break frequency can be varied by changing the value of k_{θ} , but the gain does not change. In this case, the block $H_c(t)$ in Figure 4.1 can be considered the effective CE dynamics. With the SAS active and $k_{\theta} = 0.22$, the effective CE dynamics become the DYN1 dynamics given in Table 4.2, but if the SAS fails (i.e. $k_{\theta} = 0$), the effective CE dynamics become the DYN2 dynamics given in Table 4.2. One downside of this change, however, is that the level of control activity is no longer approximately constant between the two conditions. In the DYN1 condition, the standard deviation of the control inputs needs to be approximately four times larger than in the DYN2 condition.

Table 4.2: Parameters of the CE dynamics for the two dynamics cases.

Dynamics case	k_c	ω_b [rad/sec]
DYN1	90	20
DYN2	90	0.2

The bode plots for the two conditions are shown in Figure 4.3. As can be seen, when the break frequency is 20 rad/sec, the system resembles single integrator dynamics for the majority of the crossover region. For the double integrator condition, the break frequency is 0.2 rad/sec, which is well below the crossover frequency and thus the pilot will have to generate lead to stabilize the system. The effect of keeping the gain constant at 90 is to shift the magnitude line for the DYN2 condition vertically upwards compared to the case where the gain would change to 30 for the DYN2 condition.




Figure 4.3: Bode plots for the two CE dynamics conditions.

4.1.3. Human Operator Dynamics

The remaining three blocks from Figure 4.1, all relate to the HO, though in this subsection only the $H_p(t)$ and $H_{del}(t)$ blocks will be discussed. The dynamics used to describe the HO are the ones used by McRuer in the crossover model, plus the neuromuscular system dynamics [21]. They have the form shown in Equation (4.5).

$$H_p(t) = k_p(1 + T_L s) \cdot e^{-\tau_e s} \cdot \frac{\omega_{nm}^2}{s^2 + 2\zeta_{nm}\omega_{nm} + \omega_{nm}^2} \quad (4.5)$$

Since the simulation will only be used to model the DYN12 and DYN21 transitions, the HO will not need to generate lag and thus that term is omitted. The value of the parameters in the neuromuscular system dynamics are set to $\omega_{nm} = 15$ rad/sec and $\zeta_{nm} = 0.7$. These values remain constant for both dynamics cases which is a limitation, but previous research has found that there is little variation in these parameters for the CE dynamics considered here [41].

The values for k_p , T_L , and τ_e were determined based on the verbal adjustment rules, but with some adjustments [21]¹. The first change is that for the single integrator condition, despite increasing the break frequency of the CE dynamics to 20 rad/sec, some effect of the second integrator is already felt in the crossover region so it was not possible to get the same crossover frequency and phase margin as dictated by the verbal adjustment rules. Instead, the gain k_p was tuned to get the same phase margin resulting from the verbal adjustment rules, while decreasing the crossover frequency. For the double integrator condition, this was not necessary. The second change accounts for the fact that the neuromuscular dynamics, which are not included in McRuer's model, are in the simulation. Since the phase lag caused by the neuromuscular system increases the time delay in the system, the effective time delay τ_e was decreased by the amount of delay added by the neuromuscular system

¹Note the bandwidth used for the calculations is 2.722 rad/sec which corresponds to the frequency of the fourth lowest sine wave since it was determined that beyond this point the amplitude of the sine waves drop off considerably.

at the respective crossover frequencies. Taking into account these modifications, the values for the parameters in the HO dynamics are summarised in Table 4.3, along with the corresponding crossover frequency, ω_c , and phase margin, φ_m .

Table 4.3: Parameters of the HO.

Dynamics case	k_p	T_L [sec]	τ_e [sec]	ω_c [rad/sec]	φ_m [deg]
DYN1	0.840	0	0.09	3.72	40.0
DYN2	0.008	5	0.23	3.60	23.0

A further two conditions arise post-transition in the time-varying conditions due to the constant pilot keeping the same gain and lead time constant as it had pre-transition but with different CE dynamics. The phase margin and crossover frequency for these two conditions are given in Table 4.4. As can be seen, the open-loop system with the constant pilot remains stable post-transition in the DYN21 condition (the middle row in Table 4.4), but with sub-optimal performance. On the other hand, the open-loop system with the constant pilot becomes unstable post-transition in the DYN12 condition (the bottom row of Table 4.4) due to the lack of lead generation by the pilot.

Table 4.4: Phase margin and crossover frequency for the constant pilot model post-transition in the time-varying conditions.

Condition	ω_c [rad/sec]	φ_m [deg]
DYN1 CE with DYN2 HO	0.0366	99.6
DYN2 CE with DYN1 HO	8.51	-92.0

The $H_{del}(t)$ block represents the HO delay, which is the exponential term in Equation (4.5). This was modeled using the 'variable time delay' block in Simulink, which simply delays outputting the signal value by the amount of the time delay.

4.1.4. Noise Dynamics

The final dynamics which must be modelled are the noise dynamics, which serve to model the HO's remnant. McRuer's model is a quasi-linear model, meaning it is the linearised equivalent of a nonlinear system. Humans are inherently nonlinear, so to make the simulation more realistic, noise is added to account for these nonlinearities.

While there is no clear consensus on how to model HO remnant, it was decided to use Gaussian white noise with unit variance passed through a filter as done by van Grootheest [42]. He investigated several filters of the form defined by Equation (4.6), and concluded that the second order filter was the best for fitting the remnant. Thus, m is set equal to 2. The time constant T_n is kept at a constant 0.06 sec, consistent with van Grootheest [42].

$$H_n(t) = \frac{k_n(t)}{(T_n s + 1)^m} \quad (4.6)$$

The values of the gain k_n were then solved analytically. The desired value for the ratio of the power (equivalent to variance, σ^2 , in the time-domain) in the control input due to the noise, $\sigma_{u_n}^2$, to the power in the control input, σ_u^2 , is 0.2, as defined by Equation (4.7) [43]. In Equation (4.7), $\sigma_{f_t}^2$ is the power in the control input due to the forcing function (a deterministic value), W is the white noise intensity (equal to 1 since the variance of the Gaussian white noise is 1), and $S_{f_t f_t}(\omega)$ is the power spectral density of the forcing function.

$$\frac{\sigma_{u_n}^2}{\sigma_u^2} = \frac{\sigma_{u_n}^2}{\sigma_{u_n}^2 + \sigma_{u_{f_t}}^2} = \frac{k_n^2 \int_0^\infty \left| \frac{1}{1 + H_p(\omega) H_c(\omega) H_{del}(\omega)} \right|^2 W d\omega}{k_n^2 \int_0^\infty \left| \frac{1}{1 + H_p(\omega) H_c(\omega) H_{del}(\omega)} \right|^2 W d\omega + \int_0^\infty \left| \frac{H_p(\omega) H_{del}(\omega)}{1 + H_p(\omega) H_c(\omega) H_{del}(\omega)} \right|^2 S_{f_t f_t}(\omega) d\omega} = 0.2 \quad (4.7)$$

To facilitate solving for k_n , let a equal the integral in the numerator of Equation (4.7) and let b equal the second integral in the denominator. Then, k_n can be solved for analytically as given in Equation (4.8).

$$\frac{k_n^2 a}{k_n^2 a + b} = 0.2 \rightarrow k_n^2 a = 0.2 k_n^2 a + 0.2 b \rightarrow k_n = \sqrt{\frac{1}{4} \cdot \frac{b}{a}} \quad (4.8)$$

The final expression for k_n is presented in Equation (4.9). Note that the integral represented by b is the power in the control input due to the forcing function multiplied with π . This is a deterministic value that could be calculated using the simulation by setting the noise equal to 0 and finding the variance of the control input. The integral in the denominator was solved numerically.

$$k_n = \frac{1}{2} \cdot \sqrt{\frac{\sigma_{u_{ft}}^2 \pi}{\int_0^\infty \left| \frac{1}{1+H_p(\omega)H_c(\omega)H_{del}(\omega)} \right|^2 \left| \frac{1}{T_n(j\omega)s+1} \right|^2 W(\omega) d\omega}} \quad (4.9)$$

Solving Equation (4.9) for the DYN1 and the DYN2 conditions resulted in the k_n values given in Table 4.5.

Table 4.5: Values of the noise filter gain k_n .

Dynamics case	k_n [-]
DYN1	0.00136
DYN2	0.000257

4.1.5. Simulink Implementation

The implementation of the entire system in Simulink is given in Figure 4.4. The top half of the figure is the simulation for the adaptive pilot model and the bottom half (where all of the boxes have a '1' after their name) is the simulation for the constant pilot model. The following blocks were used to construct the simulation:

- From Workspace: for importing the time trace of each of the time-varying parameters, defined in the MATLAB script
- To Workspace: for getting the outputs of the simulation (system output, control input, and error) to the workspace for analysis
- Variable time delay: for the HO delay
- Transfer Fcn: for the time-invariant portion of the transfer functions
- Integrator
- Product
- Sum

4.1.6. Simulation Verification

To verify that the simulation is correctly implemented in Simulink, it needs to be compared with a much simpler model, which is known to be correct. For this, an independent, time-invariant simulation was made using the *lsim* command in MATLAB. The time-invariant simulation, hereafter referred to as the verification model, is simpler than the complete time-varying simulation in Simulink since it is only necessary to define the forcing function, noise input, and transfer functions for the four blocks in Figure 4.1. Thereafter, the closed loop system could be defined using the *feedback* command.

First, the Simulink model was used for steady-state (time-invariant) simulations to allow for direct comparison with the verification model. All of the possible input-output relationships were checked to ensure they were correctly implemented and the parameter values were correctly defined. This involved checking only the response to the forcing function, whereby the noise input was set to 0, checking only the response to the noise input, whereby the forcing function was set to 0, and checking the response to both inputs. The similarity between the verification model and the complete Simulink model was quantified by looking at the rms. The results are summarised in Table 4.6. Note that for generating these

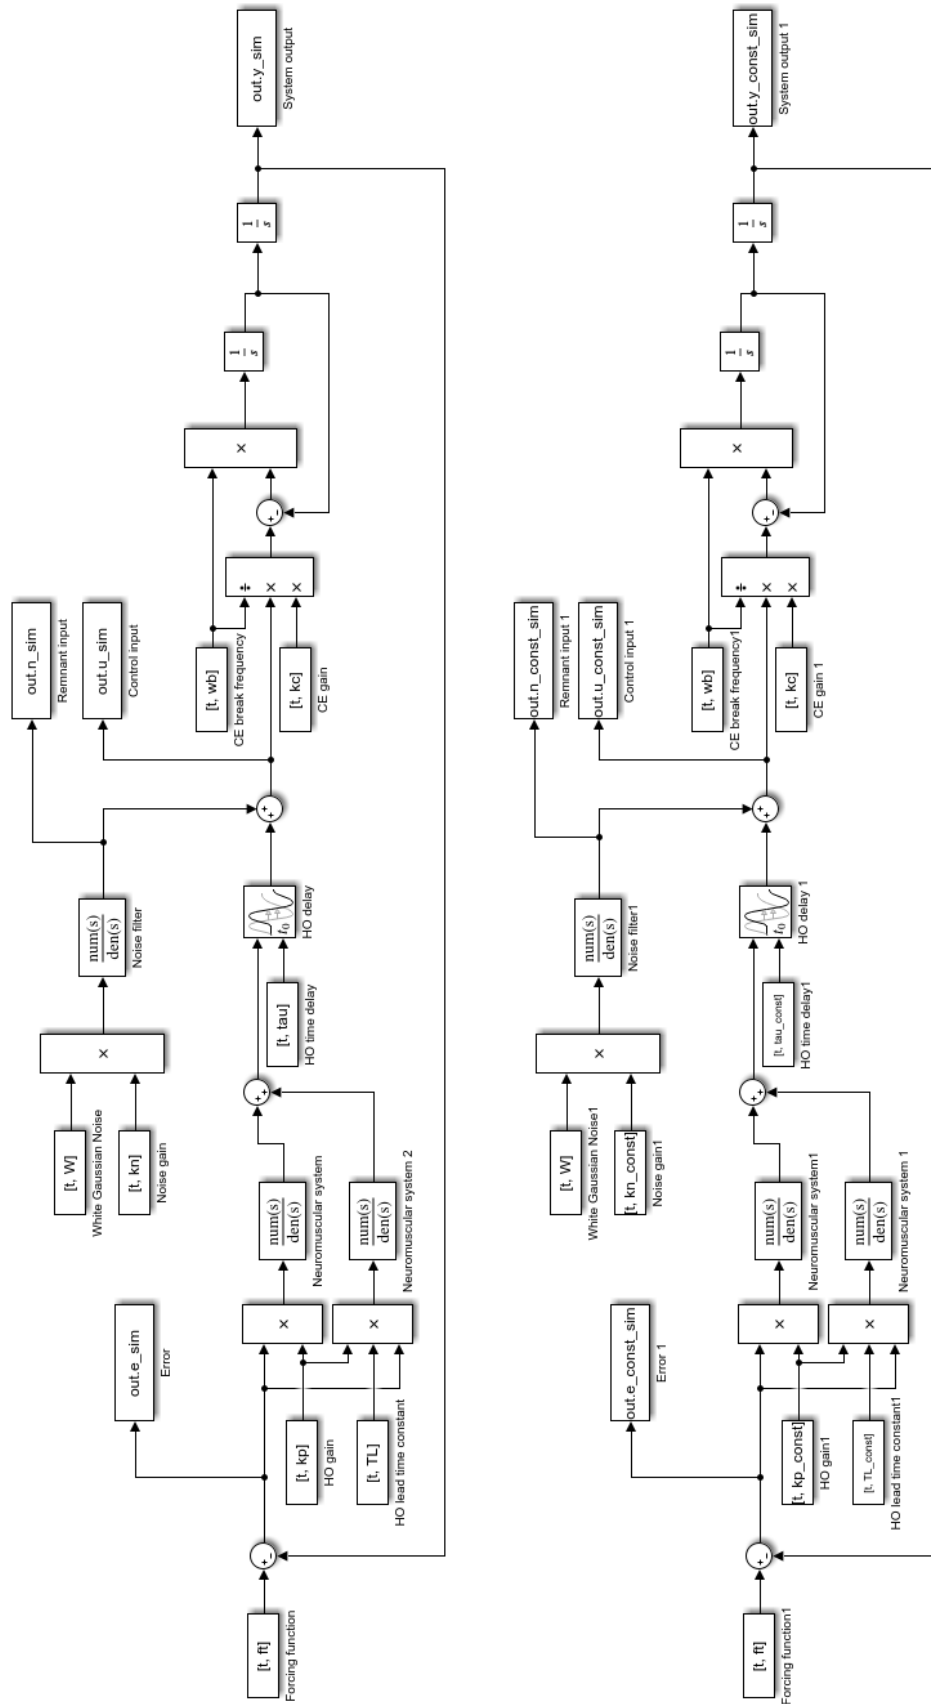


Figure 4.4: Implementation of the compensatory tracking simulation in Simulink. The top half is the simulation for the adaptive pilot, while the bottom half is the constant pilot.

results, the CE gain k_c was varied from 90 to 30 for the DYN1 and DYN2 conditions respectively, instead of keeping the gain constant at 90 which is the condition that was introduced in Subsection 4.1.2 and will be used for the analysis in Section 4.2. This is because that was the original condition implemented in the simulation. Also note that the structure of the system does not change between the single and double integrator conditions, so given that one is correct, the other one should also be correct. Both are included primarily for the purpose of checking that the parameter values have been correctly defined.

As can be seen, the difference between the two models is at most 0.27%, and in most cases significantly less than that. Such small differences can partially be explained by the fact that a third order Padé approximation was used to approximate the time delay in the verification model, while the 'Variable time delay' block was used in Simulink. Another source of differences is in the numerical solvers used by the two models. Thus, given these two factors and the fact that the differences are very small, the Simulink model can be considered verified for the time-invariant conditions based on the results of Table 4.6.

