University of Glasgow | School of Engineering

Determining the facilitative value of intermittent control systems in the direct approach to system identification

By Matteo Rochon Cocchiara, 2377128R

Coursework Declaration and Feedback Form

The Student should complete and sign this part

Student Number : 2377128	Student Name : Matteo Nicholas Rochon Cocchiara
Course Code : ENG4110P	Course Name : Individual Project 4
1st Supervisor: Henrik Gollee	2nd Supervisor: Thomas Franke
Title of Assignment : Determining the facilitative value of intermittent control systems in the direct approach to system identification	

Declaration of Originality and Submission Information

I affirm that this submission is all my own work in accordance with the University of Glasgow Regulations and the School of Engineering requirements Signed (Student) : Matteo Rochon Cocchiara	
Date of Submission : 18.04.2022	

Feedback from Lecturer to Student – to be completed by Lecturer or Demonstrator	
Grade Awarded: Feedback (as appropriate to the coursework which was assessed):	
Lecturer/Demonstrator:	Date returned to the Teaching Office:

Abstract

This study aimed to determine if the use of an intermittent controller structure, either clock-driven or event-driven, could facilitate system identification when implementing the direct approach (as defined by Forssell and Ljung, 1999). This was accomplished by simulating the behaviour of a single-segment inverted pendulum in both a stabilised and unstable form with varying amplitudes of noise. The controller setup was then alternated across sets of trials between clock-driven intermittent, event-driven intermittent and predictive continuous control. The direct approach to system identification was then implemented using the MATLAB System Identification Toolbox, in order to determine whether fits obtained when implementing an intermittent controller were more accurate than those in which a continuous controller was used. The results suggest that the intermittent controllers do indeed show some value as facilitators to the direct approach, with clock-driven control significantly outperforming continuous control.

Structural Outline

Section 1 outlines the aims of this report, and delineates the question being addressed. Given the large volume of discordant literature on the subjects, a review has been included in Section 1.3 to help contextualise the problem in the fields of system identification of human movement control. Section 2.1 provides the essential theory underpinning the rest of the project. Section 2.2 outlines the overall procedure for the investigation, and Section 2.3 provides greater detail on specific elements of the methodology. The results are displayed in Section 3, and the discussion has been grouped thematically in the different subsections. Conclusory remarks in Section 4 recapitulate the most significant findings, as well as their limitations. Finally, in Section 5, the results of this study are utilised in conjunction with the existing literature to recommend further areas of research

Acknowledgements

I would first like to thank my project supervisor Dr Henrik Gollee for his support throughout this academic year and for his engaging lectures which sparked my interest in control. Thanks also to Alberto Álvarez Martin and Abril Surisadday Díaz Vergara who were unstintingly generous with their time and knowledge when I was learning the basics of system identification with MATLAB. I am also grateful to Professor Thomas Franke for his support in the role of second supervisor.

I also owe a tremendous debt of gratitude to my family. Special thanks to my Mum and Dad for always being there for me; to my sister Anna, whose work ethic never fails to inspire me; and to my Nonna for encouraging me to study for the love of learning rather than grades, and for reminding me not to work too hard.

Contents

List of Figures

Key:
Clock-driven control (CDC), Event-Driven Control (EDC), Predictive Continuous Control, PCC, Stable system (S), Unstable system (US)

List of Tables

Section 1: Project Aims and Introduction

1.1 – Aims

This project aims to establish whether the use of an intermittent controller can enable the fitting of more accurate transfer functions when using only input and output timeseries data from a plant (the direct approach). Clock-driven control (herein CDC) and event-driven control (EDC) strategies will be implemented and compared to predictive continuous control (PCC) to determine if there is any statistically significant difference in transfer function accuracy. The impact of open-loop interval (for CDC), threshold (for EDC), and time delay (for PCC) will also be examined, in order to determine the values of these variables that lead to optimal system identification. Furthermore, zero-mean Gaussian white noise will be added to the output of the transfer function to simulate a real system more accurately; the amplitude of this noise will be varied across sets of simulations to discern the impact of noise on fitting accuracy for each controller structure. Finally, different estimation techniques will be used to determine if they have an impact on fitting efficacy.

1.2 – Introduction

Human movement is integral to daily life; however, our knowledge of its underlying control mechanism remains incomplete. Previous attempts at modelling can be broadly broken down into two groups: those that envision sustained motion as a concatenation of ballistic movements (intermittent control, or sometimes serial ballistic control), and those that argue that parts of the body are in fact continuously controlled.

The continuous control hypothesis posits that information about the current state of the body is continuously conveyed to, and corrected by, the body's control system (van der Kooij and de Vlugt, 2007). Continuous models can be separated by their inclusion or exclusion of a predictive element which is used to counteract the influence of a pure time-delay, which is often included in models of physiological systems (Gawthrop, Loram and Lakie, 2009). Some previous studies have credited predictive continuous models with being the major actor in maintenance of human balance (van der Kooij and de Vlugt, 2007).

By contrast, the intermittent control (herein IC) paradigm postulates that, although information may be continuously conveyed to the control system, control action is exerted in the form of a series of open-loop trajectories (Loram *et al.*, 2014). The calculation of such a trajectory can be done at regular time intervals (often termed clock-driven intermittent control (Loram *et al.*, 2014)) or once the error between the open-loop and closed-loop state estimates exceeds a certain threshold (event-driven intermittent control (Gawthrop and Wang, 2009)). Though experimental evidence in favour of IC is mixed, it has the advantage of explaining certain observed phenomena, such as the psychological refractory period (Gawthrop *et al.*, 2011) which cannot be explained by the continuous control paradigm.

System identification has, naturally, played an integral role in many of these studies. However, the choice of system identification method must be considered alongside the underlying control structure (van der Kooij, van Asseldonk and van der Helm, 2005). Thus, a dilemma emerges since the selection of an identification tool

is predicated on the (unknown) feedback structure. This quandary is starkly evident when one examines methods that fall under Ljung's classification of the "Direct Approach" (Forssell and Ljung, 1999), wherein input and output data are used to derive information about the plant without integrating information about the controller structure. These methods are not applicable to closed-loop continuous control systems since any attempt to relate the input and output would not solely reflect the properties of the plant (Forssell and Ljung, 1999). Additionally, the negative impact of noise is exacerbated by its constant reintroduction to the plant. However, if the serial ballistic hypothesis is indeed correct, it is possible that the input and output data could be sufficient to generate valid models.

To this end, this project aimed to establish the extent to which intermittent control can act as a facilitator to the direct approach to system identification. This was accomplished by fitting transfer function models to input-output data from an inverted pendulum model of human standing; this was done using the MATLAB System identification Toolbox. The resultant transfer functions were post-processed in order to determine whether there was a statistically significant difference in the quality of fits obtained when implementing an intermittent controller as compared to a continuous controller.

1.3 – Literature Review

This literature review has been included to provide context to the current study, and to clarify the possible implications its findings could have on the fields of system identification and human movement control. To date, there are no studies which assess the potential of intermittent control (IC) as a facilitator for the direct approach (DA) to system identification. First, this review will examine the relevant literature pertaining to the identification of closed-loop systems. Subsequently, the points from the first section will be contextualised through a brief analysis of the available research on human postural control. Finally, the findings presented in the two preceding sections will be combined in order to establish why intermittent control has not yet been explored as a facilitator to system identification.

In 1977, Gustavsson et al. (Gustavsson, Ljung and Söderström, 1977), noted that the methods of system identification could be categorised according to the assumptions made regarding the feedback mechanism. These classifications remain largely unchanged with the exception of a small tweak by Forssell and Ljung in 1999 (Forssell and Ljung, 1999). As suggested by Gevers (Gevers, 2003), the need for such designations arose from the rapid expansion in the number of papers relating to system identification that succeeded the publication of the seminal Åström-Bohlin and Ho-Kalman papers in the mid-1960s (Åström and Bohlin, 1965; Ho and Kalman, 1966).

According to the updated definitions, the direct-approach to system identification is defined as identification of the open-loop system using measures of input and output whilst ignoring feedback mechanisms (Forssell and Ljung, 1999). Therefore, the authors note, it would be inappropriate to apply this method to closed-loop systems since the output would be correlated to the input. However, the authors also note that it is a perfectly valid means of system identification in the context of open-loop control. Thus, if the view is entertained that

human movement represents a series of open-loop trajectories (i.e. is controlled intermittently), the direct approach may remain a useful analytical tool. However, as will be seen in the next few paragraphs, movement control remains an area of major debate among the scientific community.

As of this writing, the nature of the human postural control system remains an unsolved question. At the centre of this dilemma are three established models, each with experimental evidence to their credit: non-predictive continuous control (NPCC), predictive continuous control (PCC) and intermittent control (IC). Originally, these theories were suggested independently: for example, the first reported association of IC with human movement control was in a posthumously published paper by K. J. Craik (Craik, 1947). More than four decades later, a team led by Johansson et al. published a paper in which they fitted a continuous PID controller model to experimental data of perturbed standing (Johansson, Magnusson and Akesson, 1988); no mention is made of the possibility that human posture may not be continuously controlled. This would not be the last time that experimental data was framed in the context of one model without regard for the possibility of other controller configurations.

As experimental evidence mounted in favour of each theory, so too did the need to cross-examine these different models. This need was addressed by several authors (van der Kooij and de Vlugt, 2007; Gawthrop, Loram and Lakie, 2009; Gollee *et al.*, 2012), who applied both numerical and experimental methods in order to determine the true nature of the human control system. These later studies were successful in validating PCC over NPCC (Gawthrop, Loram and Lakie, 2009), but disagreement between studies meant that a unified theory of human postural control never emerged. Indeed, despite the efforts of several teams of researchers, the results remained conflicting: for example, a frequently-cited 2007 paper by van de Kooij et al. boldly stated that continuous control "dominated" human postural responses (van der Kooij and de Vlugt, 2007). However, later rebuttals (Loram *et al.*, 2014) by advocates of IC cite its capacity to masquerade as continuous control under specific conditions, as well as its ability to explain physiological phenomena that, as of this writing, are still left unaddressed by the continuous control paradigm.

The most recent studies still fall short of providing a unified view of human movement control. Some investigations, such as a 2014 paper by Elias et al. (Elias, Watanabe and Kohn, 2014) note that the structure of human physiology permits the implementation of both continuous and intermittent control mechanisms in postural control. Such thoughts have also been echoed in engineering studies (Tanabe *et al.*, 2016). There exist nonetheless some papers which continue to frame their methodologies and results in the context of particular models without acknowledging this potential duality (Hwang *et al.*, 2016).

Hence, the discordant findings of human postural studies provide an explanation for why intermittent control has thus far been overlooked as a facilitator for system identification. Given the inextricable link between controller design and identification technique, methods that did not match the control paradigm under examination were rejected out of hand. For example, when discussing the implementation of the direct

approach in previous studies, a 2005 paper by van de Kooij et al. commented that the inaptness of the method rendered the results “meaningless and confusing” (van der Kooij, van Asseldonk and van der Helm, 2005). However, if one removes the supposition that the underlying controller structure is continuous, it unlocks the possibility that the direct approach may provide a fruitful means of system identification.

One can conclude, based on the literature, that conclusively demonstrating the facilitative value of IC systems could have far-reaching benefits in the study of human movement control, and more broadly in the identification of closed-loop systems. However, progress in this area has been stifled by the conflicting results of previous studies, particularly in the area of human postural control.

Section 2: Methodology

2.1 – Essential Theory

2.1.1 – System identification and transfer functions

This section aims to clarify what is meant by the "direct approach", and to provide background for the methods outlined in Section 2.2. System identification is applied in a wide variety of contexts, and with significant variation, rendering it challenging to isolate a single, comprehensive definition. However, broadly, system identification refers to the practice of deriving a model that relates observed quantities (Ljung, 2010). Methods of system identification can be broken down by the data they use to relate these observed quantities. This report implements the set of definitions expressed by Forssell and Ljung in their 1999 paper "Closed-loop identification revisited" (Forssell and Ljung, 1999). Here, they provide three categories of system identification methods based on the information used to derive the model. Firstly, there is the direct approach (DA), which is the focus of this paper: in this method, the inputs and outputs are used to derive a model of the plant, without utilising any information about the controller structure. The alternatives to this approach are the indirect approach and the joint input-output method: since these are comprehensively reviewed in (Forssell and Ljung, 1999; van der Kooij, van Asseldonk and van der Helm, 2005), they will not be explained in detail here. However, it is important to note that the two alternative methods have shortcomings not shared by the DA. Specifically, they either require a priori knowledge of the controller, or would necessitate further experimentation on the system to derive sensitivity functions (van der Kooij, van Asseldonk and van der Helm, 2005).

In modern control, it is common to represent systems in state-space form (Equation 1), which relates the inputs (u) and outputs (y) by a set of differential equations in terms of the states (x) (Antsaklis and Michel, 2007). One of the primary reasons for its widespread use is its versatility: through restructuring using similarity transformations, one can derive valuable information about the controllability and observability of the system (Antsaklis and Michel, 2007). However, it is precisely this propensity for rearrangement that renders state-space form unsuited to comparison between models; this is because, the matrices $\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C},$ and $\boldsymbol{D}$ can take on an infinite number of configurations through transformation using non-singular matrices (Hinrichsen and Pritchard, 2005). Therefore, given that this study is entirely reliant on comparison between fitted models and the true system, an alternative means of parametric representation is required.

$$\dot{x} = \boldsymbol{A}x + \boldsymbol{B}u$$
$$y = \boldsymbol{C}x + \boldsymbol{D}u$$

Equation 1

Transfer functions (Equation 2) provide a means of representing the relationship between the inputs and outputs of a system in the frequency domain (Isermann and Münchhof, 2011). Transfer functions are comprised of the ratio of two polynomials, and the order of the numerator and denominator represent the number of zeros and poles of the system respectively (Iqbal, 2017). Unlike state-space models, transfer functions only have a

single form, rendering them more easily compared and thus ideal for the model fitting in this study. The roots of the denominator polynomial, often called the characteristic equation, are the system poles; the location of these poles on the complex plane is highly indicative of the dynamic properties of a system (Morud and Skogestad, 1996).

$$\frac{Y(s)}{U(s)} = \frac{d}{as^2 + bs + c} \qquad \text{Equation 2}$$

Note that the transfer function in Equation 2 represents a generic system with no zeros and two poles, which is analogous to the system under examination in this study.

2.1.2 – Intermittent control

This section aims to provide an understanding about the basics of intermittent control, to aid interpretation of subsequent sections. Open-loop control and closed-loop control can be considered to sit on opposing sides of a spectrum: in the former, the output is never fed back, meaning that control action is exerted without accounting for the current state of the system. By contrast, in closed-loop control, the output is being measured and used to regulate control action. For this reason, open-loop control alone is insufficient to stabilise an unstable plant (Skogestad, Havre and Larsson, 2002).

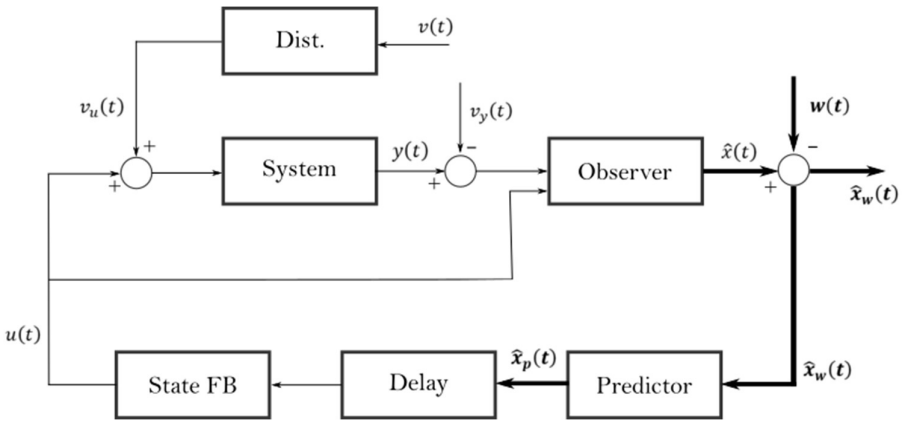

Figure 1: A block-diagram representation of a predictive continuous controller, as adapted from (Gawthrop *et al.*, 2011). The same notation has been employed, with the thin arrows representing scalar quantities, and the thick arrows representing vectors. Importantly, note the summing junction after the "System" block, where $v_y(t)$ is added; this is where the sensor noise was implemented during the simulations.

IC systems can be considered intermediates between the two extremes of open and closed loop control: instead of fastidious state observation and correction, IC systems exert control action by means of sequential, open-loop trajectories. According to some authors, this intermittency can provide a means of sparing computational resources (Gawthrop *et al.*, 2011).

By examining Figures 1 and 2, one notes that the only blocks not common to the two diagrams are those labelled "Hold" and "Trig."; these are the drivers of intermittent control action. The "Trig." block is responsible for triggering the sampling action of the controller: this updates the value of $\widehat{\boldsymbol{x}}_{\boldsymbol{w}}(\boldsymbol{t_i})$ at sample times t_i. This, in turn, updates the value of $\widehat{\boldsymbol{x}}_{\boldsymbol{h}}(\boldsymbol{t})$, thus altering the vector that is fed through the state feedback controller and changing the control trajectory. As its name suggests, the "Hold" block ensures that the sampled signal $\widehat{\boldsymbol{x}}_{\boldsymbol{p}}(\boldsymbol{t_i} - \boldsymbol{t_d})$ is defined at all values of t. Most often, the function implemented in the "Hold" block is not a zero-order hold as implemented in most digital controllers. This is because the zero-order hold is not generally effective when implemented alongside an event-driven controller (Gawthrop *et al.*, 2011). Instead, a generalised hold function is adopted, specifically the system-matched hold as outlined in (Gawthrop and Wang, 2011).

As mentioned, the two primary manifestations of intermittent control are clock-driven control (CDC) and event-driven control (EDC). In CDC, the "Trig." block prompts sampling at regular intervals, i.e., t_i is constant. By contrast, in EDC, sampling is triggered when the difference between the hold state $\widehat{\boldsymbol{x}}_{\boldsymbol{h}}(\boldsymbol{t})$ and the observer state $\widehat{\boldsymbol{x}}_{\boldsymbol{w}}(\boldsymbol{t})$ exceeds a given threshold. As such, sampling is generally irregular (t_i is not constant). However, it is critical to note that event-driven controllers also have a minimum sampling time Δ_{min} (Gawthrop, Gollee and Loram, 2018), which prevents overactivity of the controller. Therefore, in cases where the system output is driven to extremes, such as by external disturbance, the effectuation of new control trajectories will become regular. Thus, the controller will transition from event-driven to clock-driven with an open-loop interval (t_i) of Δ_{min}.

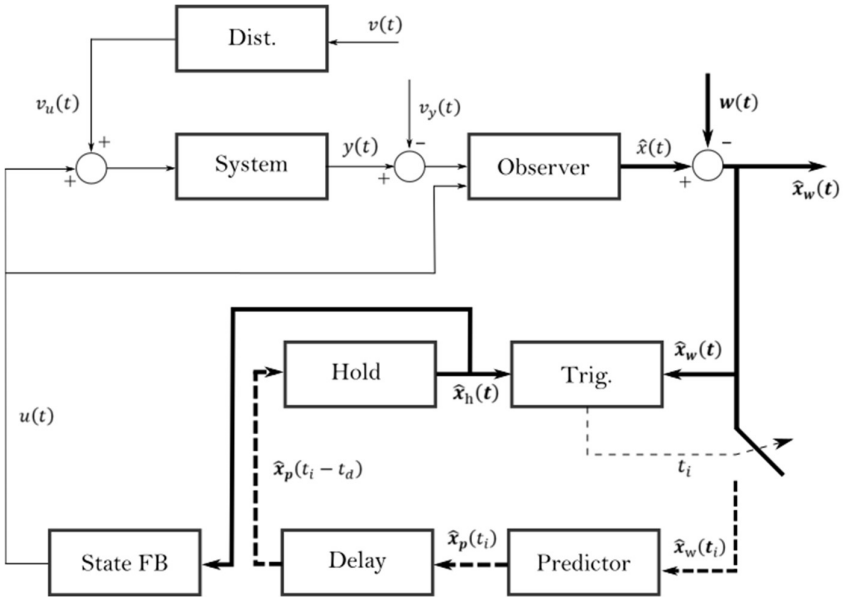

Figure 2: A block diagram representation of an intermittent controller, as adapted from (Gawthrop *et al.*, 2011). Note the introduction of several new blocks not included in the continuous controller (Figure 1). Despite these differences, the point of application of the sensor noise $v_y(t)$ remains the same. As in the original paper, the bold arrows represent vector quantities, the slim arrows represent scalar quantities, and the dashed arrows indicate values that are only defined at times when the sampling trigger is active (i.e., when $t = t_i$)

2.2 – Overview of method

2.2.1 – Background

In order to assess the facilitative value of intermittent control systems, input and output data had to be collected from a system regulated by a variety of controller setups. In the present study, this information was obtained from a MATLAB simulation of a one-segment inverted pendulum designed by Álvarez and Doublein (Álvarez and Doublein, 2021). The Álvarez and Doublein model (henceforth simply "the Model") is based on the conversion of the dynamic equations of the inverted pendulum to a second-order state-space form. The structure of the Model is detailed more thoroughly in Section 2.3.1.

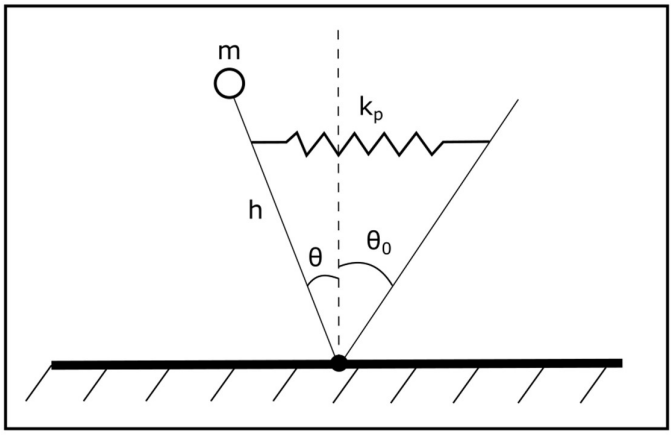

Figure 3: A schematic showing the structure of the single-segment inverted pendulum modelled by Álvarez and Doublein. More detail is provided in Section 2.3.1. Figure adapted from (Álvarez and Doublein, 2021)

The Model was selected for this investigation because of its many tuneable parameters: for example, the state-space matrices underpinning the model were able to be modified, allowing experimentation on the original (unstable) system, as well as a stabilised version (c.f., Section 2.3.1). Another critical benefit of this simulation was that the Model was capable of switching its mode of operation between intermittent and continuous control. Furthermore, the nature of the intermittent control could be specified (i.e., event-driven, or clock-driven), as could the controller properties (including open-loop interval and error threshold). Given the high flexibility of the model, the variables relevant to system identification needed to be isolated in order to ensure that the method employed would provide a valid test of the thesis. With regards to intermittent control, it is evident that the open-loop interval would likely play a critical role in the efficacy of DA system identification approaches. Since the DA is most applicable to open-loop systems (Forssell and Ljung, 1999), it follows that increasing the threshold in EDC and the open-loop interval in CDC could improve the accuracy of system identification. In the case of PCC, alteration of the time delay would likely affect the quality of identification since the fitting mechanisms employed (c.f., Section 2.3.3) do not account for a delay component. Thus, the expectation is that the quality of system identification would deteriorate, along with the quality of control.