Table 4.6: Comparison of the rms values in steady-state conditions for the two simulations.

Dynamics case	Input	Output	RMS Verification Model	RMS Simulink Model	Difference
Single integrator	f_t	y	0.0270	0.0270	<0.01%
Single integrator	f_t	u	0.00755	0.00755	<0.01%
Single integrator	W	y	0.00336	0.00336	0.01%
Single integrator	W	u	0.00384	0.00384	<0.01%
Double integrator	f_t	y	0.0286	0.0286	0.16%
Double integrator	f_t	u	0.00528	0.00528	0.14%
Double integrator	W	y	0.0109	0.0109	<0.01%
Double integrator	W	u	0.00269	0.00269	<0.01%
Single integrator	f_t & W	y	0.0270	0.0270	<0.01%
Single integrator	f_t & W	u	0.00817	0.00817	0.02%
Double integrator	f_t & W	y	0.0297	0.0297	0.17%
Double integrator	f_t & W	u	0.00554	0.00552	0.27%

For the time-varying conditions, the only difference is that the (time-varying) parameters have a different beginning and end value, which varies according to the sigmoid function defined in Equation (4.1). Therefore, the first test was to visually check the time traces of all of the time-varying parameters to make sure the beginning and end values were correct. Next, the rms values pre- and post-transition were also checked to ensure they were similar to the values during steady-state tracking. The results are shown in Table 4.7. Note that the rms values for the pre-transition cases are calculated based on the signal between $t = 0$ secs and $t = 50$ secs for both of the simulations, and the rms values for the post-transition cases are calculated based on the signal between $t = 70$ secs and $t = 90$ secs. This is to minimize the effect that transients from the transition have on the rms since it cannot be captured by the verification model. As can be seen, with the exception of the post-transition output y in the DYN12 condition, the maximum difference is 0.40%, which is sufficiently small to consider the Simulink model verified given the differences between the models previously described and the fact that some transients are still present 10 seconds following the transition. For the post-transition output in the DYN12 condition, a further check was performed to see why the difference presented in Table 4.7 is so large. By looking at the rms values for the period between $t = 75$ secs and $t = 90$ secs, the difference in rms between the two models for y decreased to 0.11%, and for u it decreased to <0.01%. Thus, the difference can be explained by a longer transient post-transition when transitioning to less stable dynamics.

Finally, a check was done to ensure the constant pilot was also correctly implemented. The first step was to check that its output was exactly the same as the adaptive pilot for the entire simulation in steady-state conditions, or up until the transition for the time-varying conditions, which was the case. The other check was to visually inspect the time traces of the HO parameters to ensure they stayed constant throughout the simulation, which was also the case. As such, the complete time-varying in Simulink was verified and can now be used for analysis.

Table 4.7: Comparison of the rms values in the time-varying conditions with both inputs.

Dynamics case	Pre- or Post-transition?	Output	RMS Verification Model	RMS Simulink Model	Difference
DYN12	Pre-transition	y	0.0275	0.0275	<0.01%
DYN12	Pre-transition	u	0.00819	0.00819	0.07%
DYN12	Post-transition	y	0.0268	0.0275	2.64%
DYN12	Post-transition	u	0.00586	0.00583	0.40%
DYN21	Pre-transition	y	0.0303	0.0303	0.17%
DYN21	Pre-transition	u	0.00534	0.00532	0.25%
DYN21	Post-transition	y	0.0279	0.0279	0.02%
DYN21	Post-transition	u	0.00806	0.00805	0.14%

4.2. Simulation Results

In this section, the results from the simulation will be presented and discussed. First, the methodology used in the analysis will be presented in Subsection 4.2.1, and then the results and conclusions for the DYN12 and DYN21 conditions will be presented in Subsection 4.2.2 and Subsection 4.2.3, respectively.

4.2.1. Methodology

As previously stated, the goal of the simulation is to get a sense of what the experiment results may look like and develop hypotheses that can be tested in the experiment. To ensure the results were not specific to a particular property in the chosen forcing function or to the remnant realization, 100 distinct forcing functions and 25 remnant realizations for each were simulated. Thus, the results to be presented in this section are based on 2500 simulations. The analysis then took part in two main steps. In the first step, the goal is to identify which signals are the most relevant for the detection phase. This is done by comparing the rms of each of the signals during different phases of a simulation, which allows for the identification of which signals have large differences in their statistical properties and thus are likely the most relevant in the detection process. In the second step of the analysis, the goal is to determine potential threshold values for those signals. This is done by calculating what the accuracy of a model with various threshold values would be. The accuracy of a given model is calculated according to Equation (4.10), where 'TP' stands for 'true positive', 'FP' for 'false positive', and 'FN' for 'false negative'. However, the way in which TPs, FPs, FNs are calculated is different for the two conditions (DYN12 and DYN21).

$$\text{Accuracy} = \frac{\text{TP}}{\text{TP} + \text{FP} + \text{FN}} \quad (4.10)$$

An example schematic showing how they are calculated for the DYN12 condition is given in Figure 4.5. In this condition, whenever the instantaneous value of the signal exceeds the threshold, it counts as a trigger. If it occurs before the CE dynamics transition (at $t = 60$ sec), it is therefore a FP and if it occurs after the transition it is a TP. As an example, if the threshold is 2σ , a FP would occur at time t_1 while a TP would occur at time t_3 . On the other hand, if the threshold is 3σ , a FP would occur at time t_2 while a TP would occur at time t_4 . Note that a TP can only occur in the expected detection period, which is 2.2-7.4 seconds post-transition. If a TP is not detected in this time period, it results in a FN (therefore, the number of FNs is 2500 minus the number of TPs). Also note that the signals from the constant pilot are used for calculating TPs and FNs since the goal is to investigate at which point a HO would detect that a change in CE dynamics occurred, so before that point, the HO would be acting like the constant pilot.

For the DYN21 condition, the same analysis is not possible since the magnitude of the oscillations in each of the signals post-transition gets smaller for the constant pilot, instead of bigger as in the DYN12 condition. As a result, it was chosen to base the triggers on when the instantaneous value of a signal remains *below* the threshold for a certain amount of time, called the window size. A schematic for this is shown in Figure 4.6. Various window sizes were experimented with in the analysis, but suppose the window size is the time $t_2 - t_1$ and the threshold is 0.5σ . If that is the case, then a FP would occur at time t_2 and a TP would occur at time t_4 . However, if the window size became any larger, there would be no FP since the period $t_2 - t_1$ is the longest the signal stays below the threshold in the pre-transition phase,




Figure 4.5: Schematic showing how the second step of the analysis was done for the DYN12 condition.

but there would still be a TP since the signal stays below the threshold for longer than $t_4 - t_3$. On the other hand, if the threshold were to be 1σ , there would be both a FP and a TP regardless of the window size since the signal always stays below that threshold. As was the case with the DYN12 condition, if there is no TP in the expected detection period it results in a FN.

4.2.2. DYN12 Results

First, the results from the DYN12 condition will be presented. As was explained in Subsection 4.2.1, the first step is to compare the rms of each of the six signals (control input, output, error, and their rates) during different phases of a simulation. The results are shown in Figure 4.7. The leftmost boxplot in each sub-figure, labeled 'Pre-transition Steady-state', is the rms calculated from $30 \text{ sec} < t < 60 \text{ sec}$, so from the end of the run-in time until the transition. The middle boxplot, labeled 'Adaptive Pilot Detection Period', is the rms calculated from $62.2 \text{ sec} < t < 67.4 \text{ sec}$, which is the expected detection period. Finally, the rightmost boxplot, labeled 'Post-transition Steady-state', is the rms of the adaptive pilot signals calculated from $70 \text{ sec} < t < 90 \text{ sec}$, so from ten seconds after the transition until the end of the simulation, where the adaptive pilot has already (mostly) finished adapting to the new CE dynamics. In all three cases, one data point is generated from each simulation, so the boxplots show the spread in the rms over all 2500 simulations. Note that the constant pilot in the expected detection period is not shown because the system becomes unstable so the signals become multiple orders of magnitude larger than in the pre-transition phase.

Several observations can be made from Figure 4.7. First, it can be seen that the primary difference between the adaptive pilot during the expected detection period and in the steady-state post-transition phase is a significant decrease in the spread of the rms values, which is to be expected given that the transient behaviour introduces extra variability. The median also slightly decreases in all cases as any built-up errors in the transition phase are reduced to the steady-state tracking values. The spread of the rms values can also be seen to increase between the steady-state pre-transition phase and the steady-state post-transition phase in the majority of cases, due to the fact that the DYN2 condition is more difficult to control than the DYN1 condition. This leads to greater variability in y (Figure 4.7a), \dot{y} (Figure 4.7b), e (Figure 4.7e), and \dot{e} (Figure 4.7f). On the contrary, the spread in the rms values for u (Figure 4.7c) and \dot{u} decreases because the HO has to be less aggressive with the control inputs as a result of the system being closer to instability.




Figure 4.6: Schematic showing how the second step of the analysis was done for the DYN21 condition.

Finally, it can also be seen that there is a significant difference in the pre-transition steady-state rms and post-transition steady-state rms of the adaptive pilot (although not explicitly shown here, the same is true for the constant pilot since it becomes unstable). This is primarily visible for u (Figure 4.7c) and \dot{u} (Figure 4.7d) where the median steady-state post-transition rms is 77% and 76%, respectively, lower than the median pre-transition rms. Smaller, but still big, differences can also be seen in e (Figure 4.7e) and \dot{e} (Figure 4.7f). This leads to the first hypothesis, which is, **all participants will adapt their control strategy to avoid the same degradation in performance observed with the constant pilot, instead tending towards the adaptive pilot behaviour.**

Having identified that there is a significant difference in the rms of all six signals when comparing the steady-state pre-transition values with the constant pilot in the detection period (since it becomes unstable), that concludes the first step of the analysis. Since it is impossible to tell which of the signals is the first to diverge given that they are all in a closed loop together, the second step of the analysis will be done using e and \dot{e} since this has been the main focus in previous work and has been successful, as concluded in Section 3.3. The number of TPs, FPs, and FNs calculated according to the method from Figure 4.5 can be seen in Figure 4.8.

As can be seen, no matter how high the threshold is, there is a TP in 100% of the simulations for both e (Figure 4.8a) and \dot{e} (Figure 4.8b), and thus there are 0 FNs for both cases. Again, this is due to the system becoming unstable and the signals growing to several orders of magnitude larger than in the pre-transition phase. For thresholds below 2σ in e and 2.1σ in \dot{e} , a FP occurs in 100% of the simulations, which is undesired so the threshold must be higher than that. At 4σ for e and 4.4σ for \dot{e} , the number of FPs is reduced to 0, meaning the accuracy (defined in Equation (4.10)) is 100%. If 4σ were to be used as the threshold for both signals, the accuracy would be 100% and 99.9% for e and \dot{e} , respectively. This is very similar to the limit of 3.9σ proposed by van Ham [26] from her experiment results and thus suggests that the simulation is capable of predicting actual HO control behaviour. This leads to the second hypothesis, which is, **a significant increase in the standard deviation and maximum of the error and error rate signals will trigger the HO to detect a transition in the CE dynamics.**

4.2.3. DYN21 Results

Next, the results for the DYN21 condition will be presented. Again, the first step is to look at the rms of all six signals during different phases of a simulation. The results are presented in Figure 4.9. Note that in comparison to the DYN12 condition (Figure 4.7), there is an additional boxplot in each sub-figure, labeled 'Constant Pilot Detection Period'. This is the rms of the constant pilot signals calculated during the period $62.2 \text{ sec} < t < 67.4 \text{ sec}$, so it allows for comparison between the adaptive pilot and the constant pilot immediately after the transition.

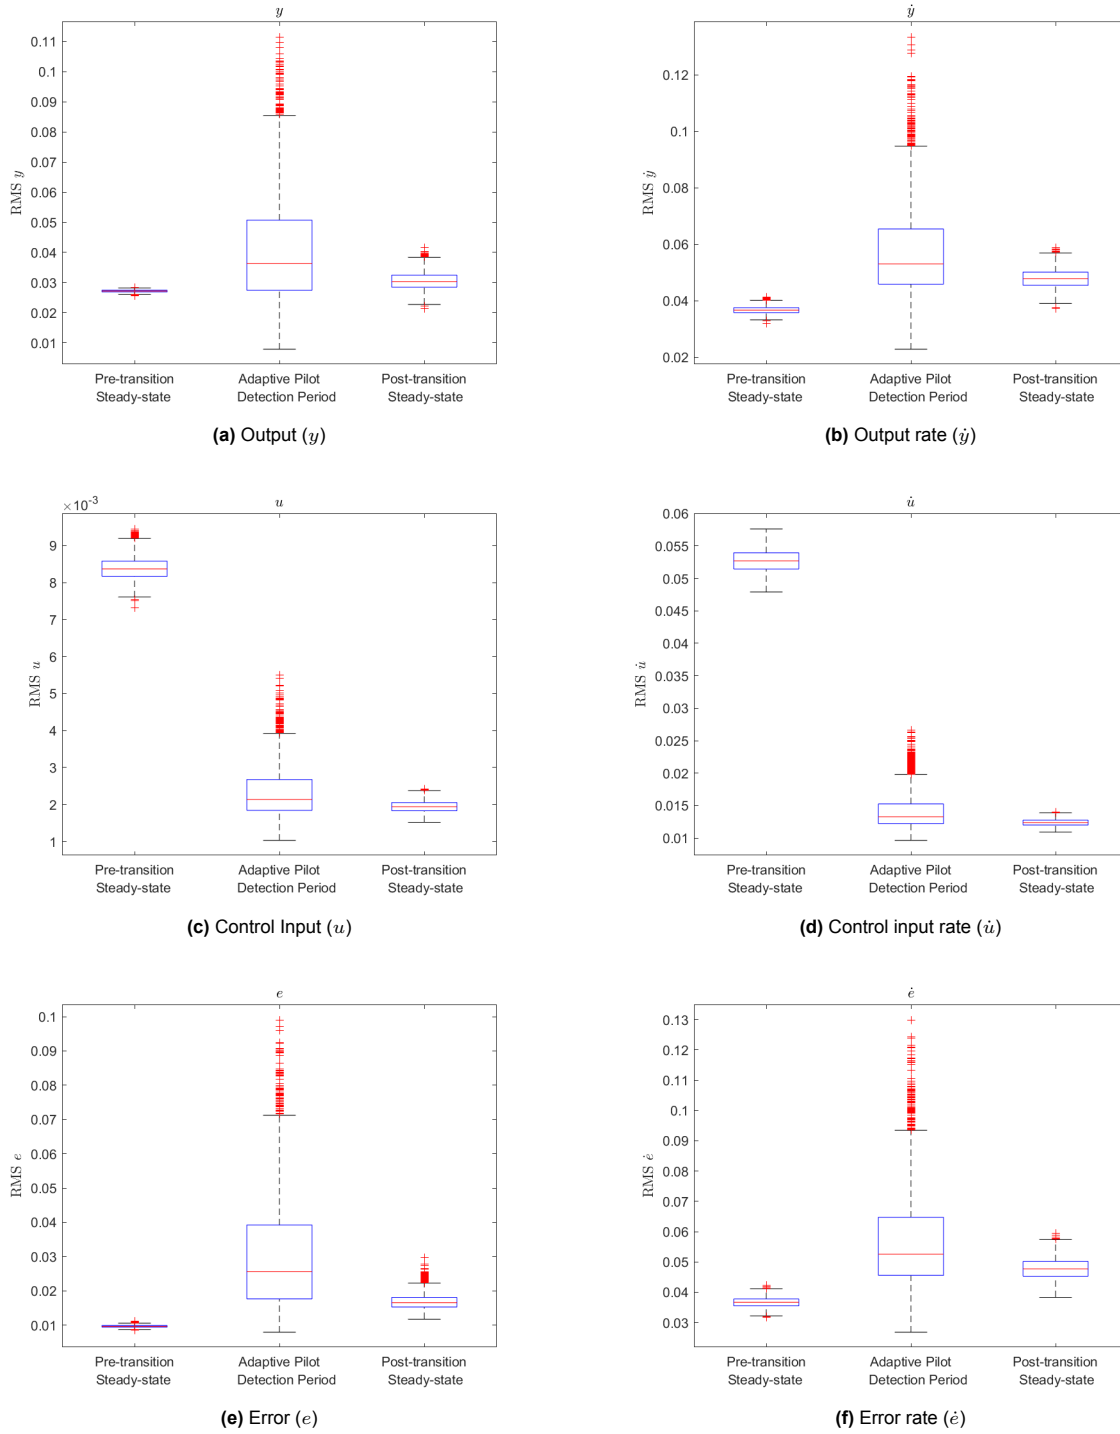


Figure 4.7: Comparison of the rms of each of the signals in the steady-state pre-transition phase with the adaptive pilot in the expected detection period, and the adaptive pilot in the steady-state post-transition phase for the DYN12 condition.