2.2.2 – Important note: Variables of Interest

Throughout Section 2, the shorthand "Variable of Interest" is used. The reason for this is that not all variables are relevant to each control structure: for example, clock-driven intermittent control has no relation to thresholds, so there would be no sense to varying the threshold value in a set of simulations that used a clock-driven controller. Hence, the designation "Variable of Interest" was devised to enable emphasis to be placed on the most relevant variables for each controller structure. This would allow the derivation of more meaningful, robust conclusions about the impact of each variable on system identification. The Variables of Interest are as follows for each controller structure:

- Clock-Driven Intermittent Control (CDC) – Open Loop Interval (t_i)
- Event- Driven Intermittent Control (EDC) – Error Threshold
- Predictive Continuous Control (PCC) – Time Delay

For more details on the default parameters, including the default inputs, see Appendix A.

2.2.3 – Overview of method

The Model simulated the behaviour of an inverted pendulum over a period of time that could be defined by the user. The default duration was thirty seconds, and this was retained for this study since it provided time for the pendulum to respond to multiple step inputs (c.f. Appendix A). Additionally, the sampling time was kept at 1ms, since this would provide a large dataset from which to obtain accurate models.

The initial conditions of the simulation were consistent across all trials: the velocity of the mass at the pendulum tip was equal to $+0.1ms^{-1}$, with the pendulum angle equal to 0°. Following implementation of these conditions, the simulation would run for thirty seconds. At each sampling point, the angle of the

pendulum and the control input were saved to two separate arrays. These arrays were then used as arguments for the MATLAB "iddata()" function along with the sampling time. The output of this operation yielded an "iddata" object that could be used as an input to the "tfest()" function.

Besides the "iddata" object, the "tfest()" function demanded inputs of the desired number of poles and zeros (2 and 0 respectively, in this instance). Furthermore, because this study sought to determine the facilitative value of IC under a range of conditions, the "tfest()" function was given a further argument to alter the algorithm that was used to estimate the initial parameter values; more detail is provided in Section 2.3.3. The output of the "tfest()" function was a transfer function with numerator and denominator orders as specified in the function input. The poles of this fitted transfer function were found using the "pole()" method, and compared to the true poles as detailed in Section 2.3.4.

Two further pieces of data were collected for each trial: the median open-loop interval, and the median error. The median open-loop interval was found by finding the time intervals between which a new controller trajectory was initiated using the MATLAB "findpeaks()" function. Finding the differences in peak locations would yield an array of data, the median of which was the open-loop interval; of course, this was only relevant for clock-driven and event-driven control (and in the former, all open-loop intervals would be equal to the median). The median error was obtained by finding the absolute error between the reference input $w(t)$ and the output of the system $y(t)$: this would provide an array of data the size of the original timeseries. The median of these values was the median error, and could be used to obtain an indication of the quality of control. Granted, this would be an imperfect metric, but this report was unable to find an established alternative in the literature. Using the notation of Figures 1 and 2, the error at the i-th timestep can be calculated as shown in Equation 3. In this case, T represents the total simulation time, and h represents the sampling interval.

$$\epsilon = |w(t_i) - y(t_i)| \qquad 0 \leq i \leq \frac{T}{h} \qquad \text{Equation 3}$$

The paragraphs above described a single trial. As mentioned, each simulation included a function that simulated sensor noise; this was added to the true output as outlined in Section 2.3.2. In order to generate this noise, a random number generator function was used; to control the output of this random number generator, a "seed" could be selected by specifying an argument for the MATLAB "rng()" function. An identical seed would ensure that the noise generated across trials was uniform. Hence, to prevent overfitting to the noise, each simulation was processed eight times implementing a different integer between 1 and 8. This constituted a result set. This is important because, in Section 3, references will be made to the median value of a particular result set. To remove any ambiguity, one result set implements the same type of controller, the same noise scale factor, and the same parameter initialisation method.

Then, these variables were altered in turn in order to explore how they each affected the accuracy of system identification. Each result set was then repeated three times, once for each type of controller: event-driven, clock-driven and predictive continuous. Subsequently, the amplitude of noise was increased by an increment of 0.1 (c.f. Section 2.3.2). The minimum scale factor was 0.1, which would mean that 68% of noise values

would be between 13% and 40% of the magnitude of the set-point signal. The maximum scale factor used was 2.0, since this was found to be the value at which the fitting performance of all controller types deteriorated significantly.

As mentioned, each controller structure had an assigned Variable of Interest (c.f. Section 2.2.2), which was varied between the ranges of 0.1 and 0.8. This range was selected to enable the obtention of sufficient data for the derivation of robust conclusions. Furthermore, early experimentation with the model suggested a propensity for system identification to break down more frequently when the variables of interest were assigned values of more than 0.8; as such, this provided an appropriate upper limit. Finally, the algorithm used to estimate the initial parameter values was varied as outlined in Section 2.3.4. Note that only one estimation algorithm was found to be appropriate for the unstable system, whereas four were found to be suitable for the stabilised pendulum.

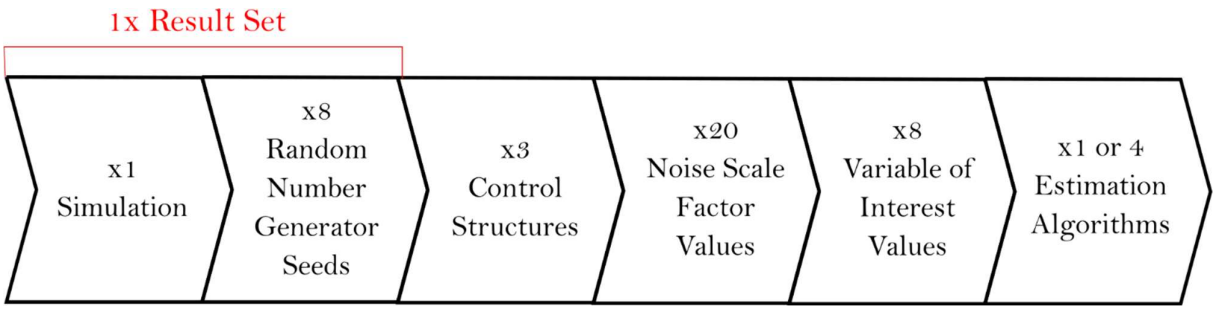

Figure 4: A succinct explanation of how the different parameters were varied between sets of simulations.

2.2.4 – Post-processing

Once the polar distances had been obtained, further refinement was required in order to definitively gauge whether intermittently controlled systems enabled more accurate fitting of transfer functions. Additionally, clarification was sought as to the conditions under which each of the controllers performed best.

Prior to the analysis of trends, anomalous polar distances were filtered out of the dataset. Initially, a value was designated as an anomaly if it sat outside three mean absolute deviations of the dataset median. However, since the fitting methods implemented would occasionally converge to values orders of magnitude greater than the next highest value, they would skew the statistical parameters to the extent that this was not always a valid identifier; thus, it became necessary to implement a manual threshold to counteract this. A threshold of fifty units of polar distance was then additionally imposed, and this was effective in filtering out any extraneous fits.

The resultant refined data was incredibly rich in information, so the analysis occurred in several stages with each step focused on a different aspect. First, it was determined if any of the controller setups enabled accurate identification in a way that was more robust to noise. To do this, the median polar distance of a dataset was

plotted against the noise scale factor for each of the three controllers. If the median polar distance increased more rapidly with noise amplitude for one controller type, it would be suggestive of a lower noise robustness. The results of this analysis are discussed in Section 3.1.

Secondly, the influence of parameter estimation algorithm on identification accuracy was determined by conducting a Kruskal-Wallis test. For each type of controller, the data would be separated into categories depending on the initial estimation algorithm that was implemented (c.f. Section 2.3.4). A Kruskal-Wallis test would then be conducted in order to determine whether there was a statistically significant difference between the quality of fits obtained when each algorithm was used. The Kruskal-Wallis test was selected over ANOVA because it could not be presumed that the data would be normally distributed. The results of this analysis are discussed in Section 3.2.

The influence of the open-loop interval on identification accuracy for intermittent controllers was then assessed. Due to its close relationship to threshold in event-driven control, the analysis of the two variables was combined in Section 3.3. This was accomplished by plotting several graphs showing the variation of median polar distance with both threshold and median open-loop interval. The impact of time delay on continuous controller identification accuracy was similarly deduced by plotting median polar distance against the applied time delay. The results of this analysis are displayed in Section 3.4.

Finally, a comprehensive analysis was done to determine if intermittent controllers enabled the production of statistically more accurate fits than the PCC controller. This was accomplished by ranking all the fits for each controller type from the lowest median polar distance to the highest. The intermittent controller distance was then subtracted from the continuous polar distance to obtain the accuracy deficit D (c.f. Equation 4). A positive accuracy deficit would indicate that the poles of the models obtained using intermittent controllers sat closer to the true poles than those fitted to data where a continuous controller was used (whereas a negative deficit would prove the reverse). A one-sample t-test was then done to determine if the accuracy deficit was statistically significantly different from zero, which would provide evidence of a consistent difference between fit quality achieved when implementing the two controller types.

$$D = \hat{p}_c - \hat{p}_i \qquad \text{Equation 4}$$

2.3 – Details of methodological aspects

2.3.1 – Conversion between stable and unstable systems

The Model consists of a second-order state-space system based on the dynamics of an inverted pendulum. The equations underpinning the pendular motion are based on classical Newtonian mechanics, so are not directly

relevant to a report centred on control. Nonetheless, explanations of their derivation are available in the original paper (Álvarez and Doublein, 2021). Reconfiguration of the dynamic equations into state-space form yields:

$$\dot{x} = \boldsymbol{A}x + \boldsymbol{B}u$$
$$y = \boldsymbol{C}x + \boldsymbol{D}u$$

Equation 1

In the case of the original (unstable) system, the matrices are as below:

$$\boldsymbol{A} = \begin{bmatrix} -\frac{V}{J} & \frac{mgh(1-c)}{J} \\ 1 & 0 \end{bmatrix}, \boldsymbol{B} = \begin{bmatrix} \frac{cmgh}{J} \\ 0 \end{bmatrix}, \boldsymbol{C} = [1 \quad 0], \boldsymbol{D} = 0$$

Here, as in the original paper, V is the ankle viscosity, J is the inertia of the pendulum, c is the ratio between tendon stiffness and load stiffness, h is the pendulum length, m is the pendulum mass, and g is the acceleration due to gravity. The constant values used in the simulations are provided in Appendix A. In order to stabilise the system, both values in the top row of the $\boldsymbol{A}$ matrix were made negative; this would ensure that all eigenvalues of the matrix fall in the left half plane, which is the criterion that governs a stable dynamic response. As such, in the stabilised version, the matrices are as follows:

$$\boldsymbol{A} = \begin{bmatrix} -\frac{V}{J} & -\frac{mgh(1-c)}{J} \\ 1 & 0 \end{bmatrix}, \boldsymbol{B} = \begin{bmatrix} \frac{cmgh}{J} \\ 0 \end{bmatrix}, \boldsymbol{C} = [1 \quad 0], \boldsymbol{D} = 0$$

2.3.2 – Addition of noise and disturbance to the system

The implementation of protocols to simulate sources of noise was central to this study. This is because, as has been noted in the literature, systems that are noise-free and void of nonlinearities are substantially easier to identify (Forssell and Ljung, 1999), and do not reflect the properties of real systems.

All the simulations in this study included external inputs that simulated the action of sensor noise. In Figures 1 and 2, this is denoted by $v_y(t)$. This was done by adding a zero-mean Gaussian white noise value to each datapoint of the output timeseries ($y(t)$ in the Figures 1 and 2). In MATLAB R2021b, a pseudo-random number can be drawn from a zero-mean normal distribution with a standard deviation of 1 using the "randn()" function. High and low noise levels were distinguished through multiplication of the randomly generated value by a scale factor. Thus, the sensor noise can be described mathematically as:

$$y_j'(t_i) = y_j(t_i) + A * v_{y,j}(t_i) \qquad 0 \leq i \leq \frac{T}{h}$$

Equation 5

Where $y_i'(t)$ represents the noisy output, $y(t)$ represents the original output value, A is the scale factor, $v_y(t)$ is the pseudo-randomly generated noise value, T is the total simulation time, and h is the sampling time. Here,

the subscript "j" is used to denote the different random number generator seeds, which will alter the noise characteristics across simulations.

2.3.3 – Model fitting

Among the outputs of the Model were the control signal ($u(t)$ in Figures 1 and 2) and the system output ($y(t)$). Thus, without integrating any further knowledge about the structure of the controller, the direct approach to system identification can be used to convert this input and output data to a transfer function. The coefficients of the transfer function (a, b, c, and d in Equation 2, reprinted below) were designated as the fitting parameters, and the methods built into the MATLAB "tfest()" function were used to converge to a set of optimal values. In this case, the order of the plant was presumed to be known, so a parametric method of fitting was appropriate. In future studies, it may be beneficial to establish the facilitative value of IC in contexts where the order of the plant transfer function is unknown; however, given that this study is a proof-of-concept, it is reasonable to use a known plant structure since this will allow for easier comparison of the fits obtained from IC and PCC.

$$\frac{Y(s)}{U(s)} = \frac{d}{as^2 + bs + c} \qquad \text{Equation 2}$$

As outlined, the fitting of transfer functions to the input and output data was done using the MATLAB "tfest()" function (MathWorks United Kingdom, 2022a). By default, the inputs to the function were the desired number of poles and zeros (2 and 0 respectively in this example). Additionally, one could specify the algorithm used for initial estimation of the transfer function coefficients: these included the Instrument Variable (herein IV) method, the State Variable Filters (SVF) method, the Gauss-Poisson Moment Function (GPMF) approach, the Numerical algorithms for System Identification (N4SID) method, and a Frequency Domain (FD) approach which is an inbuilt MATLAB algorithm based on a prediction-error approximation model (Ljung, 2003). Only the last method necessitated conversion of the input-output dataset to the frequency domain; the others operated exclusively from the timeseries data, though each implemented a degree of prefiltering. A full exploration of the distinctions between these methods is outside the scope of this report, however more detail is provided in the MATLAB "tfest()" function documentation (MathWorks United Kingdom, 2022a). Following the initial estimation of the model parameters, the "tfest()" function would iteratively re-estimate these parameter values using a mixture of four least-squares gradient descent methods. Again, a full exploration is not relevant to this study, but more information is available in (MathWorks United Kingdom, 2022b)

Originally, this study intended to implement all five algorithms for both the stabilised and unstable systems. However, following preliminary studies, it became rapidly apparent that not all initialisation algorithms provided accurate fits for this system. For example, parameter initialisation using the N4SID approach provided significantly worse fits than all other fitting approaches, to the extent that it obscured trends that were otherwise evident (c.f. Figure 5). Similarly, when the unstable system was used, only the FD estimation algorithm enabled the obtention of sufficiently accurate fits to warrant post-processing.

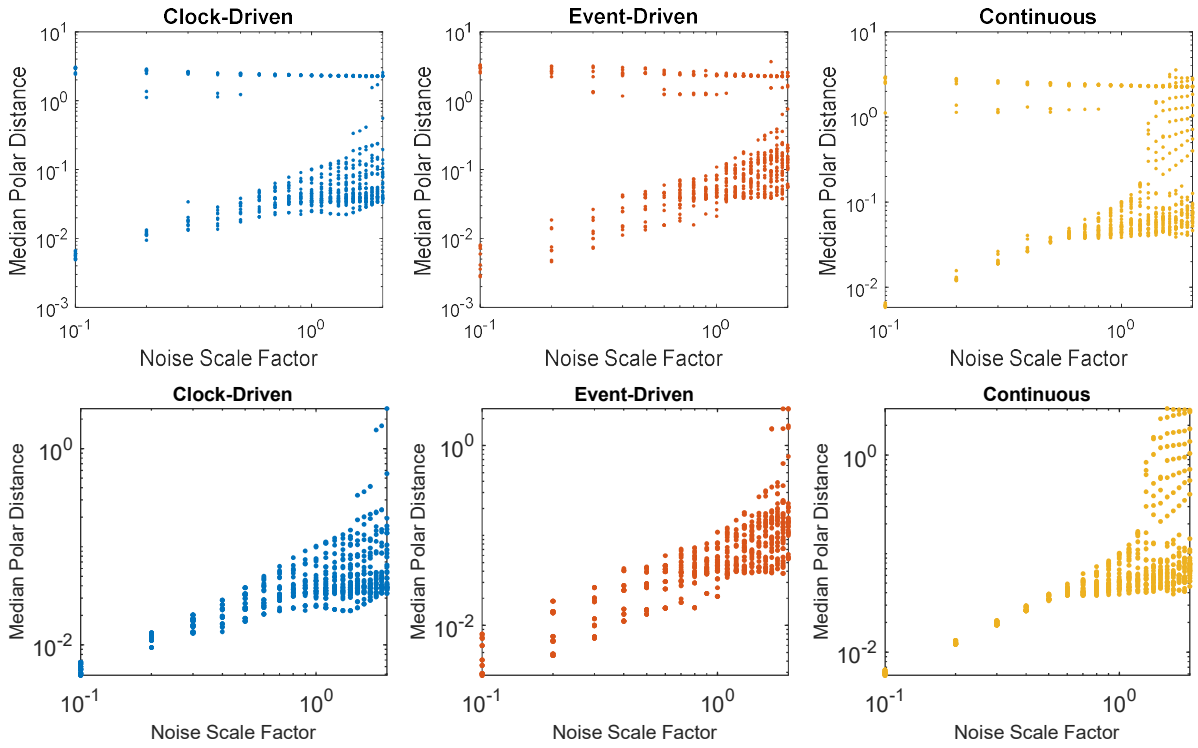

Figure 5: A set of plots justifying the exclusion of the N4SID dataset during post-processing. The top row of subplots shows the unaltered data: note the stray line of anomalous points at an almost constant value of median polar distance. After removal, the bottom set of plots are obtained, with a much more easily discernible trend.

2.3.4 – Comparing fit quality – minimum polar distance

As previously emphasised, this project was centred around the identification of transfer functions from input and output data. Specifically of interest was the proximity of these fitted transfer functions to the true transfer function of the system. However, as of this writing, there is no standard protocol to be found in the literature for likeness comparison of transfer functions. Thus, prior to running the simulations, a suitable metric of similarity had to be derived.

Of course, it is important to note that the superficial features of the function (such as the coefficient values) are less important than what the transfer function represents i.e., the dynamic behaviour of the system. Hence, what is desired is a means of comparing the similarity of two transfer functions *in the context of the dynamic behaviour of the system*. Furthermore, the metric of choice should be independent of whether the model was fitted in the time-domain or the frequency domain since both methods would be used in the fitting process (c.f. Section 2.3.3).

The system being modelled was also very simple: it had only two poles which were conjugate complexes of one another, meaning that there was no pole that dominated the dynamic behaviour of the response. Additionally, the lack of zeros further simplified comparison: since zeros would also influence the dynamic behaviour of a system (Morud and Skogestad, 1996), their absence means that the system dynamics can be

inferred primarily by examining the poles. Thus, on the basis of these factors, it is reasonable to designate proximity of the transfer function poles as a proxy for similarity in dynamic behaviour.

In MATLAB, this was implemented by finding the absolute distance between each pole resulting from a fit, and each true pole of the transfer function. The minimum polar distance from each true pole would be found for each fitted pole, and the sum of these would constitute the similarity metric for the study. This is shown schematically in Figure 6.

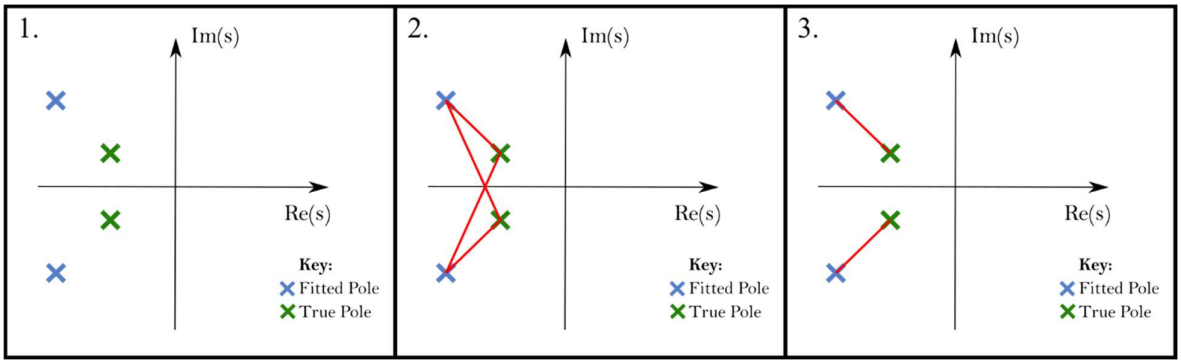

Figure 6: A schematic representation of the procedure for finding minimum polar distance. The influence of polar position on dynamic properties is outlined in (Morud and Skogestad, 1996). Thus, it is reasonable that greater proximity to the true poles is indicative of greater similarity in dynamic properties.

Mathematically, the distance between a fitted pole and a true pole can be written as:

$$d = |\hat{p} - p| \qquad \text{Equation 6}$$

Here, the letters designate a complex value plotted on a two-dimensional plane: $\hat{p}$ indicates a fitted pole, and p indicates a true pole. Thus, for the n-th fitted pole, one must calculate the distance to each of the m true poles, hence:

$$d_{n,m} = |\hat{p}_n - p_m| \qquad m = 1,2 \qquad \text{Equation 7}$$

Hence, the minimum polar distance for the n-th fitted pole can be written as:

$$d_{min,n} = min\{d_{n,m}\} \quad \text{m} = 1,2 \qquad \text{Equation 8}$$

Finally, the indicator of fit quality, i.e., the sum of the minimum polar distances can be denoted by Q, hence:

$$Q = \sum_{n=1}^{2} d_{min,n} \qquad \text{Equation 9}$$

Section 3: Results and Discussion

3.1 – Impact of noise on fit quality

3.1.1 – Stabilised pendulum

Figures 7a, 7b, and 7c show the impact of noise on fit quality for the three different types of controller structure in the case of the stabilised pendulum. None of the three control systems are able to completely mitigate the deterioration in fit quality caused by the increasing amplitude of the noise, and this is reflected in the positive gradients of the regression lines of the three graphs. However, as is evident from Table 1, clock-driven intermittent control exhibits a substantially more gradual deterioration in fit quality compared to event-driven and continuous control. Furthermore, the confidence interval displays no overlap with either of the other two control structures, suggesting its superiority in this regard is statistically significant.

Interestingly, the high gradient of the event-driven regression line would suggest that it is the least robust to noise, even when compared to continuous control. However, this interpretation seems superficial when one contrasts Figures 7b and 7c, since a visual inspection shows a marked decay in fit quality for the continuously controlled system above noise scale factors of 1.3. This is further reinforced by the confidence intervals displayed in Table 1, which are overlapping for event-driven and continuous control. As such, the results of these simulations are insufficient to validate event-driven control over continuous control.