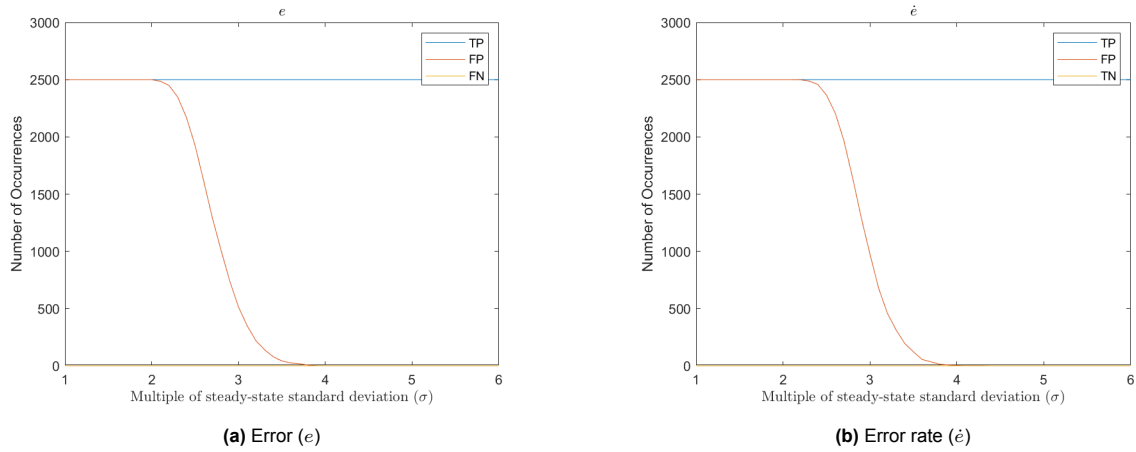


Figure 4.8: Number of TPs, FPs, and FNs for the constant pilot with various thresholds in the DYN12 condition.

The first thing that can be seen in Figure 4.9 is that there is once again a significant difference in the rms values between the constant pilot and the adaptive pilot in the detection region as well as between the two steady-state phases (pre- and post-transition). This reinforces the first hypothesis introduced in Subsection 4.2.2 which says, **"all participants will adapt their control strategy to avoid the same degradation in performance observed with the constant pilot, instead tending towards the adaptive pilot behaviour"**. Comparing the two steady-state phases, the reverse effect than what was observed in the DYN12 condition can be seen, namely that the spread in rms values decreases post-transition for y (Figure 4.9a), \dot{y} (Figure 4.7b), e (Figure 4.9e), and \dot{e} (Figure 4.9f) and increases post-transition for u (Figure 4.9c), and \dot{u} (Figure 4.9c). The cause of this is the nature of the CE dynamics, specifically that the DYN2 condition is more difficult to control, as was explained in Subsection 4.2.2.

The most significant differences between the pre-transition steady-state rms and constant pilot in the detection period occurs for \dot{y} , u , and \dot{e} . For these three signals, the median rms of the constant pilot is 89%, 36%, and 39% lower than the median pre-transition steady-state rms, respectively. Thus, these are the three signals that will be further analysed in the second step of the analysis. Before that, however, it is worth mentioning that there is a large spread in the rms values of e (Figure 4.9e) for the constant pilot in the detection period, and there is also some overlap with the pre-transition steady-state rms. This observation helps to explain why previous attempts at developing a model for the detection phase based on error for the DYN21 transition have been less successful. Figure 4.10 shows the spread in the rms of e for the constant pilot in the detection period over all 100 forcing function realizations. As can be seen, there are significant differences between forcing function realizations, which makes it very difficult to develop a model that performs well consistently. The spread in rms values for the constant pilot in the detection period for the three signals that will be further analysed in the next step is much lower, as seen in Figure 4.9.

For the second step of the analysis, in addition to which multiple of the steady-state standard deviation to use as the threshold, the window size is introduced as an additional variable. The window size determines how long a signal needs to stay below the threshold for it to cause a trigger and thus be classified as either a TP or a FP, as explained in Subsection 4.2.1. Therefore, pseudocolor plots were generated to determine which combination of threshold and window size results in the highest accuracy model (calculated according to Equation (4.10)). The results for \dot{y} , u , and \dot{e} are shown in Figure 4.11. Interestingly, despite large differences in rms values identified from Figure 4.9, no matter what combination of threshold and window size is used, it is impossible to achieve an accuracy close to 1 for u (Figure 4.11b) or \dot{e} (Figure 4.11c). This result helps support why previous models for the detection phase in the DYN21 condition based on \dot{e} do not perform well. It also suggests that modeling the detection phase for the DYN21 condition using a compensatory display can be troublesome since only the error, error rate, and control input are known to the HO. Given that e has already been determined not

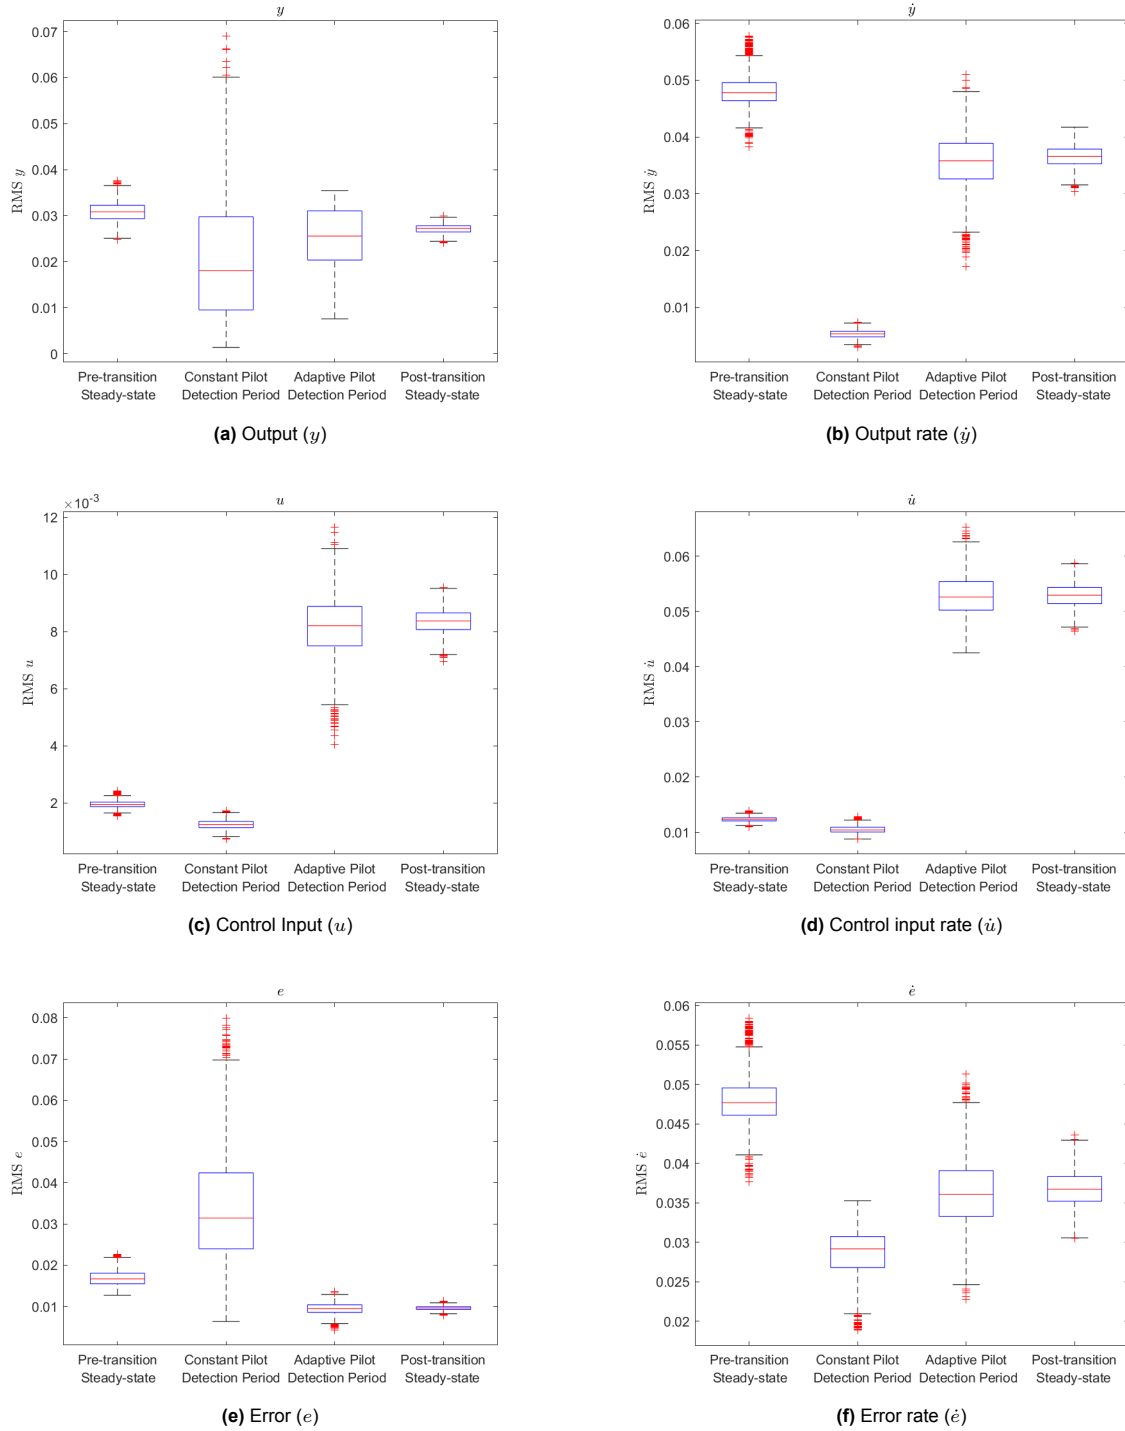


Figure 4.9: Comparison of the rms of each of the signals in the steady-state pre-transition phase with the constant pilot in the expected detection period, the adaptive pilot in the expected detection period, and the adaptive pilot in the steady-state post-transition phase for the DYN21 condition.




Figure 4.10: Spread in rms for the constant pilot in the detection region over all 100 forcing function realizations.

to be a good signal to look at, now that u and \dot{e} have also been shown not to be good signals to look at, it leaves no further signals on which to base a model or for the HO to use to determine when the CE dynamics change. On the other hand, there are several combinations of thresholds and window sizes towards the top left of the graph of \dot{y} in Figure 4.11a that result in an accuracy of 1. This means that with explicit knowledge of the output, which is the case when using a pursuit display, it may be easier to model the detection phase.

Choosing a threshold of 0.5σ for \dot{y} can be seen to result in an accuracy of 1 for many window sizes in Figure 4.11a. Thus, a vertical slice of the pseudocolor plot at 0.5σ was taken to see how the number of TPs, FPs, and FNs vary with different window sizes. The results are shown in Figure 4.12. As expected, the number of FPs falls off quickly for \dot{y} (Figure 4.12a) and the number of TPs stays at 2500, resulting in an accuracy of 1. If a window size of 3 seconds is taken, the accuracy would be 99.8% which leads to the third hypothesis, **the output rate staying below 0.5 standard deviations for 3 seconds is a trigger for a HO to detect a change in CE dynamics**. On the contrary, it can be seen that for u (Figure 4.12b) and \dot{e} (Figure 4.12c), the number of FPs and TPs reduce at a very similar rate, and thus it is not possible to achieve an accuracy over 0.5. The same is true for thresholds other than 0.5σ , which explains why the accuracy of u and \dot{e} never increases above 0.5 in Figure 4.11b and Figure 4.11c, respectively.

4.3. Conclusions

Overall, two main conclusions can be drawn from the simulation results. First, they were able to help develop hypotheses that could be used to optimize the experiment design. The three hypotheses discussed in Subsection 4.2.2 and Subsection 4.2.3 are:

1. Both conditions: All participants will adapt their control strategy to avoid the same degradation in performance observed with the constant pilot, instead tending towards the adaptive pilot behaviour.
2. DYN12: A significant increase in the standard deviation and maximum of the error and error rate signals will trigger the HO to detect a transition in the CE dynamics.
3. DYN21: The output rate staying below 0.5 standard deviations for 3 seconds is a trigger for a HO to detect a change in CE dynamics.

The first is motivated by the observation that there is a significant difference in the statistical properties

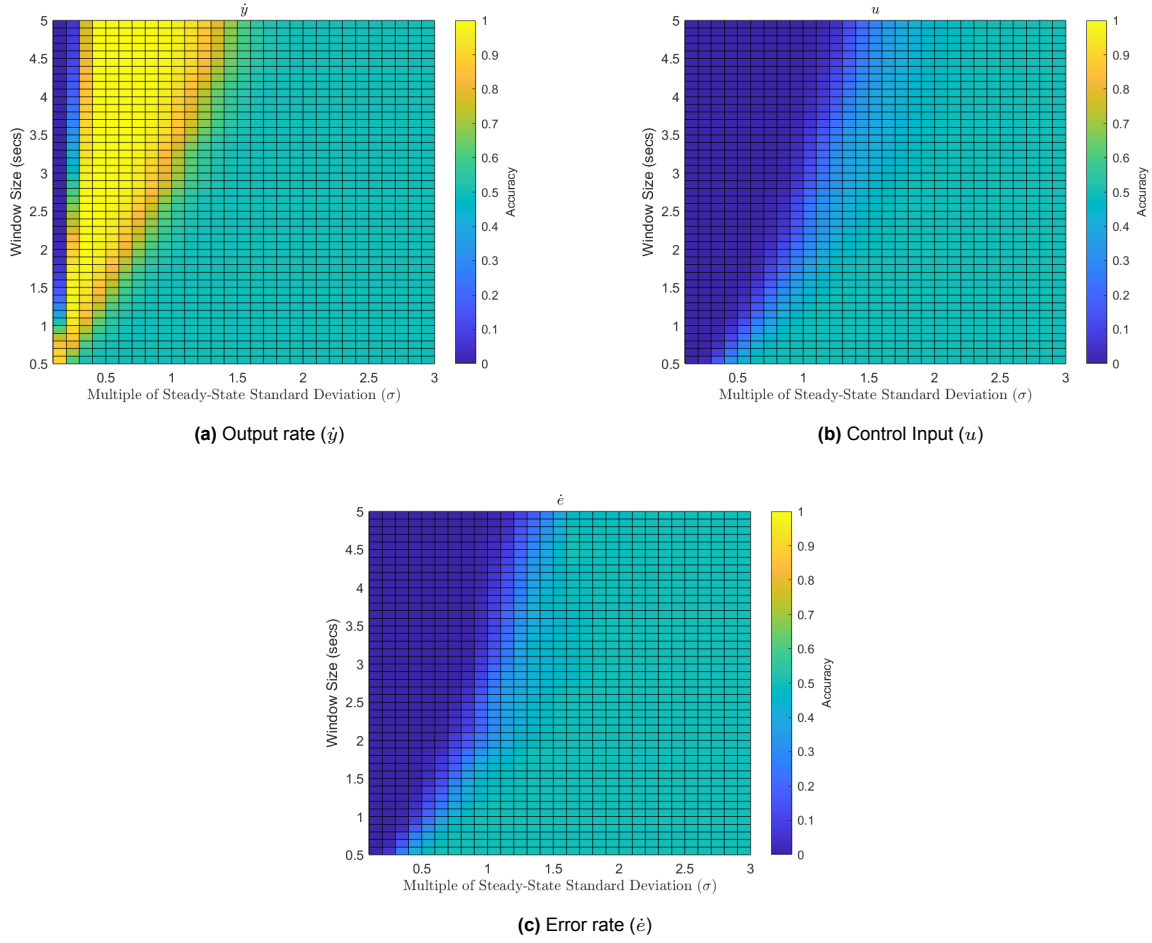


Figure 4.11: Accuracy plotted as a function of the threshold (multiple of steady-state standard deviation) and window length for the three signals that were identified to be most relevant for the detection phase.