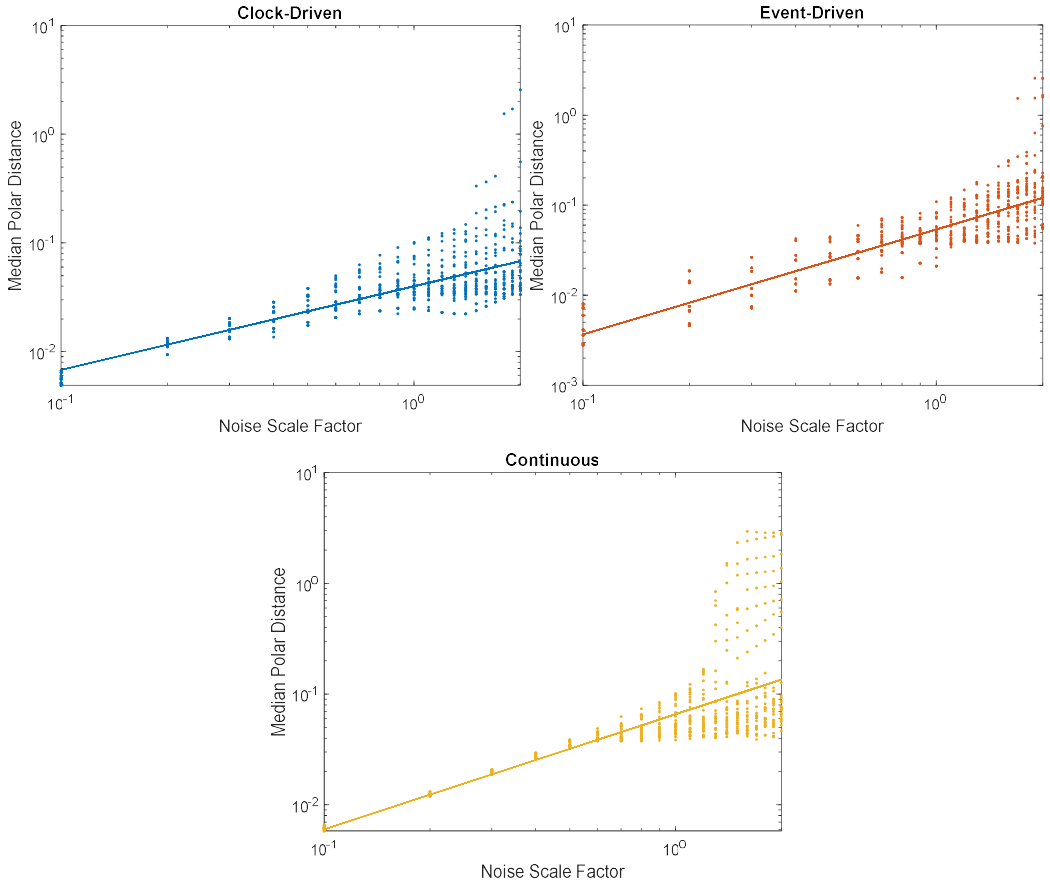

Figure 7a (Top Left), 7b (Top Right), and 7c (Bottom): Logarithmic plots showing the impact of the noise scale factor on median polar distance for the stabilised pendulum.

Table 1: A tabulation of the coefficients of the least-square regression lines for each of the datasets in Figures 7a, 7b, and 7c (above).

	x^1 coefficient	Lower Bound x^1 CI	Upper Bound x^1 CI
Clock-Driven	0.769	0.7196	0.8178
Event-Driven	1.164	1.1073	1.2209
Continuous	1.039	0.9624	1.1158

Nonetheless, it is worth contemplating why fit accuracy of the event-driven controller is significantly less robust to noise than that of the clock-driven controller. This report believes that it is an unfortunate by-product of the increased computational efficiency of the event-driven controller.

As mentioned, event-driven controllers only update the control trajectory if the error exceeds the threshold set by the operator and if their minimum time interval has elapsed. By contrast, a clock-driven controller will enact a new control trajectory after every t_i, regardless of the amount of correction required. When noise is added to the system, this skews the controller's perception of the state of the plant. This means that, for an event-driven controller, the impact is twofold: the control action will be faulty (because the $\hat{x}_h$ vector being acted on by the state feedback controller will be affected), and the system will remain in the open-loop state for longer. Therefore, the performance of the controller is impeded in two ways. By contrast, the clock-driven controller will initiate new control trajectories independently of the output, so its performance is only hindered to the extent that the $\hat{x}_h$ vector is affected. This twofold impact of noise on event-driven controllers could explain why their ability to facilitate accurate fitting is less robust to noise.

3.2.2 – Unstable pendulum

The fits obtained for the unstable system provide a decidedly less compelling case for the noise robustness of IC systems. The deterioration of fit quality with increasing noise scale factor is more pronounced, and this is reflected in the augmented gradients for all controller types. This is most evident in the case of event-driven control, where the gradient increases by more than 40% (1.680 for the unstable system compared to 1.164 for the stabilised system). Interestingly, Figure 8b shows that the linear fit is much sparser for the unstable case. Furthermore, the orderly fanning out of datapoints at higher scale factors is absent, replaced by a more uneven pattern. Both these factors provide support for the suggestion that EDC is less robust to noise than the other two control systems. A potential explanation for this was provided in the previous section.

Additionally, as shown by the overlapping confidence intervals in Table 2, there is no statistically significant difference in gradient between the least-squares lines of clock-driven and continuous control. This contradicts the earlier finding and suggests that, in the context of unstable systems, the use of an intermittent controller provides no robustness advantage for system identification in the presence of noise. However, it should be noted that, as outlined in Section 2.3.3, only one parameter estimation algorithm proved effective for unstable systems. Hence, one cannot guarantee that these trends hold true for other estimation methods, so these results should be considered with caution.

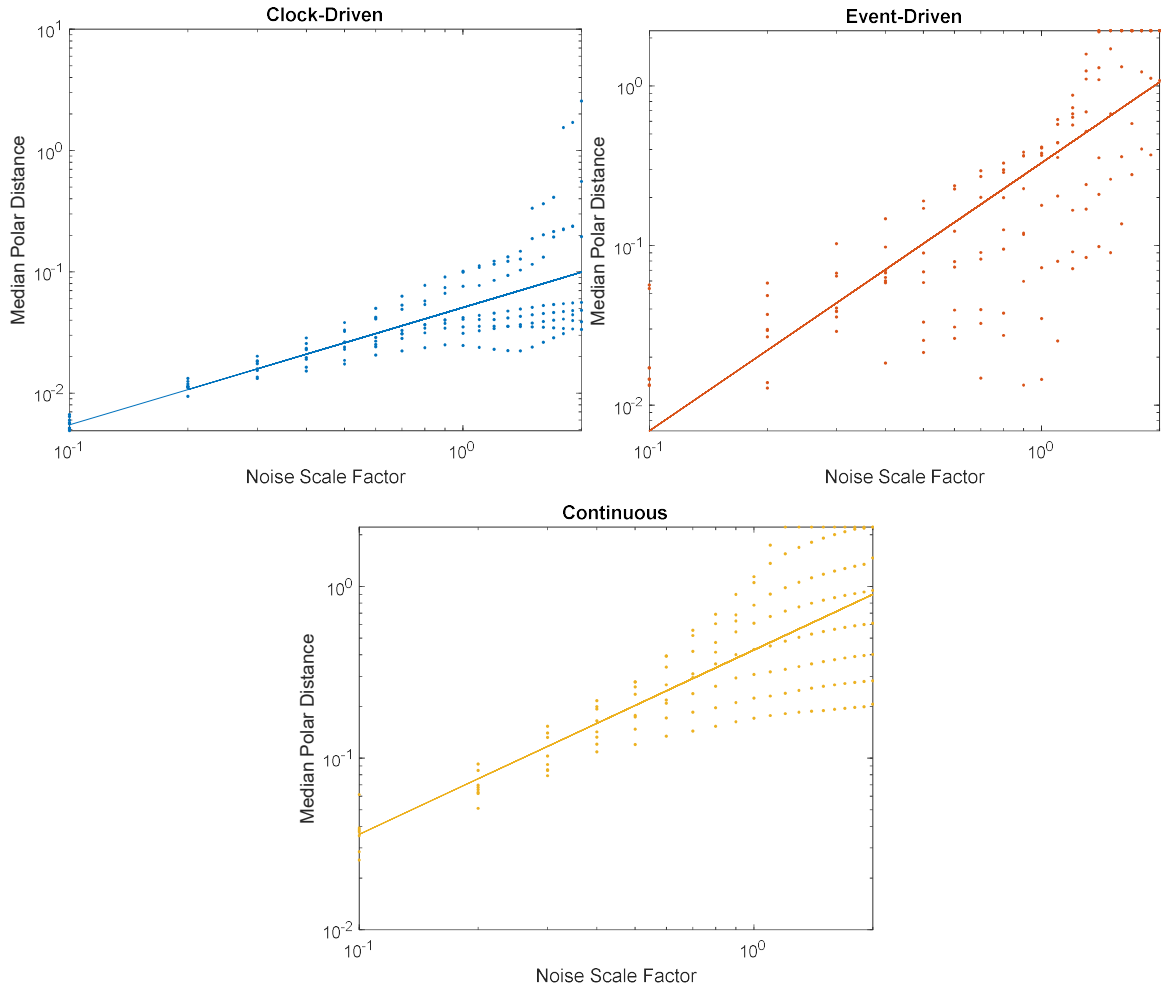

Figure 8a (Top Left), 8b (Top Right), and 8c (Bottom): Logarithmic plots showing the impact of the noise scale factor on median polar distance for the unstable pendulum.

Table 2: A tabulation of the coefficients of the least-square regression lines for each of the datasets in Figures 8a, 8b, and 8c (above). Note the contrast when compared to Table 1: there is no statistically significant boost to noise robustness with the implementation of a clock-driven controller.

	x^1 coefficient	Lower Bound x^1 CI	Upper Bound x^1 CI
Clock-Driven	0.967	0.823	1.120
Event-Driven	1.680	1.477	1.883
Continuous	1.074	0.941	1.206

3.2 – Impact of initial estimation algorithm: stabilised pendulum only

The results of the Kruskal-Wallis tests show that there are almost no statistically significant differences in efficacy between estimation algorithms for the stabilised pendulum model. As shown in Table 3, the only fitting method which could be considered to fall short is the FD approach and this is especially notable when implemented alongside a continuous controller. This is starkly evident in Figure 9c, which shows the vast

range of polar distances produced by fits of the continuously controlled system. However, it is challenging to pin down why this is the case: across all sets of simulations, the FD approach was far more prone to anomalies than were the other algorithms. However, the breakdown of the FD approach in the continuous controller is distinct, with its interquartile range dwarfing even the full ranges of polar distances produced by the other estimation methods (c.f. Figure 9c).

The reason for this discrepancy is challenging to isolate due to the large number of interacting variables. However, this report believes that, perhaps, the performance of the FD estimation method could be ameliorated by the use of a periodic input; as noted in Appendix A, the input was entirely aperiodic, and was selected on the basis of its capacity to produce high-quality fits during preliminary trials. However, an examination of the literature reveals that periodic inputs are often considered ideal when working with frequency-domain identification methods (Ljung, 2006). This, combined with the lack of significant difference between the alternative initialisation algorithms, means that this report is not positioned to provide conclusive recommendations for selection of estimation method.

Table 3: A table comparing the fit accuracy of the different estimation methods for each of the controller setups. The numbers in each of the cells represents the p-value. Note that the reading should be done starting with the row e.g., the cell marked * should be read as "Frequency-domain identification provided statistically less accurate fits than the instrument variable approach, with a p-value of 0.032".

	Clock-Driven				Event-Driven				Continuous			
	FD	GPMF	IV	SVF	FD	GPMF	IV	SVF	FD	GPMF	IV	SVF
FD	X	0.297	0.310	0.300	X	0.256	0.032*	0.315	X	0.000	0.000	0.000
GPMF	0.297	X	1.000	1.000	0.256	X	0.808	0.999	0.000	X	0.830	0.998
IV	0.310	1.000	X	1.000	0.032	0.808	X	0.740	0.000	0.830	X	0.729
SVF	0.300	1.000	1.000	X	0.315	0.999	0.740	X	0.000	0.998	0.729	X
Key		Statistically more accurate				Statistically less accurate				No statistical difference		

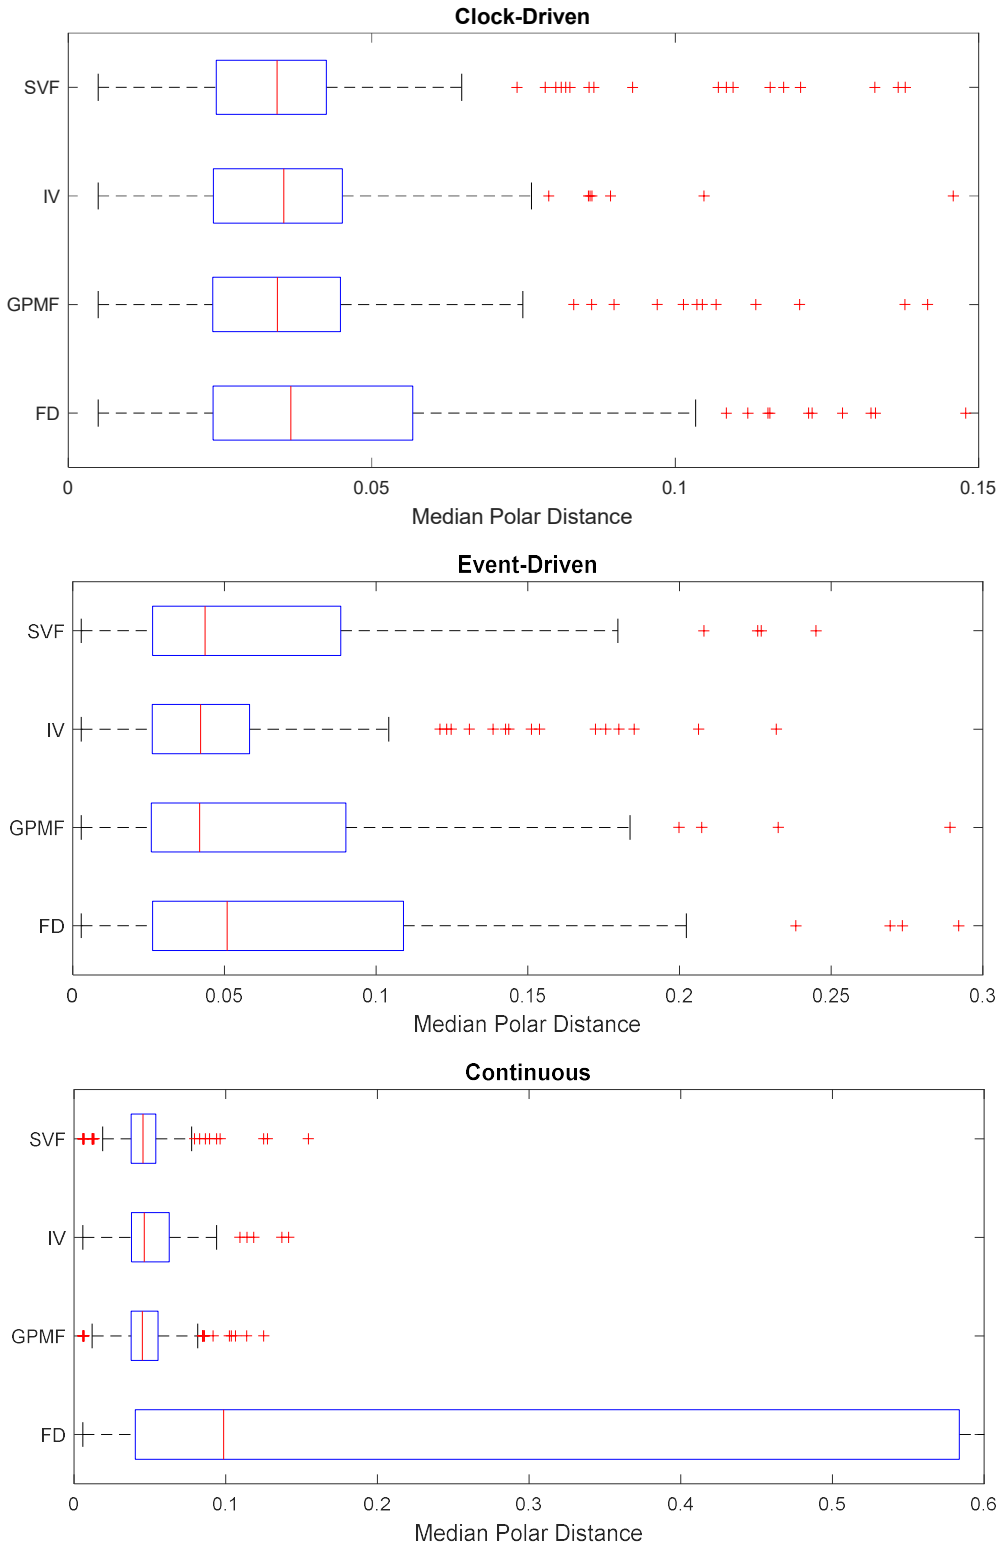

Figures 9a (Top), 9b (Middle), and 9c (Bottom): Figures 9a, 9b, and 9c show a set of boxplots of the median polar distances obtained when implementing each of the different parameter initialisation algorithms for the stable system. All diagrams have been truncated for ease of reading; in every case, the FD approach exhibited a propensity to produce models with abnormally high polar distances, with the largest for all controller types nearing 3.0. Note the differing scales on the horizontal axes.

3.3 – Impact of open-loop interval on fit quality: stabilised and unstable pendulums

Figures 10a and 10b show the influence of median open-loop interval on fit quality for the stable pendulum. Both systems clearly display a propensity for fit quality to improve with larger values of t_i, as shown by the negative gradients of the regression lines. This trend was very similar for the unstable system, so the graphs for the unstable system have not been included here (however they are available in Appendix B). Furthermore, Table 4 demonstrates that this improvement in fit quality exceeds the threshold for statistical significance, since the 95% confidence intervals do not extend to zero or positive gradients. Interestingly, the decline is more salient in the case of EDC which one can note by contrasting Figures 10a and 10b, and this is even more pronounced in the unstable system (c.f. Figures 18a and 18b in Appendix B).

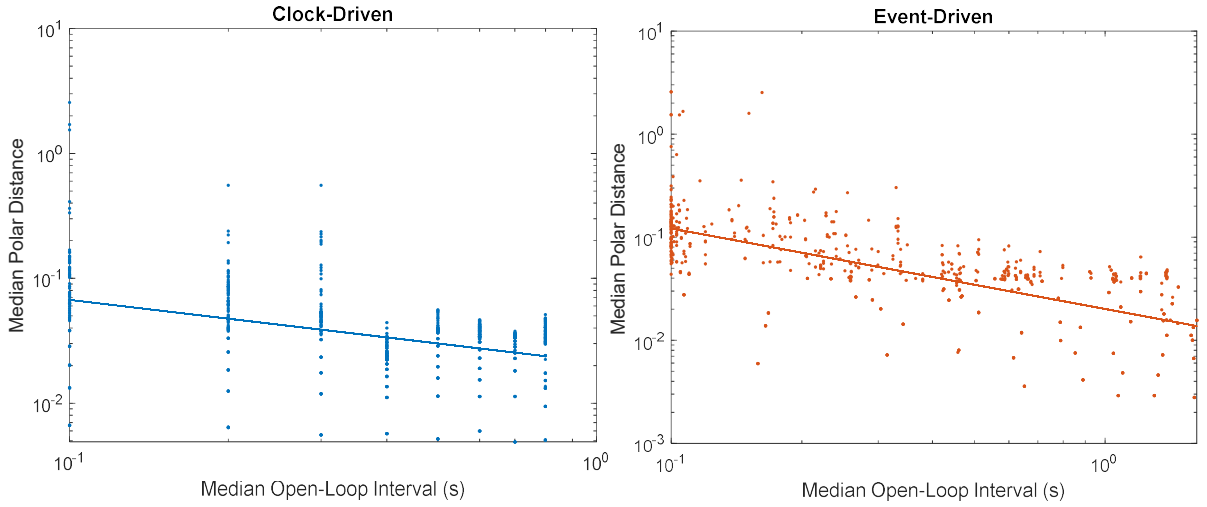

Figure 10a (Left), 10b (Right): Logarithmic plots showing the impact of median open-loop interval on median polar distance for the stabilised pendulum. Note that the open-loop intervals in clock-driven control are evenly spaced since t_i for any single controller is constant.

Table 4: A tabulation of the regression line gradients for Figures 10a and 10b (above), with confidence intervals.

	x^1 coefficient	Lower Bound x^1 CI	Upper Bound x^1 CI
Clock-Driven	-0.500	-0.585	-0.415
Event-Driven	-0.781	-0.849	-0.713

The fact that longer open-loop intervals tend to elicit superior fits is reasonable from the earlier theory: since the direct approach is more suitable for open-loop systems (Forssell and Ljung, 1999), the longer that an IC system spends in an open-loop state, and the more effective the DA should become. However, the question remains as to why this improvement is more pronounced for EDC systems.

The explanation for this observation could be linked to control quality: in CDC, improvements in fitting resultant from augmented t_i will be offset by the accompanying deterioration in control quality, and this leads to the shallower gradient seen in Figures 10a and 18a. Indeed, if one examines Figure 11a, it is visible that there is a dramatic increase in median polar distance as control quality deteriorates. By contrast, this is not necessarily the case in event-driven control, since the system is remaining in the open-loop state for the longest possible time before the error threshold is exceeded; thus, longer open-loop intervals need not necessarily come

at the expense of control quality. Overall, this longer open-loop interval without loss of control quality could enable more accurate estimates to be obtained.

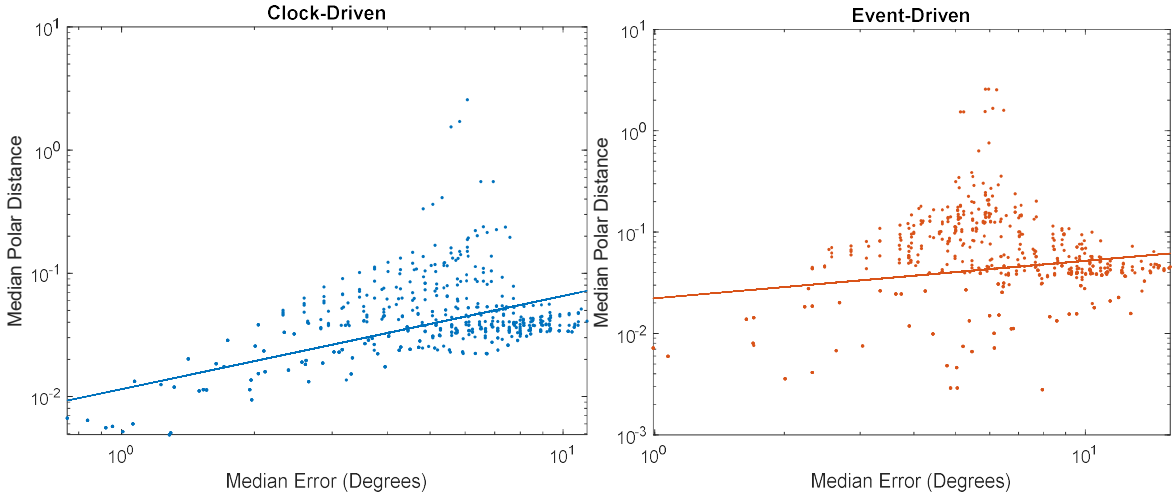

Figure 11a (Left) and 11b (Right): Logarithmic plots showing the relationship between median error and median polar distance. Note that, for CDC, as the median error increases (i.e. control quality decays), the median polar distance increases. However, this deterioration in fit quality is not matched in EDC: the influence of control quality appears to be less significant. This may be because higher thresholds lead to both worse control quality and longer open-loop intervals, so the detriment to fitting quality of the former offsets the benefit of the latter.