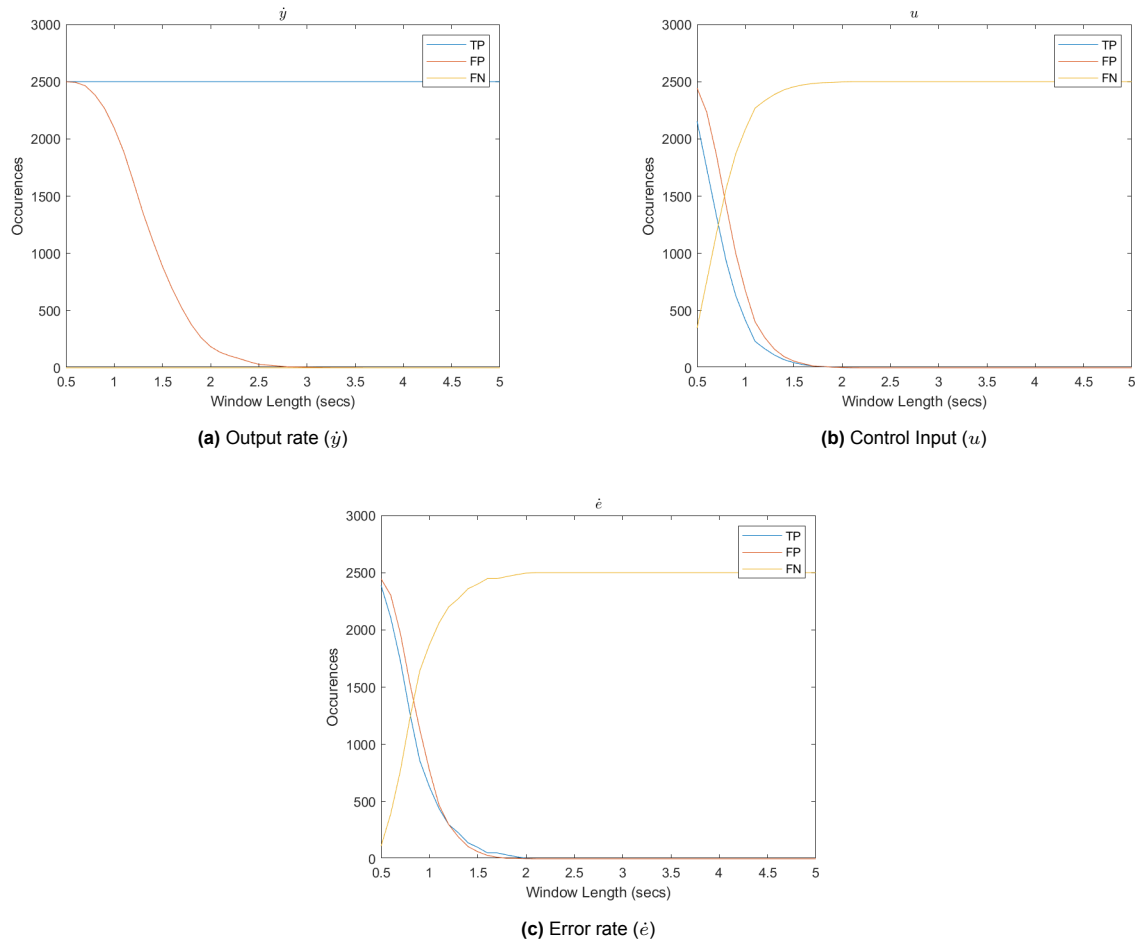


Figure 4.12: Number of TPs, FPs, and FNs over varying window sizes for the constant pilot in the DYN21 condition and a threshold of 0.5σ .

of at least some signals between the pre-transition steady-state phase and the post-transition steady-state phase. The second hypothesis is motivated by the fact that the constant pilot system becomes unstable post-transition forcing the HO to detect the transition, and the fact that previous research efforts have been successful in modeling the transition using the error and error rate signals. Finally, the third hypothesis is motivated by the fact that the output rate was the only signal that showed a reliable and significant difference in statistical properties when comparing the pre-transition steady-state phase with the constant pilot in the detection period, and for which it was possible to find a combination of threshold and window size for which the accuracy is close to 100%.

The second main conclusion that can be drawn is that the simulation results help explain why previous research efforts focusing on the error and error rate signals to model the detection phase in the DYN21 condition have been less successful. It was found that the spread of rms values of the constant pilot's error signal during the detection period is very large, overlaps with the rms values of the steady-state pre-transition phase, and varies significantly for different forcing function realizations. For the control input and error rate, it was impossible to find a combination of threshold and window size for which the accuracy was close to 1. Since these three signals are the only ones known to a HO when using a compensatory display, it helps highlight the difficulty of making a model for the DYN21 transition. However, when using a pursuit display, the output rate is explicitly known to the HO, which will likely make the modeling process easier, as suggested by the third hypothesis.

5

Future Work

Having concluded the literature review and analysis of the simulation results, this chapter will outline the future work that will be done. First, the experiment design will be introduced in Section 5.1, and then the analysis to be done on the experiment data will be discussed in Section 5.2.

5.1. Experiment Design

In this section, an answer to the fourth sub-research question ("how can the experiment best be designed to help develop and validate the model?") will be provided by explaining and motivating the experiment design. This includes an introduction to the control task in Subsection 5.1.1, a discussion on the CE dynamics in Subsection 5.1.2, the choice of forcing functions in Subsection 5.1.3, the choice of participants in Subsection 5.1.4, and the hypotheses that will be tested during the experiment in Subsection 5.1.5.

5.1.1. Control Task

The experiment will involve participants conducting a pitch-tracking task using a pursuit display. A sketch of this is given in Figure 5.1. Since it is a pursuit display, participants will be able to explicitly see their own output (i.e. the controlled element output) and the forcing function, which is the target signal. Their goal will be to minimize the error throughout the runs. Each run will last between 95 and 105 seconds. Data will be collected for exactly 90 seconds in each run, but the run length can vary because the run-in time will be randomly chosen between 5 and 15 seconds to avoid predictability, since the CE dynamics transition will always occur 30 seconds after the end of the run-in time. The 90 second run length allows there to be exactly three periods of the forcing function. The first period can be used to model steady-state tracking behaviour for the first condition, the second period serves as the transition period, and the final 30 seconds can be used to model steady-state tracking behaviour for the second condition. If participants notice a change in the CE dynamics, they will be instructed to immediately press a button on a different side-stick. It was chosen to use a different side-stick for the button press because van Ham [26] concluded that having to do two tasks with the same hand (keeping the tracking error low and pressing the button) led to decreased performance in one of the two tasks. To avoid this, the button will be on a different side-stick which participants can press with their other hand.

5.1.2. CE Dynamics

The two CE dynamics will be the same ones that were used in the simulation. For convenience, they are repeated here. The DYN1 condition is given in Equation (5.1) and the DYN2 condition is given in Equation (5.2). The break frequency (which is time-varying) will vary according to the sigmoid function given in Equation (5.3). ω_{b_1} is the break frequency of the first condition, ω_{b_2} is the break frequency of the second condition, G is the maximum rate of change and M is the moment in time where the maximum rate of change occurs. G will be set to 100 sec^{-1} to simulate a step-like change, representative of a failure of the SAS as explained in Subsection 4.1.1. M will be set to 30 sec (the run-in time counts as negative time so this will always be 30 seconds after the end of the run-in time).




Figure 5.1: Sketch of the pursuit display to be used in the experiment.

$$H_{CE_1} = \frac{90}{s(s + 20)} \quad (5.1)$$

$$H_{CE_2} = \frac{90}{s(s + 0.2)} \quad (5.2)$$

$$\omega_b(t) = \omega_{b_1} + \frac{\omega_{b_2} - \omega_{b_1}}{1 + e^{-G(t-M)}} \quad (5.3)$$

5.1.3. Forcing Functions

The properties of the forcing function in the transition region have been shown by Jakimovska [44] and van Ham [26] to significantly impact how long it takes HOs to detect a transition. Specifically, larger gradients in the forcing function in the transition region generally lead to quicker detection times because they require larger control inputs, thus exciting the system more and making a change in the CE dynamics more noticeable. To avoid biasing the results due to the choice of forcing function, a total of nine different forcing functions will be used in the experiment. Out of the nine, six will be used for developing the model, and three will be kept for validating the model. The nine forcing functions were chosen such that three of them had low power in the two-second window centered around the transition, three of them had medium power in the same window, and three of them had high power. Further details are given in Appendix A.

For the six forcing functions that will be used to generate training data, one run will be performed in each of DYN1, DYN12, DYN2, and DYN21 (so four runs per forcing function). For the three forcing functions that will be used for validation, one run will be performed in both DYN12 and DYN21 (so two runs per forcing function). Thus, there will be a total of 30 runs performed in the experiment, excluding training runs where the participants can familiarize themselves with the tracking task and the two types of CE dynamics.

The experiment will be split into two, whereby the fifteen runs in either DYN1 or DYN12 will be performed sequentially, and the fifteen runs in either DYN2 or DYN21 will be performed sequentially. Within the two halves, the order of the 15 runs will be determined by an incomplete Latin square (one for the DYN1 and DYN12 conditions, and a separate one for the DYN2 and DYN21 conditions). As a result, to avoid any order effects, a multiple of fifteen participants will be required.

5.1.4. Participants

The choice of participants is very important, as it is critical that they maintain a high level of concentration and attempt to keep the tracking error as low as possible for the duration of the experiment. In Terenzi's research [37] this was not the case, with the majority of participants having a crossover frequency lower than 1 rad/sec. In turn, this made it difficult to draw conclusions based on the data and generalize the findings. To avoid a repetition of this, participants who are known to be good at tracking tasks will be invited to the experiment, and if their performance is significantly worse than the rest of the participants, their results will be excluded and a replacement participant will be found.

5.1.5. Hypotheses

A total of four hypotheses have been developed for the experiment, as listed below.

1. Both conditions: All participants will adapt their control strategy to avoid the same degradation in performance observed with the constant pilot, instead tending towards the adaptive pilot behaviour.
2. DYN12: A significant increase in the standard deviation and maximum of the error and error rate signals will trigger the HO to detect a transition in the CE dynamics.
3. DYN21: The output rate staying below 0.5 standard deviations for 3 seconds is a trigger for a HO to detect a change in CE dynamics.
4. Both conditions: Participants will be quicker to detect transitions in the CE dynamics when the gradient of the forcing function around the moment of transition is larger.

The first three have already been introduced and motivated in Chapter 4. The fourth and final one is motivated by the findings in Jakimovska [44] and van Ham's [26] research, as briefly discussed in Subsection 5.1.3.

5.2. Analysis to be Done


After completing the experiment, the final two steps are developing the model and validating it. For developing the model, a similar analysis to what was done with the simulation results in Section 4.2 will be performed. The goal will be to find a relationship between when participants noticed a change in the CE dynamics and a specific property in one or more of the signals, such that the model triggers for both transitions in the CE dynamics and has similar detection lags to the participants. For this, the moment participants pressed the button to indicate they noticed a change in the CE dynamics will be used to help narrow down where to look for said properties in the six signals that will be analyzed. The model with the highest accuracy will then be further investigated to ensure its robustness with a sensitivity analysis. Finally, the model will be validated using the data collected in the experiment reserved for validation, as explained in Subsection 5.1.3. Once this is complete, the model can be used for future research on the adaptive behaviour of HOs.

References

- [1] Mulder, M. et al. *Manual Control Cybernetics: State-of-the-Art and Current Trends*. Oct. 2018. DOI: 10.1109/THMS.2017.2761342.
- [2] *Loss of Control In-flight (LOC-I)*. URL: <https://www.iata.org/en/programs/safety/operational-safety/loss-of-control-inflight/#:~:text=LOC%2DI%20refers%20to%20accidents,failures%2C%20icing%2C%20or%20stalls..>
- [3] Drop, F. "Control-Theoretic Models of Feedforward in Manual Control". In: (2016). DOI: 10.4233/uuid:7c1f62db-9a5a-4e02-8f11-488d6a299500. URL: <https://doi.org/10.4233/uuid:7c1f62db-9a5a-4e02-8f11-488d6a299500>.
- [4] Young, L. R. et al. "Adaptive Dynamic Response Characteristics of the Human Operator in Simple Manual Control". In: *IEEE Transactions on Human Factors in Electronics* HFE-5.1 (1964), pp. 6–13. ISSN: 21682852. DOI: 10.1109/THFE.1964.231648.
- [5] Friston, K. "What is optimal about motor control?" In: *Neuron* 72.3 (Nov. 2011), pp. 488–498. ISSN: 08966273. DOI: 10.1016/j.neuron.2011.10.018.
- [6] Bastian, A. J. "Moving, sensing and learning with cerebellar damage". In: *Current Opinion in Neurobiology* 21.4 (Aug. 2011), pp. 596–601. DOI: 10.1016/j.conb.2011.06.007.
- [7] Kawato, M. "Internal Models for Motor Control and Trajectory Planning". In: *Current Opinion in Neurobiology* 6.5 (1999), pp. 718–727. URL: <http://www.erato.atr.co.jp/DB/>.
- [8] Miall, R. C. and Wolpert, D. M. "Forward Models for Physiological Motor Control". In: *Neural Networks* 9.8 (1996), pp. 126–1279.
- [9] Wolpert, D. M., Miall, R. C., and Kawato, M. "Internal models in the cerebellum". In: *Trends in Cognitive Sciences* 2.9 (1998), pp. 338–347. DOI: 10.1016/S1364-6613(98)01221-2.
- [10] Shadmehr, R. and Mussa-Ivaldi, F. A. "Adaptive Representation of Dynamics during Learning of a Motor Task". In: *Journal of Neuroscience* 14.5 (1994), pp. 3208–3224.
- [11] Lackner, J. R. and Dizio, P. "Gravitoinertial Force Background Level Affects Adaptation to Coriolis Force Perturbations of Reaching Movements". In: *Journal of Neurophysiology* 80.2 (1998), pp. 546–553.
- [12] Lackner, J. R. and Dizio, P. "Rapid adaptation to Coriolis force perturbations of arm trajectory". In: *Journal of Neurophysiology* 72.1 (1994), pp. 299–313.
- [13] Stoodley, C. J. *The Cerebellum and Neurodevelopmental Disorders*. Feb. 2016. DOI: 10.1007/s12311-015-0715-3.
- [14] Wolpert, D., Ghahramani, Z., and Jordan, M. I. "An Internal Model for Sensorimotor Integration". In: *Science* 269 (1995), pp. 1880–1882. URL: www.sciencemag.org.
- [15] Davidson, P. R. et al. "Detecting adaptive inverse models in the central nervous system". In: *Annual Reports of the Research Reactor Institute, Kyoto University*. Vol. 1. 2001, pp. 853–856. DOI: 10.1109/iembs.2001.1019076.
- [16] Imamizu, H. et al. "Human cerebellar activity reflecting an acquired internal model of new tool". In: *Nature* 403 (2000), pp. 192–195.
- [17] Doya, K., Kimura, H., and Miyamura, A. "Motor Control: Neural Model And Systems Theory". In: *International Journal of Applied Mathematics and Computer Science* 11.1 (2001), pp. 77–104.
- [18] Neilson, P. D. and Neilson, M. D. "Adaptive Model Theory: Modelling the Modeller". In: *Seminal Contributions to Modelling and Simulation*. 2016, pp. 19–22. DOI: 10.1007/978-3-319-33786-9{_}4.
- [19] Neilson, P. D., Neilson, M. D., and O'Dwyer, N. J. "Adaptive Optimal Control of Human Tracking". In: *Motor Control and Sensory Motor Integration: Issues and Directions* (2011), pp. 97–140.