Table 5: A tabulation of the regression line gradients for Figures 11a and 11b (above), with confidence intervals.

	x^1 coefficient	Lower Bound x^1 CI	Upper Bound x^1 CI
Clock-Driven	0.757	0.677	0.837
Event-Driven	0.369	0.211	0.528

This hints at a more nuanced view of the original inference that larger t_i facilitates system identification; longer open-loop intervals do allow for better system identification, provided that they do not undermine the quality of control. However, this conjecture is not directly provable since it is not possible to extricate the impact of each of these interlinked variables. For example, if it is assumed that open-loop interval does not directly affect control quality, there should ideally be no association between the two variables. Despite this, as shown in Figure 12, there exists a positive association between median error and open-loop interval. Without context, it is tempting to assume a causal link. Nonetheless, it is critical to remember that threshold value can act as a confounding variable: high threshold values lead to both poorer quality of control and longer open-loop intervals.

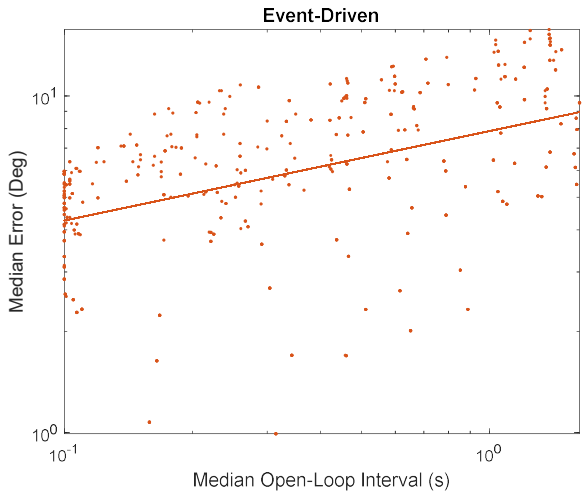

Figure 12: A logarithmic plot showing the relationship between median open-loop interval and median error. If taken at face value, this graph contradicts the reformulated explanation, since it could suggest a decay in control quality with longer open-loop intervals. This, by extension, would imply that loss of control quality is not to blame for the less dramatic improvement of fit accuracy of CDC systems with increased t_i. However, this neglects the existence of threshold value as a confounding variable.

Table 6: A tabulation of the regression line gradient for Figure 12 (above), with confidence intervals.

	x^1 coefficient	Lower Bound x^1 CI	Upper Bound x^1 CI
Event-Driven (Error/t_i)	0.266	0.227	0.305

This new explanation is again validated when one examines the impact of threshold value on median polar distance (c.f., Figure 13 and 18c). High values of threshold have a less pronounced impact on fit quality than median open-loop interval (as shown by the shallower gradient). Framing this in the context of the reformulated explanation, it could be that this is because the improvement arising from longer open-loop intervals is offset by the accompanying decay in control quality.

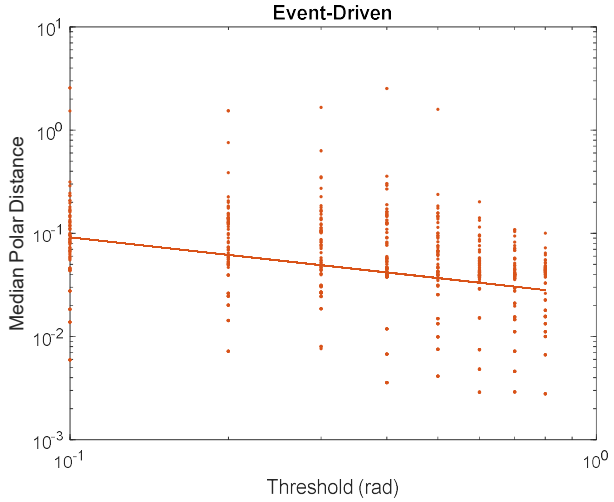

Figure 13: A logarithmic plot showing the relationship between median polar distance and threshold.

Table 7: A tabulation of the regression line gradient for Figure 13 (previous page), with confidence intervals.

	x^1 coefficient	Lower Bound x^1 CI	Upper Bound x^1 CI
Event-Driven	-0.567	-0.688	-0.446

3.4 – Impact of time delay on PCC fit quality: stabilised and unstable pendulums

The relationship between time delay and fit quality for the stabilised pendulum is shown in Figures 14a and 14b. As can be seen, the variation in median polar distance is highly dependent on estimation method: for ease of reading, the results obtained with the FD algorithm have been separated from the rest. Starting with Figure 14a, it is clear that the increase of time delay not only improves fit quality but also dramatically reduces the spread of results. Indeed, the range decreases from approximately 3.0 at a time delay of 0.1s to less than 0.4 at a time delay of 0.8s. By contrast, the improvement in fit and reduction in result spread is significantly less evident in all other initialisation algorithms (c.f. Figure 14a): the decrease in range is irregular, and the decrease in minimum polar distance is modest (though is, nonetheless, statistically significant with a p-value of 4.06E-08).

This result is surprising since, as mentioned, no parameter was available in the model to allow for the introduction of a time-delay. As such, given the absence of such a provision, the quality of fitting should have decayed significantly as its magnitude was augmented. Additionally, the quality of control deteriorated with increasing time-delay: this is expected but further cements the impression that the quality of fitting should have suffered as the time delay increased. However, the difference in behaviour between the FD algorithm and the timeseries initialisation algorithms points towards an explanation based on the different influences of time delay in the time and frequency domains. This is supported by the comparatively strong performance of the FD algorithm for EDC and CDC systems (as covered in Section 3.2). However, as of this writing, the author has been unable to find a satisfactory explanation in the literature to justify this conjecture.

Nonetheless, the author believes that this trend merits further exploration, since the quality of fitting of the continuous system at high time-delays occasionally rivalled that obtained when the intermittent control systems were implemented. Nonetheless, this is covered more thoroughly in the next section.

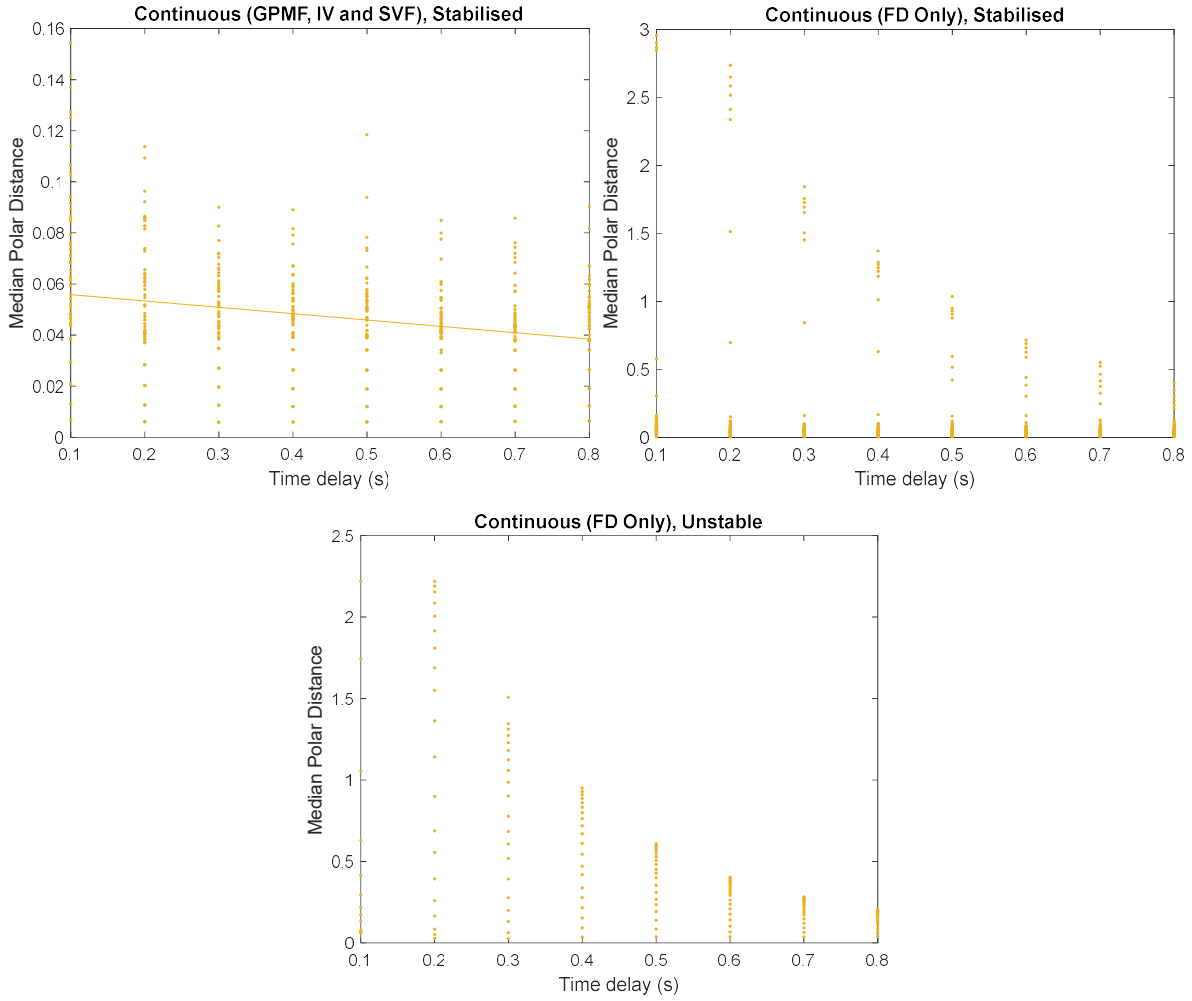

Figure 14a (Top Left), 14b (Top Right), and 14c (Bottom): Figures 14a and 14b show the variation in median polar distance with increasing time delay for the stabilised system. The separation into two plots was made necessary by the unusual behaviour of the frequency-domain estimation method (note the differences in scale of the vertical axes). For the unstable system, only the FD approach was able to consistently identify the plant transfer function to a good degree of accuracy; the impact of time delay on the median polar distance is shown in Figure 14c.

3.5 – Comprehensive analysis: intermittent control vs predictive continuous control

3.5.1 – Stable System

Based on the available data, there is good evidence to suggest that clock-driven intermittent controllers can act as facilitators to the direct approach to system identification for stable systems. This is most clearly visible in Table 8: for every method of system identification, the median polar distance was statistically smaller for fits of pendulums controlled by CDC systems than for continuously controlled pendulums. This remained true when data from all fitting methods was combined, as shown by the positive accuracy deficit of 0.065.

By contrast, EDC systems lacked the definite facilitative value of their clock-driven counterparts. Upon initial inspection of Table 8, it may appear that they can enable more accurate system identification, since the aggregated data shows a positive accuracy deficit (suggesting that the fitted transfer function poles were closer

to the true poles if an event-driven controller was implemented compared to a continuous controller). However, closer inspection shows that this result is dominated by the polar distances fitted using the FD method; those obtained using any other algorithm demonstrated that continuously controlled pendulums produced closer fits than those controlled by EDC systems. Indeed, when one recalculates the mean and p-value in the absence of the FD data, the continuous controller is proven to enable statistically significantly better fits, with a mean accuracy deficit of -0.0112 and a p-value of 1.46E-17.