- [20] Mulder, M. et al. "Neuroscience Perspectives on Adaptive Manual Control with Pursuit Displays". In: *IFAC-PapersOnLine* 55.29 (2022), pp. 160–165. DOI: 10.1016/j.ifacol.2022.10.249.
- [21] McRuer, D. T. and Jex, H. R. "A Review of Quasi-Linear Pilot Models". In: *IEEE Transactions on Human Factors in Electronics* HFE-3.3 (1967), pp. 231–249. DOI: 10.1109/THFE.1967.234304.
- [22] Young, L. R. and Stark, L. *Biological control system - a critical review and evaluation, developments in manual control*. Tech. rep. In collab. with NASA, Mar. 1965. URL: https://archive.org/details/nasa_techdoc_19650009660/page/n107/mode/2up.
- [23] Elkind, J. I. and Miller, D. C. "Process of Adaptation by the Human Controller". In: *Second Annual NASA University Conference on Manual Control* (1966), pp. 47–63.
- [24] Young, L. R. "On Adaptive Manual Control". In: *Ergonomics* 12.4 (July 1969), pp. 635–674. ISSN: 13665847. DOI: 10.1080/00140136908931083.
- [25] Phatak, A. V. and Bekey, G. A. "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 (1969), pp. 72–80. DOI: 10.1109/TMMS.1969.299886.
- [26] Van Ham, J. M. *Adaptive Manual Control - The Human Response to Sudden Changes in Controlled Element Dynamics (MSc thesis)*. 2021. URL: <https://surfdriive.surf.nl/files/index.php/s/p3sY6Ur0IFjsj1X>.
- [27] Thorpe, S., Fize, D., and Marlot, C. "Speed of processing in the human visual system". In: *Nature* 381 (1996), pp. 520–522.
- [28] Niemela, R. J. and Krendel, E. S. "Detection of a Change in Plant Dynamics in a Man-Machine System". In: *Correspondence* (1975).
- [29] Hess, R. A. "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.
- [30] Hess, R. A. "Modeling pilot control behavior with sudden changes in vehicle dynamics". In: *Journal of Aircraft* 46.5 (2009), pp. 1584–1592. ISSN: 15333868. DOI: 10.2514/1.41215.
- [31] Hess, R. A. "Modeling Pilot Detection of Time-Varying Aircraft Dynamics". In: *Journal of Aircraft* 49.6 (2012), pp. 2100–2104. DOI: 10.2514/1.C031805.
- [32] Proske, U. and Gandevia, S. C. "The Proprioceptive Senses: Their Roles in Signaling Body Shape, Body Position and Movement, and Muscle Force". In: *Physiol Rev* 92 (2012), pp. 1651–1697. DOI: 10.1152/physrev.00048.2011.-This. URL: www.prv.org.
- [33] Hess, R. A. "A Rationale for Human Operator Pulsive Control Behavior". In: *Journal of Guidance and Control* 2.3 (May 1979), pp. 221–227. ISSN: 0162-3192. DOI: 10.2514/3.55864.
- [34] Jakimovska, N. et al. "Using the Hess Adaptive Pilot Model for Modeling Human Operator's Control Adaptations in Pursuit Tracking". In: American Institute of Aeronautics and Astronautics (AIAA), Jan. 2023. DOI: 10.2514/6.2023-0541.
- [35] Plaetinck, W. *Online Time-Varying Identification and Detection of Operator Adaptation with Recursive ARX (MSc thesis)*. 2018. URL: <https://surfdriive.surf.nl/files/index.php/s/4iazNZmhttZEeBK>.
- [36] Jiao, J. et al. "Identifying Pilot Control Adaptations to Sudden Changes in Aircraft Dynamics". In: *Journal of Guidance, Control, and Dynamics* (Feb. 2023), pp. 1–8. ISSN: 0731-5090. DOI: 10.2514/1.g007358.
- [37] Terenzi, L. *Adaptive Motor Control A Predictive Coding Approach (MSc Thesis)*. 2021. URL: <https://surfdriive.surf.nl/files/index.php/s/ha0eHsh1IrucAgP>.
- [38] Tohidi, S. S. and Yildiz, Y. "Adaptive human pilot model for uncertain systems". In: *2019 18th European Control Conference, ECC 2019*. Institute of Electrical and Electronics Engineers Inc., June 2019, pp. 2938–2943. ISBN: 9783907144008. DOI: 10.23919/ECC.2019.8795847.
- [39] Tohidi, S. S. and Yildiz, Y. "A Control Theoretical Adaptive Human Pilot Model: Theory and Experimental Validation". In: *IEEE Transactions on Control Systems Technology* 30.6 (Nov. 2022), pp. 2585–2597. ISSN: 15580865. DOI: 10.1109/TCST.2022.3164237.

- [40] Zaal, P. M.T. “Manual Control Adaptation to Changing Vehicle Dynamics in Roll–Pitch Control Tasks”. In: *Journal of Guidance, Control, and Dynamics* 39.5 (May 2016), pp. 1046–1058. ISSN: 0731-5090. DOI: 10.2514/1.G001592.
- [41] Zollner, H. G. H. et al. “The Effects of Controlled Element Break Frequency on Pilot Dynamics During Compensatory Target-Following”. In: *AIAA Modeling and Simulation Technologies Conference 2010*. 2010. ISBN: 9781624101526. DOI: 10.2514/6.2010-8092.
- [42] Van Grootheest, H. A. *Human-Operator Identification with Time-Varying ARX Models (MSc Thesis)*. 2017. URL: <https://surfdrive.surf.nl/files/index.php/s/HZdqnoKHMisFSL>.
- [43] Van Der El, K., Pool, D. M., and Mulder, M. “Analysis of Human Remnant in Pursuit and Preview Tracking Tasks”. In: *IFAC-PapersOnLine*. Vol. 52. 19. Elsevier B.V., 2019, pp. 145–150. DOI: 10.1016/j.ifacol.2019.12.165.
- [44] Jakimovska, N. *Detection of Human Adaptation to Sudden Changes in Controlled Element Dynamics (MSc thesis)*. 2022. URL: <https://surfdrive.surf.nl/files/index.php/s/N7tCfo4HTByB89h>.



Experiment Forcing Functions

As was explained in Subsection 5.1.3, the nine forcing functions to be used in the experiment were chosen such that three had low power in the two-second window centered around the transition, three had medium power in the same window, and three had high power. The methodology for choosing the forcing functions will be described below.

The first step was to determine the distribution of the power in a two-second window. For this, the power in every two-second window of 1,000 forcing functions, each having an average crest factor (as defined by Equation (4.3), and briefly outlined in Subsection 4.1.1), was calculated. The resulting distribution is given in Figure A.1¹. A power below $2.5 \cdot 10^{-5} \text{ rad}^2$ (to the left of the leftmost red line in Figure A.1) is considered low power, a power between the two red lines is considered medium power, and a power above $7.5 \cdot 10^{-5}$ (to the right of the rightmost red line) is considered high power.

Having determined the distribution, the next step is to choose a set of nine forcing functions out of the 1,000 that were generated. A good forcing function for the experiment, in addition to having an average crest factor, should be roughly normally distributed (a true normal distribution is impossible to achieve with a sum of sine waves due to the nature of a sine wave). Thus, nine forcing functions which are roughly normally distributed were chosen from the set of 1,000. As an example of what is considered a good distribution for the forcing function, the Q-Q plot of one of the forcing functions that will be used in the experiment is given in Figure A.2.

Finally, the chosen forcing functions were shifted in time such that there would be three forcing functions in each of the power levels (low, medium, and high) for the two-second window centered around the transition. The forcing functions were then numbered in order of power in that two-second window (1 being the lowest, 9 being the highest) and the middle power in each of the three power levels was chosen as the validation forcing function (see Subsection 5.1.3), so forcing function numbers 2, 5, and 8.

The phase shifts (in radians) for the nine forcing functions to be used in the experiment are given in Table A.1 along with the power in the two-second window centered around the transition² (note the subscript on ϕ indicates the sine wave it corresponds to, with 1 being the one with the largest amplitude and 10 being the one with the smallest amplitude).

¹Note a time step of 0.01 was used to generate Figure A.1, so the two-second windows were 0-2 secs, 0.01-2.01 secs, 0.02-2.02 secs, ..., 27.99-29.99 secs, and 28-30 secs.

²Note the amplitudes of the ten sine waves are always the same and are given in Table 4.1.




Figure A.1: Distribution of power in the forcing function over a two-second window.




Figure A.2: Q-Q plot of a forcing function that will be used in the experiment.

Table A.1: Phase shifts and power in the two-second window centered around the transition for the nine forcing functions to be used in the experiment.

	Forcing Function Number								
	1	2	3	4	5	6	7	8	9
ϕ_1	6.065	3.454	0.885	4.315	5.608	3.671	1.040	3.219	6.172
ϕ_2	0.223	3.073	1.120	3.308	6.027	2.757	1.958	5.615	2.662
ϕ_3	4.232	5.440	4.806	0.988	3.106	5.220	2.037	4.049	0.903
ϕ_4	1.880	1.296	4.251	3.925	3.109	4.162	4.768	4.820	2.290
ϕ_5	4.757	4.062	1.260	6.259	5.506	3.685	1.522	4.160	6.112
ϕ_6	1.072	3.464	4.764	6.236	5.560	5.188	2.820	0.440	2.162
ϕ_7	2.754	0.153	5.692	0.609	5.668	5.261	1.520	2.525	4.821
ϕ_8	1.593	1.678	3.003	4.782	2.511	5.699	1.577	3.368	0.101
ϕ_9	5.404	4.585	0.116	2.776	1.259	5.566	0.114	0.213	6.090
ϕ_{10}	4.101	0.583	1.338	4.126	2.843	0.776	0.755	1.461	1.853
Power (rad ²)	$6.2 \cdot 10^{-6}$	$1.2 \cdot 10^{-5}$	$1.6 \cdot 10^{-5}$	$4.4 \cdot 10^{-5}$	$5.8 \cdot 10^{-5}$	$6.3 \cdot 10^{-5}$	$1.3 \cdot 10^{-4}$	$2.2 \cdot 10^{-4}$	$2.3 \cdot 10^{-4}$

B

Padé Approximation in Time-Varying Simulation

Ultimately, the time-varying simulation used in this report was made using MATLAB and Simulink, as described in Section 4.1. However, an earlier version of the simulation was made using only MATLAB and had one crucial difference compared to the final simulation, namely that it used a Padé approximation for the HO time delay. A Padé approximation linearises the exponential term, the HO time delay, and represents it as a transfer function. Using the 'Variable time delay' block in Simulink, there is an option to linearise it using a Padé approximation but it is not required and can also be used as a pure time delay, which is what was ultimately done for the final simulation. The reason for this is that the use of the Padé approximation caused significant problems in modeling the quickly time-varying CE dynamics used in this report, as will be illustrated below.

In the steady-state conditions, the use of a third order Padé approximation led to very similar results compared to using the 'variable time delay' block in Simulink, as can be seen in Figure B.1-B.2. In Figure B.1, it can be seen that the output of the two models very closely matches, to the point that the two lines are indistinguishable. The difference between the two models is shown in Figure B.2, where it can be seen that it is in the order of 10^{-6} . Thus, for time-invariant simulations, a third order Padé approximation works very well. Note that Figure B.1-B.2 show the results for the DYN1 condition, but similar results are obtained for the DYN2 condition so will not be repeated.

However, in the time-varying simulations, particularly the DYN12 transition, the Padé approximation introduces significant and long-lasting transients at the moment the transition occurs ($t = 60$ sec), as can be seen in Figure B.3. After approximately 30-40 seconds, the transients die out, and steady-state tracking resumes. However, during that time, the control input and output grow to extremely high numbers and do not correspond to each other. The control input is negative during the entire transient, while the output remains positive the entire time, which is not possible. Keeping a constant time delay, for example the time delay corresponding to the DYN1 condition for the DYN12 transition, significantly reduces the magnitude of the oscillations in the control input and the output, as can be seen in Figure B.4. However, the oscillations in the control input and the length of time it takes for the transients to die out make it impossible to perform any meaningful analysis.

In the DYN21 condition, Figure B.5, the transient is much smaller, albeit still present.

In the final model, there are still some transients at the moment the transition occurs, as can be seen in Figure B.6. However, the oscillations in the control input are significantly less aggressive and always die out within one second, usually less than that. Furthermore, the effect on the output is much less noticeable as the transients primarily impact the control input. As a result of this, it was determined that the Simulink model is good enough to use for analysis and the use of Padé approximations in quickly time-varying systems, at least with this implementation, is not appropriate.




Figure B.1: Comparison of the output of the MATLAB only model and the model using Simulink in the DYN1 condition.




Figure B.2: Difference between the outputs of the two models for the DYN1 condition.




Figure B.3: Transient in the control input and output as a result of using a third order Padé approximation in the DYN12 condition.




Figure B.4: Transient in the control input and output when the time delay is kept constant in the DYN12 condition.




Figure B.5: Transient in the control input and output in the DYN21 condition.





Figure B.6: Control input and output in the transition region for the final model used in this report.

Part III

Scientific Article Appendices



Full Simulation Analysis Results

The full results for the first step of the analysis (comparing the standard deviations over the different phases of the simulation) in the DYN12 condition are given in Figure A.1. Note that the constant HO in the detection period is not shown because the system with the constant HO becomes unstable. The accuracy over a range of thresholds for a model based on each of the signals (step two of the analysis) is given in Figure A.2. Note how it is possible to achieve an accuracy close to 100% for each signal. This is a result of the system with the constant HO becoming unstable post-transition, thus causing all of the signals to grow to abnormally large values.

The full results for the first step of the analysis in the DYN21 condition are given in Figure A.3.

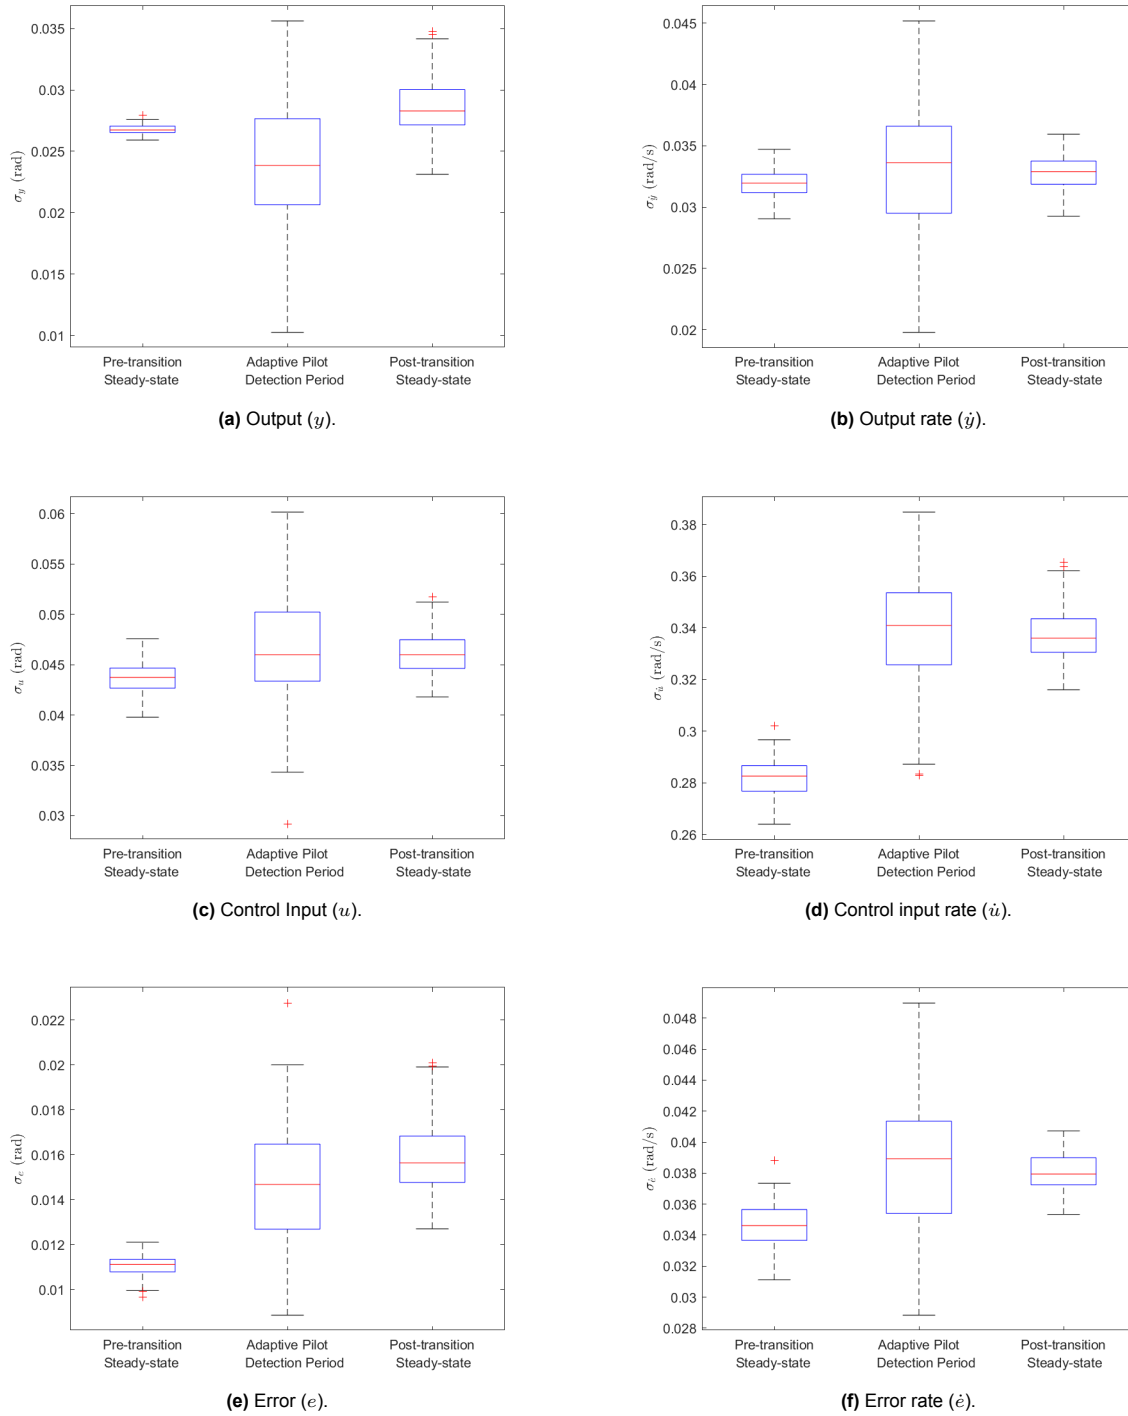


Figure A.1: Comparison of the standard deviations of each of the signals in the steady-state pre-transition phase with the adaptive HO in the detection period and the adaptive HO in the steady-state post-transition phase for the DYN12 condition. $N = 135$ in all of the box plots.




Figure A.2: Accuracy of a model based on each of the candidate signals in the DYN12 condition.

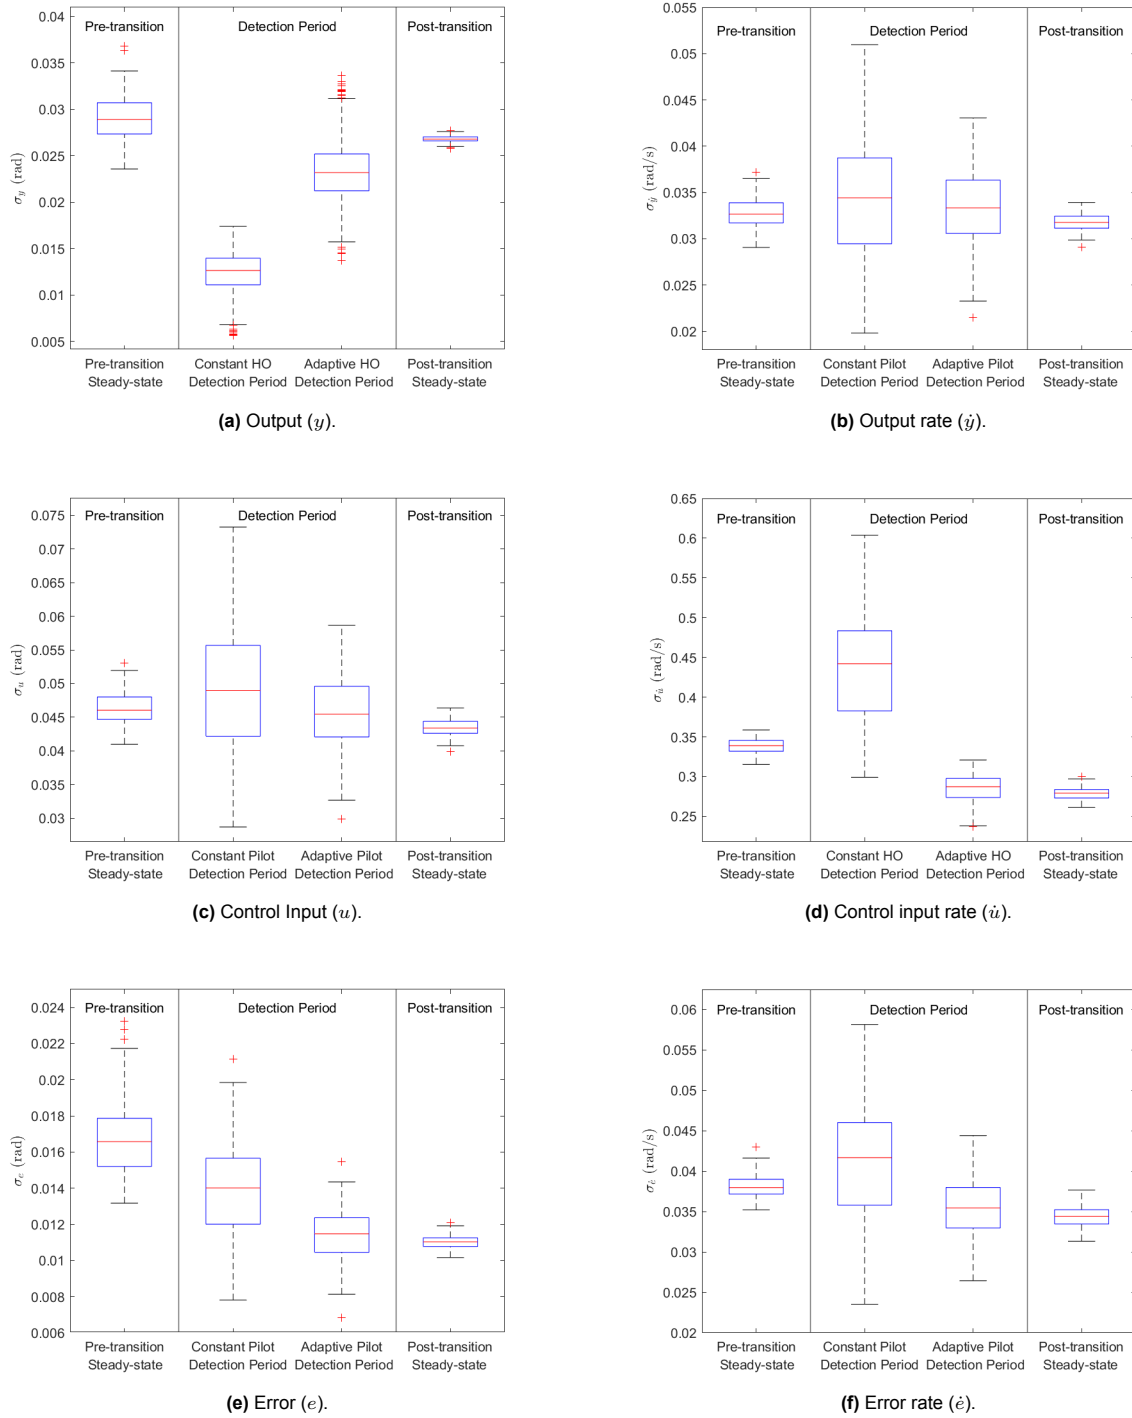


Figure A.3: Comparison of the standard deviations of each of the signals in the steady-state pre-transition phase with the constant HO in the detection period, the adaptive HO in the detection period, and the adaptive HO in the steady-state post-transition phase for the DYN21 condition. $N = 135$ in all of the box plots.

B

Participant 8 Data

The data (performance metrics) for Participant 8, whose data were not included in the results presented in the scientific article, are presented here. In each figure, the mean from the other fifteen participants is presented as a reference. As can be seen, the performance metrics for this participant indicate that they were not at the level of tracking or detection performance of the other participants, particularly in the DYN2 and DYN21 trials. Thus, the decision was made to find a replacement participant.




Figure B.1: RMSE over the six time-invariant trials in DYN1 and DYN2.




Figure B.2: Crossover frequencies calculated using the first 30 s of the measurement time for the six time-invariant trials in DYN1 and DYN2.




Figure B.3: Detection lags for the six DYN12 and DYN21 trials that did not use the validation forcing functions.

C

Performance Metrics Sorted by Forcing Function and Order in Experiment

The performance metrics sorted by forcing function (except the detection lags since they were presented in Part I) and by order in which the conditions were performed in the experiment will be presented in this chapter. In all figures except Figure C.5, there is one data point per participant in each box plot, and thus, $N = 15$. Due to a FP or FN, $N = 14$ in the following box plots:

- Trial number 1 in Figure C.5a
- Trial number 1 in Figure C.5b
- Trial number 8 in Figure C.5b
- Trial number 9 in Figure C.5b

Overall, it can be seen in all of the figures presented in this chapter that the forcing function and the order in which the conditions were performed did not have a significant impact on the performance metrics. For the latter, it can therefore be concluded that learning effects were kept to a minimum in this experiment.




Figure C.1: RMSE sorted by forcing function in the DYN1 and DYN2 trials.

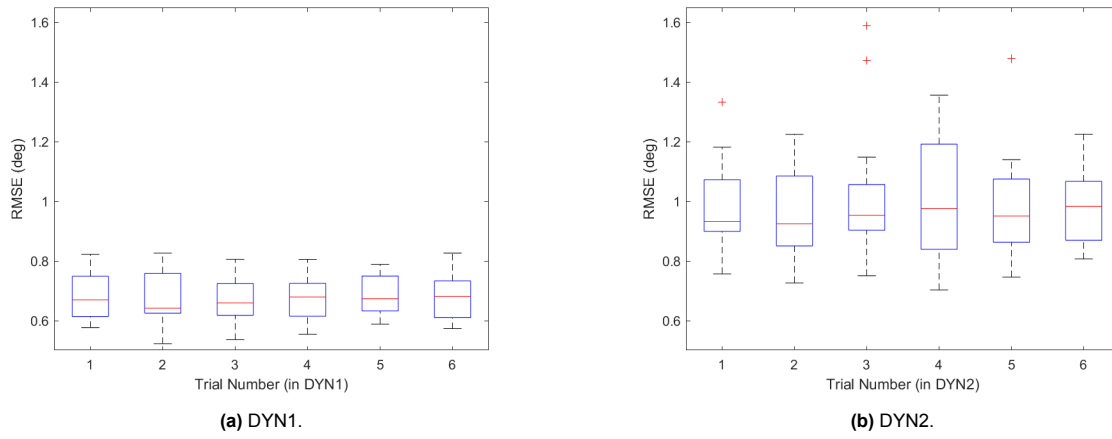


Figure C.2: RMSE sorted by the order in which the DYN1 and DYN2 trials were performed.




Figure C.3: Crossover frequencies sorted by forcing function in the DYN1 and DYN2 trials.

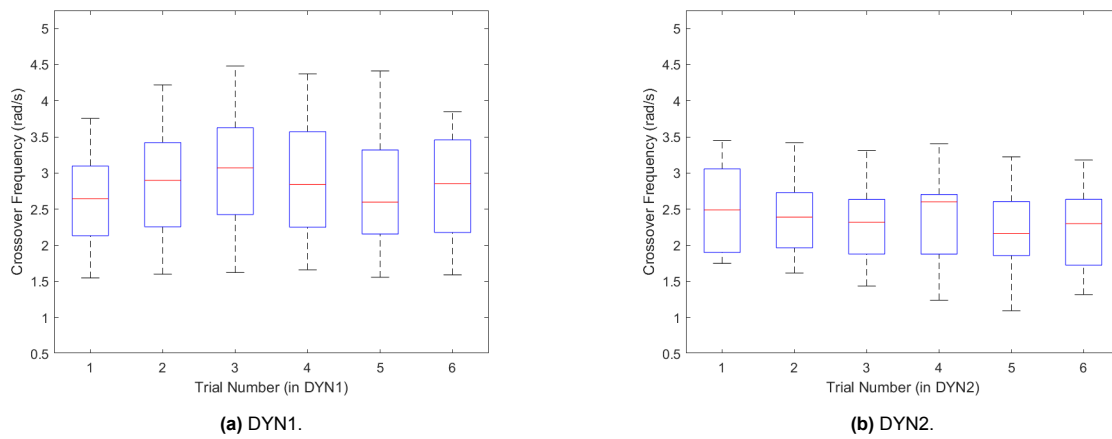


Figure C.4: Crossover frequencies sorted by the order in which the DYN1 and DYN2 trials were performed.




Figure C.5: Detection lags sorted by the order in which the DYN12 and DYN21 trials were performed.

D

Details of Experiment FPs and FN

Table D.1 summarizes the details of the FPs and the FN in the experiment. The column called "Realized Mistake?" represents whether the participant made a comment (either during the run or immediately after the run) that they believed the button push was a mistake. However, it should be noted that participants were not instructed to make such a comment in case they believed they made a mistake, so it is possible that there are more cases where the participant realized they made a mistake but did not say anything. Furthermore, it is interesting to note that Participant 3, Participant 5, and Participant 16 all made (at least) one mistake in the first three trials of the experiment, but then did not have any further FPs or FNs for the remainder of the experiment.

Table D.1: Summary of the details surrounding the FPs and the FN in the experiment.

Participant	Error Type	CE Dynamics	Forcing Function	Trial Number	Realized Mistake?
2	FP	DYN2	9	18	No
2	FP	DYN21	9	28	No
3	FP	DYN2	9	1	Yes
3	FP	DYN2	3	3	Yes
5	FP	DYN2	7	1	No
5	FP	DYN2	1	3	No
12	FP	DYN1	1	8	No
12	FP	DYN2	6	18	No
16	FP	DYN12	2	1	Yes
17	FP	DYN1	9	1	No
17	FP	DYN21	9	16	No
17	FP	DYN2	7	20	No
17	FN	DYN21	6	29	No

E

Including u in Model Threshold

For this research, an attempt was made to include the control input u in the threshold for the DYN21 model. This was done according to Eq. (E.1), where w is a weight. Thus, this approach exclusively *increases* the threshold by an amount proportional to the power (equivalent to the variance) in u in a running 0.2 s window. The 0.2 s window was chosen because Miller and Elkind [23] proposed to include u in the threshold by looking at the magnitude of the change in u from one time step to the next. Thus, a small time, similar to a quick reaction time, was chosen.

$$\text{Threshold including } u(t) = \text{Original threshold} + w \cdot \text{var}(u(t - 0.2s : t)) \quad (\text{E.1})$$

The goal of this approach was to improve the model's approximation of the participant detection lags and therefore only the TPs are considered here. Also, only the DYN21 model is considered since the DYN21 model detections (i.e., when the threshold is exceeded) often occur extremely quickly. The desired effect of increasing the threshold was to increase the detection lags by a few seconds such that they are a better approximation of the participants' detection lags. However, this was not successful because, in most cases, increasing the threshold either did nothing to the model's detection lags, delayed them by only a few time steps (of 0.01 s), or resulted in a FN. To illustrate this, consider the figures below.

As a reference, the TPs from the original DYN21 model (i.e., the case where $w = 0$ in Equation E.1) are presented in Figure E.1. It is clear that many of the model detections occur very quickly and before the participants pressed the button. The results for $w = 50$, $w = 100$, $w = 150$, and $w = 200$ are presented in Figure E.2. As can be seen, increasing the weights does not help in delaying the model detections that occur right around $t = 45$ s (immediately after the transition) by a few seconds. Instead, it either delays it by too much (e.g., to the point where the model detections occur after $t = 55$ s) or leads to a FN. Changing the period of time considered in the calculation of the power in u (i.e., using 0.5 s (Figure E.3) or 1 s (Figure E.4) instead of 0.2 s as in Eq. (E.1)) results in the same problem.

Thus, it appears that there is no way to improve the estimation of the detection lags by including u in the threshold according to Eq. (E.1).







Figure E.1: Comparison of the model's TPs with the participant's TPs in the original model without u in the threshold ($w = 0$).




(a) $w = 50$.



(b) $w = 100$.



(c) $w = 150$.



(d) $w = 200$.

Figure E.2: Comparison of the model's TPs with the participant's TPs over various weights w when considering a 0.2 s window for calculating the power in u .

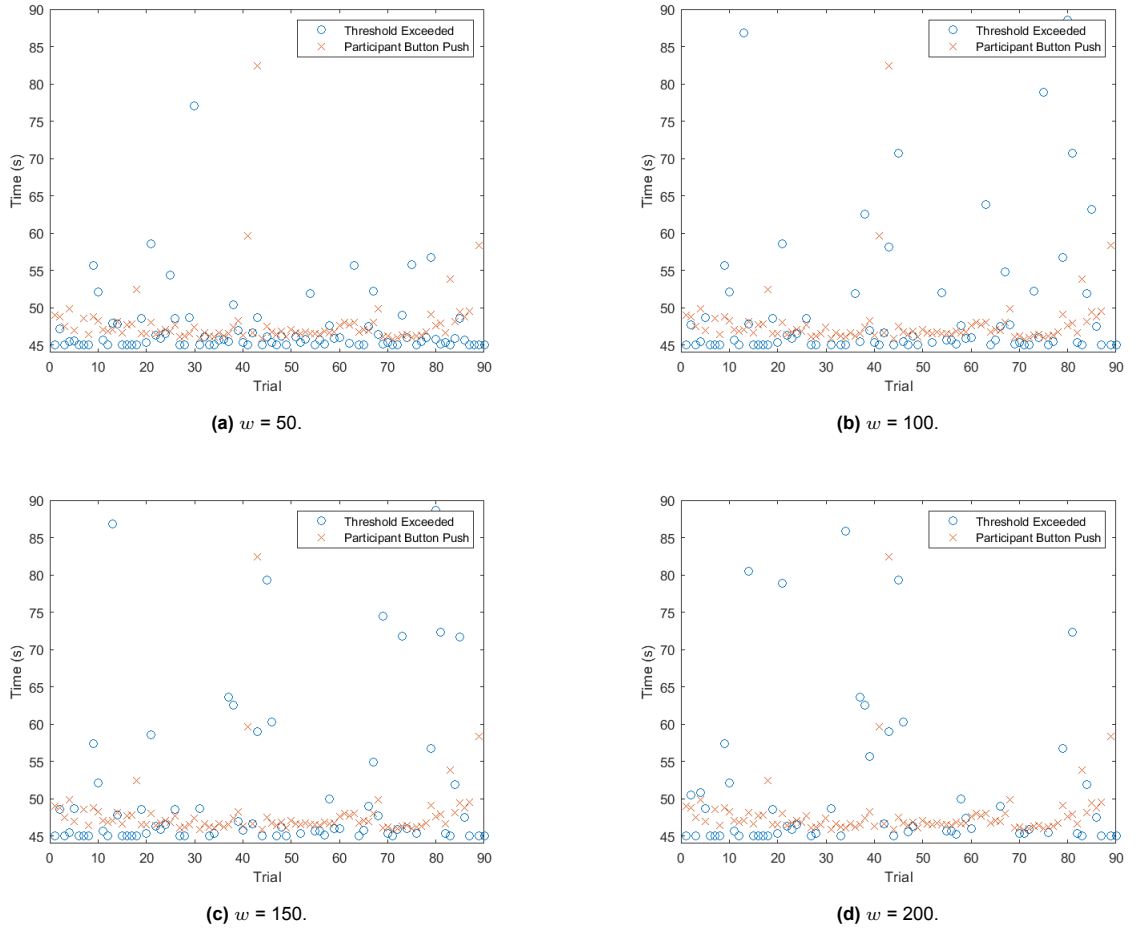


Figure E.3: Comparison of the model's TPs with the participant's TPs over various weights w when considering a 0.5 s window for calculating the power in u .

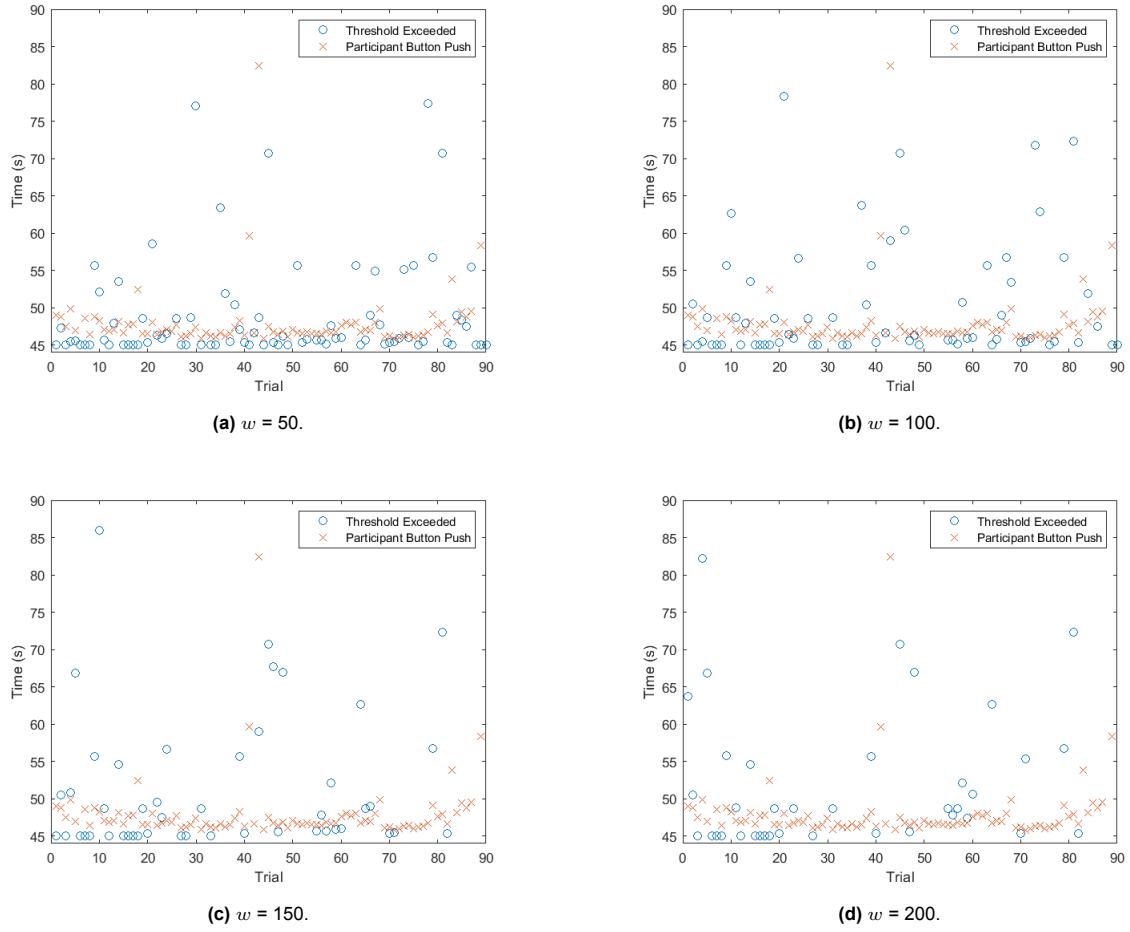



Figure E.4: Comparison of the model's TPs with the participant's TPs over various weights w when considering a 1 s window for calculating the power in u .



DYN21 Transition Time Trace Analysis

The following figures were generated to investigate the role played by the relationship between the control inputs and the system's response (y in this case). The blue-shaded area represents the mean $\pm 1\sigma$ averaged over all three periods of the respective forcing function in the DYN2 trials and over all participants. Every participant's DYN21 trial with that forcing function is then plotted on top of the blue-shaded area. The red crosses indicate when participants pushed the button. In general, what these figures show is that as the gradients in the forcing function become larger (i.e., the forcing function number increases), an increasing number of button pushes occur outside the blue-shaded area in either y or u . Therefore, as the gradients in the forcing functions increase, the difference between the expected response (tracking the forcing function) and the observed response increases, which can cause the participant to detect the change in CE dynamics. On the other hand, when the gradients in the forcing function are smaller, many of the detections occur within the blue-shaded area of both y and u . This suggests the relationship between the control inputs and the system's response becomes more important when the gradients in the forcing function are small.

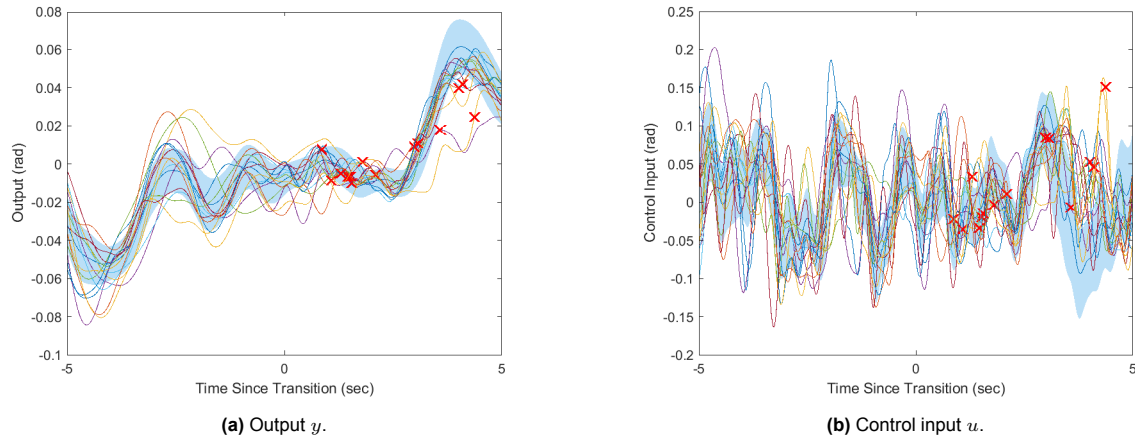


Figure F.1: Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 1. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 1.

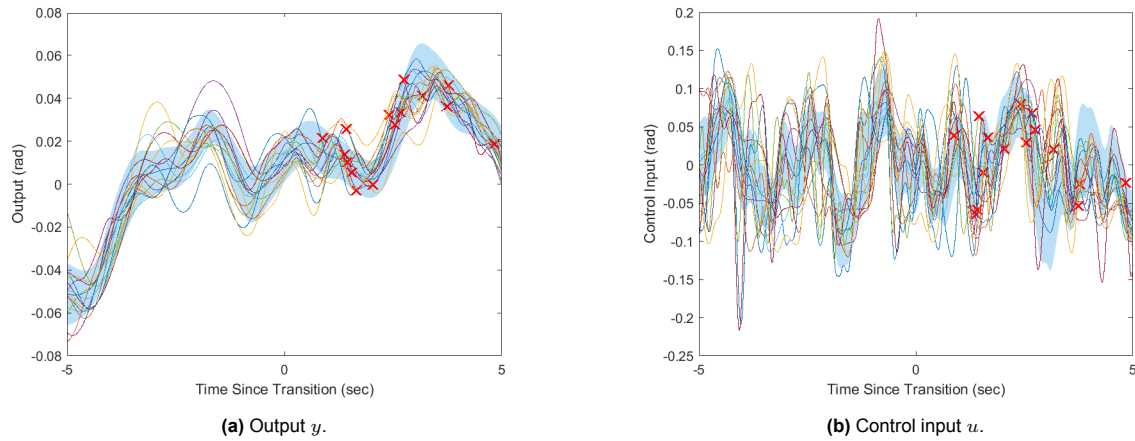


Figure F.2: Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 3. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 3.

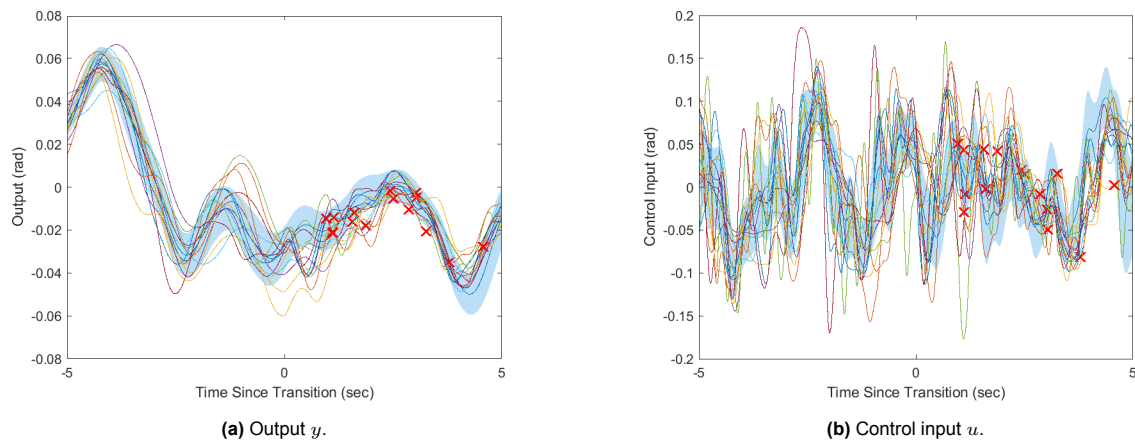


Figure F.3: Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 4. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 4.

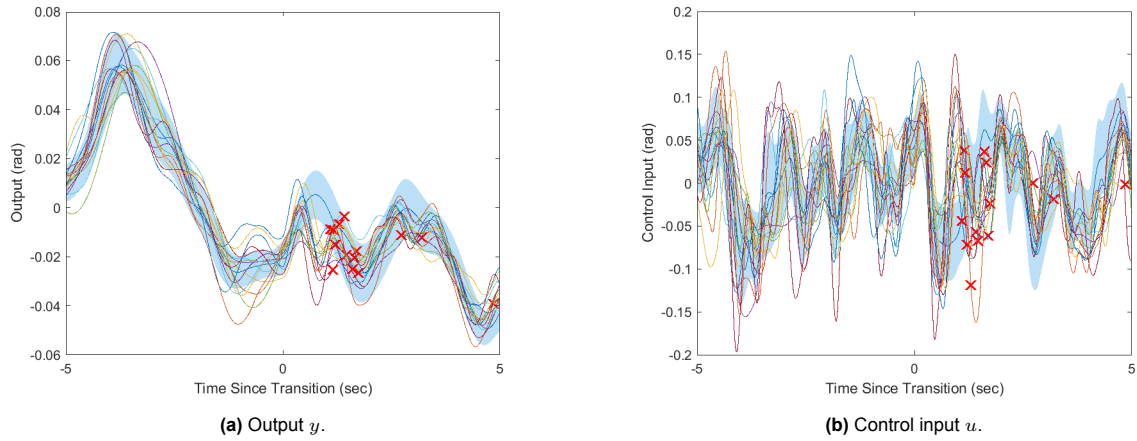


Figure F.4: Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 6. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 6.

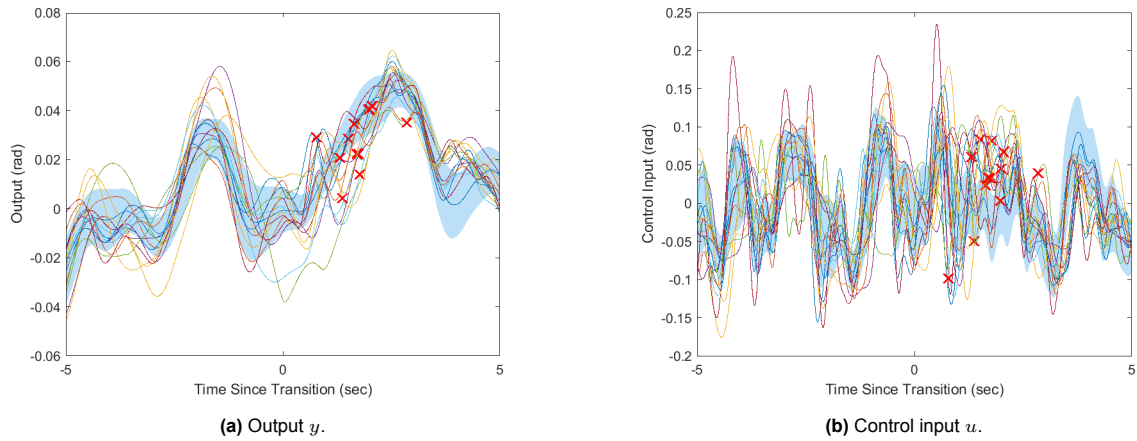


Figure F.5: Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 7. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 7.

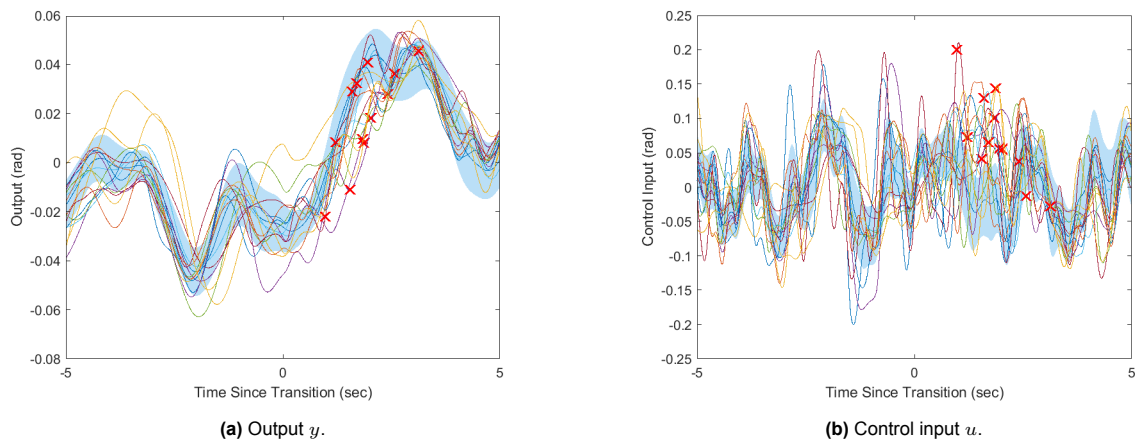


Figure F.6: Output y and control input u of all participants in the ten-second window centered around the transition for the DYN21 condition using forcing function number 9. The shaded blue area represents the mean $\pm 1\sigma$ of all the participants' output and control input over the three periods of the same ten-second window in the DYN2 condition using forcing function number 9.

G

Experiment Latin Squares

The Latin square used for the DYN1 and DYN12 conditions is given in Eq. (G.1). For brevity, the conditions are numbered in the following way (note 'FoFu' stands for 'Forcing Function'):


1. DYN1, FoFu 1
2. DYN12, FoFu 1
3. DYN12, FoFu 2
4. DYN1, FoFu 3
5. DYN12, FoFu 3
6. DYN1, FoFu 4
7. DYN12, FoFu 4
8. DYN12, FoFu 5
9. DYN1, FoFu 6
10. DYN12, FoFu 6
11. DYN1, FoFu 7
12. DYN12, FoFu 7
13. DYN12, FoFu 8
14. DYN1, FoFu 9
15. DYN12, FoFu 9

$$\begin{pmatrix} 12 & 2 & 14 & 15 & 10 & 13 & 8 & 9 & 1 & 5 & 11 & 7 & 6 & 3 & 4 \\ 7 & 12 & 9 & 10 & 5 & 8 & 3 & 4 & 11 & 15 & 6 & 2 & 1 & 13 & 14 \\ 2 & 7 & 4 & 5 & 15 & 3 & 13 & 14 & 6 & 10 & 1 & 12 & 11 & 8 & 9 \\ 13 & 3 & 15 & 1 & 11 & 14 & 9 & 10 & 2 & 6 & 12 & 8 & 7 & 4 & 5 \\ 8 & 13 & 10 & 11 & 6 & 9 & 4 & 5 & 12 & 1 & 7 & 3 & 2 & 14 & 15 \\ 5 & 10 & 7 & 8 & 3 & 6 & 1 & 2 & 9 & 13 & 4 & 15 & 14 & 11 & 12 \\ 14 & 4 & 1 & 2 & 12 & 15 & 10 & 11 & 3 & 7 & 13 & 9 & 8 & 5 & 6 \\ 10 & 15 & 12 & 13 & 8 & 11 & 6 & 7 & 14 & 3 & 9 & 5 & 4 & 1 & 2 \\ 6 & 11 & 8 & 9 & 4 & 7 & 2 & 3 & 10 & 14 & 5 & 1 & 15 & 12 & 13 \\ 9 & 14 & 11 & 12 & 7 & 10 & 5 & 6 & 13 & 2 & 8 & 4 & 3 & 15 & 1 \\ 4 & 9 & 6 & 7 & 2 & 5 & 15 & 1 & 8 & 12 & 3 & 14 & 13 & 10 & 11 \\ 15 & 5 & 2 & 3 & 13 & 1 & 11 & 12 & 4 & 8 & 14 & 10 & 9 & 6 & 7 \\ 11 & 1 & 13 & 14 & 9 & 12 & 7 & 8 & 15 & 4 & 10 & 6 & 5 & 2 & 3 \\ 1 & 6 & 3 & 4 & 14 & 2 & 12 & 13 & 5 & 9 & 15 & 11 & 10 & 7 & 8 \\ 3 & 8 & 5 & 6 & 1 & 4 & 14 & 15 & 7 & 11 & 2 & 13 & 12 & 9 & 10 \end{pmatrix} \quad (G.1)$$

The Latin square used for the DYN2 and DYN21 conditions is given in Eq. (G.2). Again, the conditions are numbered in the following way:

1. DYN2, FoFu 1
2. DYN21, FoFu 1
3. DYN21, FoFu 2
4. DYN2, FoFu 3
5. DYN21, FoFu 3
6. DYN2, FoFu 4
7. DYN21, FoFu 4
8. DYN21, FoFu 5
9. DYN2, FoFu 6
10. DYN21, FoFu 6
11. DYN2, FoFu 7
12. DYN21, FoFu 7
13. DYN21, FoFu 8
14. DYN2, FoFu 9
15. DYN21, FoFu 9

$$\begin{pmatrix}
 9 & 8 & 14 & 1 & 5 & 12 & 10 & 2 & 13 & 7 & 6 & 3 & 15 & 4 & 11 \\
 14 & 13 & 4 & 6 & 10 & 2 & 15 & 7 & 3 & 12 & 11 & 8 & 5 & 9 & 1 \\
 7 & 6 & 12 & 14 & 3 & 10 & 8 & 15 & 11 & 5 & 4 & 1 & 13 & 2 & 9 \\
 11 & 10 & 1 & 3 & 7 & 14 & 12 & 4 & 15 & 9 & 8 & 5 & 2 & 6 & 13 \\
 10 & 9 & 15 & 2 & 6 & 13 & 11 & 3 & 14 & 8 & 7 & 4 & 1 & 5 & 12 \\
 1 & 15 & 6 & 8 & 12 & 4 & 2 & 9 & 5 & 14 & 13 & 10 & 7 & 11 & 3 \\
 15 & 14 & 5 & 7 & 11 & 3 & 1 & 8 & 4 & 13 & 12 & 9 & 6 & 10 & 2 \\
 2 & 1 & 7 & 9 & 13 & 5 & 3 & 10 & 6 & 15 & 14 & 11 & 8 & 12 & 4 \\
 6 & 5 & 11 & 13 & 2 & 9 & 7 & 14 & 10 & 4 & 3 & 15 & 12 & 1 & 8 \\
 3 & 2 & 8 & 10 & 14 & 6 & 4 & 11 & 7 & 1 & 15 & 12 & 9 & 13 & 5 \\
 4 & 3 & 9 & 11 & 15 & 7 & 5 & 12 & 8 & 2 & 1 & 13 & 10 & 14 & 6 \\
 12 & 11 & 2 & 4 & 8 & 15 & 13 & 5 & 1 & 10 & 9 & 6 & 3 & 7 & 14 \\
 13 & 12 & 3 & 5 & 9 & 1 & 14 & 6 & 2 & 11 & 10 & 7 & 4 & 8 & 15 \\
 5 & 4 & 10 & 12 & 1 & 8 & 6 & 13 & 9 & 3 & 2 & 14 & 11 & 15 & 7 \\
 8 & 7 & 13 & 15 & 4 & 11 & 9 & 1 & 12 & 6 & 5 & 2 & 14 & 3 & 10
 \end{pmatrix} \quad (G.2)$$



Documents for Human Research Ethics Committee

The following documents were submitted to the Human Research Ethics Committee for approval under application number 3307. The informed consent form is presented first, followed by the experiment briefing.

Experiment Consent Form

Modeling the Human Operator's Detection of a Change in Controlled Element Dynamics

I hereby confirm, by ticking the box, that:

1. I volunteer to participate in the experiment conducted by the researcher (-), under supervision of -, from the Faculty of Aerospace Engineering of TU Delft. I understand that my participation in this experiment is voluntary and that I may withdraw ("opt-out") from the study at any time, for any reason. ☐
2. I have read the briefing document and I understand the experiment instructions, and have had all remaining questions answered to my satisfaction. ☐
3. I understand that my participation involves performing manual tracking tasks in the HMI-Lab simulator at TU Delft. I understand that only the pseudonymized recorded time traces of the tracking tasks are saved and used for data analysis. ☐
4. I confirm that the researcher has provided me with detailed safety and operational instructions for the HMI-Lab simulator (simulator setup, electro-hydraulic side stick, emergency procedures) used in the experiment. Furthermore, I understand the researcher's instructions for guaranteeing the experiment's compliance with current COVID-19 guidelines, and that this experiment shall at all times follow these guidelines. ☐
5. I understand that the researcher will not identify me by name in any reports or publications that will result from this experiment, and that my confidentiality as a participant in this study will remain secure. Specifically, I understand that any demographic information I provide (gender, age range, **see next page**) will only be used for reference and always presented in aggregated form in scientific publications. ☐
6. I understand that this research study has been reviewed and approved by the TU Delft Human Research Ethics Committee (HREC). To report any problems regarding my participation in the experiment, I know I can contact the researchers using the contact information below. ☐

My Signature

Date

My Printed Name

Signature of researcher

Contact information researcher:

-

Contact information research supervisor:

-

Participant Demographic Information

Modeling the Human Operator's Detection of a Change in Controlled Element Dynamics

Age range:

- ☐ 18 – 19
- ☐ 20 – 24
- ☐ 25 – 29
- ☐ 30 – 34
- ☐ 35 – 39
- ☐ 40 – 44
- ☐ 45 – 49
- ☐ 50 – 55
- ☐ 55+

Gender: _____

Participant number: _____
(filled out by the researcher)

Contact information researcher:

-

Contact information research supervisor:

-

Experiment Briefing

Modeling the Human Operator's Detection of a Change in Controlled Element Dynamics

First of all, thank you for your participation! This experiment is part of an MSc thesis research project that aims to understand and model how a human operator (you) detects that the controlled element dynamics have changed. The experiment consists of a simple tracking task and will be performed in the Human-Machine Interaction Laboratory (HMI-Lab) at TU Delft's Faculty of Aerospace Engineering. This briefing will give an overview of the experiment and explains what is expected from you as a participant. Please read this document carefully. Should any questions or comments remain, always feel free to discuss these with the researcher conducting the experiment.

Experiment Objective

The goal of this experiment is to investigate how human operators detect changes in the controlled element dynamics. Data from this experiment will be used to develop and validate a model that can predict the moment a human operator detects a change in controlled element dynamics and consequently can be used in future models on adaptive manual control behavior.

Experiment Set-up

The experiment will take place in the HMI-Lab (Fig. 1), a fixed-base simulator set-up at TU Delft's Faculty of Aerospace Engineering. The task you will be carrying out is a tracking task with a pursuit display (i.e. you can see the target signal and the controlled element output), as shown in Fig. 2. Please take a seat on the chair on the right and use the side-stick on your right-hand side to control the system by moving the side-stick to the left and right. There is another side-stick on your left-hand side. The purpose of this side-stick will be explained in the next section.



Figure 1: Illustration of HMI-Lab. The participant will be sitting on the right (blue) seat and controls the side-stick.




Figure 2: Sketch of the HMI Lab pursuit display.

Experiment procedure

Before starting the data collection, some training runs will be performed so that you can familiarize yourself with the different scenarios and controlled element dynamics. In the experiment, you will control dynamics that respond like a single integrator (rate control) or like a double integrator (acceleration control), and transitions between these different dynamics. Once your performance is

stable, the data collection phase will begin. There is no pre-defined number of tracking runs that are needed for collecting the data: the experimenter will simply notify you when sufficient data has been collected. All individual tracking runs last 100 seconds.

It is important that you continue to focus on keeping the error (the distance between the target signal and the controlled element output) as low as possible during the entire run by continuously controlling the system with the side-stick on your **right**. At the end of each run, your score will be communicated to you by the researcher. During the runs, the controlled element dynamics may change. If you detect a change in the controlled element dynamics, please **immediately** indicate this by pressing the button on the side-stick to your **left**. Please note that only **some** of the runs will contain a change in the controlled element dynamics, so it is possible to go through an entire run without needing to press the button.

Short breaks can be taken between runs to alleviate any discomfort that might occur due to fatigue, controlling the side-stick, or sitting in a fixed position for a prolonged period of time. A longer break of approximately 10 minutes will be taken after the first hour. The experiment will last approximately two hours. Should more breaks be required, you can request them at any moment.

For each subsequent trial, the following procedure will be applied:

1. The researcher applies the settings for the next run.
2. The researcher checks that the participant is ready to proceed and initiates the run after a countdown from 3 (3-2-1-go).
3. The participant performs the tracking task.
4. The participant will be notified of their performance in the run in terms of the error score that will appear on the display after completing the run.

Your Rights & Consent

Experiment participation is voluntary. Should you feel uncomfortable, you can decide to stop your participation at any time. By participating in the experiment you agree that the collected data may be published. Your personal data will remain confidential and anonymous, only the researcher can link the collected data to a specific participant. To ensure you understand and comply with the conditions of the experiment, you will be asked to sign an informed consent form.

<u>Contact information researcher:</u> -	<u>Contact information research supervisor</u> -
---	---

Thank you again for participating!