Figure 15 provides a good summary of these results: the median polar distances obtained using clock-driven control (the blue points) consistently sit closer to zero than both the other datasets. Conversely, the polar distances obtained when implementing an event-driven controller (the red points) are significantly more variable, alternating above and below the line of continuously controlled data. Notably, at the very lowest point on the plot, the fits obtained when implementing an event-driven controller outperform those obtained when implementing any other controller type.

This, combined with the information from Section 3.1.1 (in which the impact of noise was explored), provides an explanation for the contrast in performance of the two intermittent controller variations. As outlined earlier, the ability of event-driven control to produce accurate fits is less robust to noise than that of clock-driven control. This is due to the variation in trajectory intervals combined with the inaccuracy of state estimates output by the observer. This explains why, as the noise level gets higher (shown by larger datapoints in Figure 15), the quality of fits obtained decreases dramatically for pendulums controlled by EDC systems. However, at very low noise levels, the event-driven controller outperforms both other controller types in terms of its ability to enable accurate fits: this is because, unimpeded by excessive noise, the event-driven controller can sustain longer open-loop intervals without the loss of control quality demonstrated by CDC systems (as discussed in Section 3.3).

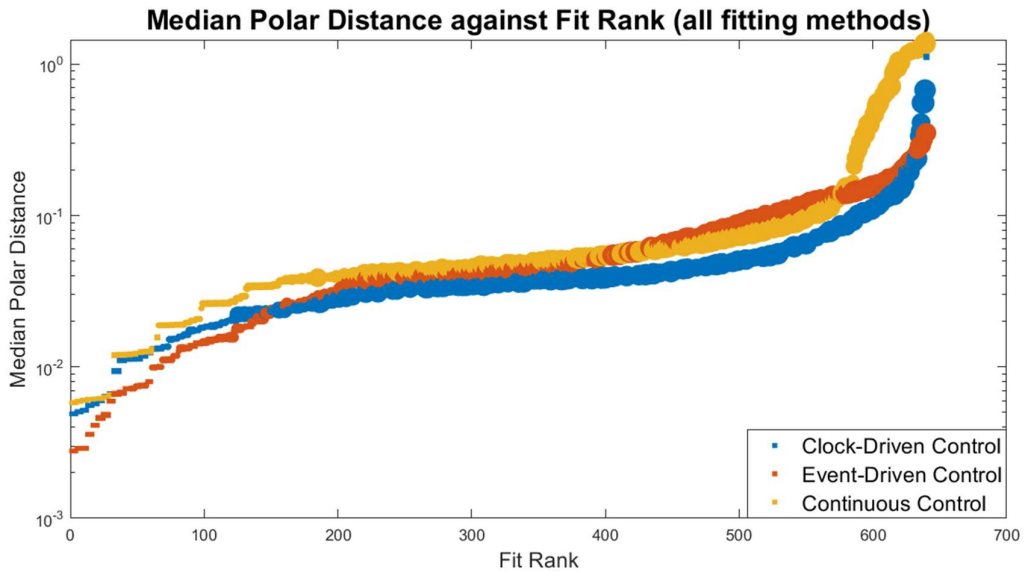

Figure 15: A plot of median polar distance against fit rank for all fitting methods of the stable system (FD, GPMF, IC and SVF). The size of a data point represents its noise scale factor. Note that the vertical axis is logarithmic.

Table 8: A tabulation of the results of the comprehensive analysis: for each fitting method, the mean accuracy deficit (i.e., difference between the continuous and intermittent median polar distances) was calculated and displayed in the rightmost column. The p-values indicate the confidence that the mean accuracy deficit in each case is significantly different from zero. As can be seen, both clock-driven and event-driven intermittent control outperform continuous control when all fitting methods are considered. However, in the case of event-driven intermittent control, this is due to a large difference in the frequency-domain initialisation algorithm. Indeed, when it is removed, the continuous controller is statistically superior to the event-driven controller with a mean accuracy deficit of -0.0112 and a p-value of 1.46E-17.

Clock-Driven			
Estimation Algorithm	**Reject Null?**	**p**	**Mean Accuracy Deficit**
All	Yes	1.95E-17	0.065
FD	Yes	3.82E-13	0.374
GPMF	Yes	1.77E-12	0.006
IV	Yes	1.29E-54	0.011
SVF	Yes	8.38E-07	0.005
Event-Driven			
Fitting	**Reject Null?**	**p**	**Mean Accuracy Deficit**
All	Yes	2.57E-11	0.053
FD	Yes	4.31E-14	0.298
GPMF	Yes	1.31E-08	-0.015
IV	No	1.23E-01	-0.002
SVF	Yes	3.02E-09	-0.016

3.5.2 – Unstable System

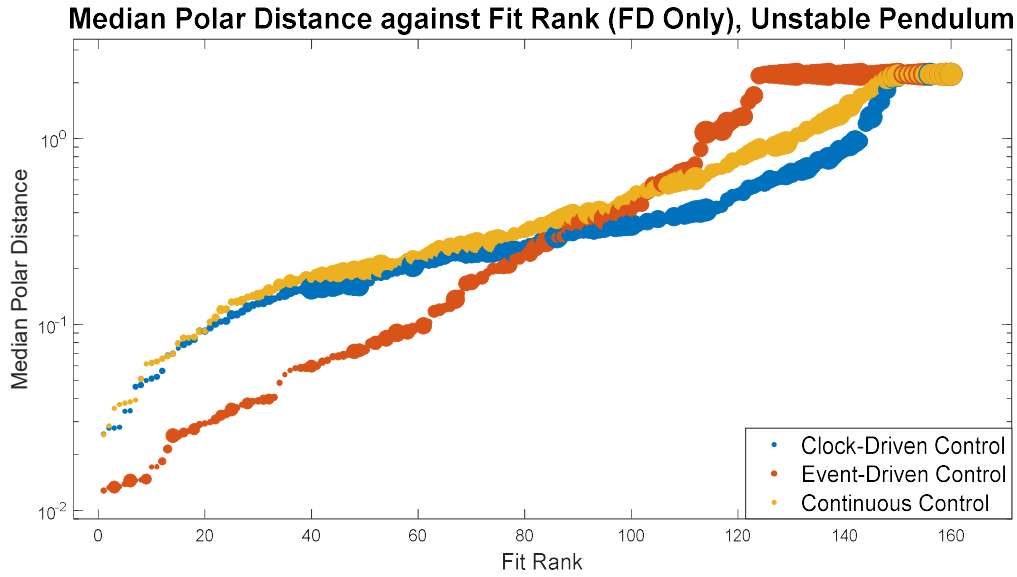

Figure 16: A plot of median polar distance against fit rank for the unstable system. The size of a data point represents its noise scale factor. Note that the vertical axis is logarithmic. Once again, EDC begins as the control structure that enables fitting of the most accurate models, however it is rapidly overtaken by both other control structures.

As shown in Figure 16, the results derived from fitting models to the unstable system are aligned with those obtained from the stabilised pendulum. At the lowest levels of noise, implementation of an event-driven controller enables the best fits to be obtained; indeed, the contrast is even more marked in Figure 16, with the line of EDC fits sitting far below the other two datasets. However, increases in output noise cause a marked degradation of the fit quality, and the event-driven controller is shortly outperformed by both the other two setups. This echoes the results seen in Section 3.1.2, which showed that, when modelling the unstable system, the median polar distance degraded even more rapidly in response to increasing noise, particularly for event-driven control. Thus, the EDC performed statistically worse than PCC as a facilitator of system identification, a fact further demonstrated by the negative mean accuracy deficit in Table 9.

As found when modelling the stable system, the clock-driven controller enabled consistently more accurate models to be fitted than the continuous controller. Table 9 further cements this point by demonstrating that the clock-driven controller enabled distinctly more accurate modelling, and that this improvement far exceeded the threshold for statistical significance ($p = 6.93E\text{-}18$). This again suggests that, although both intermittent controller variations may have facilitative value, this quality is not robust to noise in the case of event-driven control.

Table 9: A tabulation of the results of the comprehensive analysis for the unstable system. Recall from Section 2.3.3 that the only method that consistently produced high-quality fits for the unstable system was the FD approach. Note that, as with the stabilised system, CDC enabled more accurate system identification than PCC, whereas EDC did not.

Clock-Driven			
Fitting	**Reject Null?**	**p**	**Mean Accuracy Deficit**
FD	Yes	6.93E-18	0.121
Event-Driven			
Fitting	**Reject Null?**	**p**	**Mean Accuracy Deficit**
FD	Yes	1.06E-04	-0.126

Section 4: Conclusions and Limitations of the Current Study

Section 4.1 – Limitations

This study had a narrow focus: to determine the facilitative value of IC systems when implementing the direct approach to system identification in the presence of noise. As such, several other factors demand consideration prior to implementation in real system identification problems, especially in the context of human movement. For example, all simulations were done in the absence of disturbance. This was decided because, given the constraints on computing power and time, the impact of disturbance could not be tested sufficiently rigorously to enable the establishment of robust conclusions. However, disturbance features prominently in real-world systems: for example, as outlined by Tanabe et al. in their 2016 study, adaptation to external disturbances is essential to the maintenance of balance during quiet standing. Thus, the results outlined herein may not perfectly reflect the properties of systems subject to disturbance.

However, the most important limitation of this study was that, as a simulation based on empirical dynamic equations, it represents an ideal case for system identification. For example, the number of poles and zeros were known in advance, enabling the fitting of perfectly representative models. This is rarely the case in instances of system identification, and several model types are tested with the risk of over or underfitting. Furthermore, sources of noise and disturbance are not always identifiable or measurable: this is especially relevant when considering implementation of an event-driven controller since it may not be clear whether the level of noise would render it an acceptable choice. Finally, as a fully computerised simulation, all parameters were modifiable, and there was no risk of damage to the system or alteration of system properties. Therefore, in systems that require very fine control, the findings of this study may not be applicable due to the loss of control quality that results from the replacement of a continuous control system with an intermittent one.

Section 4.2 – Conclusions

To conclude, this report believes that intermittent control structures do have facilitative value in the direct approach to system identification, and this is supported by the findings of Section 3.5. Indeed, for all initialisation algorithms, the clock-driven controller enabled the fitting of more accurate models than the predictive continuous controller. The event-driven controller may have some value as an enabler of the direct approach, but its efficacy in this role is confined to low-noise environments since, as established, its ability to facilitate accurate fittings lacks robustness in this regard. These results are consistent for both stable and unstable systems suggesting that, if human movement operates intermittently, the direct approach may constitute a valid means of system identification. Nonetheless, further studies will be required prior to its implementation in this context, as will be discussed in the next section.

Section 5: Opportunities for Further Study

As mentioned previously, the implementation of these findings in real-world applications (especially in the context of human standing) is contingent on the outcomes of future studies. These investigations are needed to confirm the validity of the current findings in alternative contexts, and to devise optimal strategies for the application of intermittent controllers to system identification problems.

Building on this latter point, it would be tremendously beneficial to have a study examining whether there exists an optimal input type for identification using intermittent control systems. Such studies do exist in the literature, although they are focused exclusively on continuously-controlled systems (Narasimhan and Bombois, 2012; Krinitsyn *et al.*, 2016); as such, further validation is required to confirm that the same trends exist for IC systems. This is especially relevant if a frequency-domain method to system identification is being considered since, as outlined in Section 3.2, this study did not use a periodic input so would not be representative of the optimal performance in such contexts.

Furthermore, as more studies examine the influence of controller selection on system identification, it would be tremendously valuable to have standardised measures with which to compare the quality of a model fit. A potential starting point was provided in this study, in the form of minimum polar distance. However, as research is extended to more complex systems, methods will need to be devised that account for the presence of zeros, or poles that have an outsized impact on the dynamic response of the system (sometimes called dominant poles). For the latter, perhaps a weighted sum could be used to represent this disparity in influence. Nonetheless, the development of versatile comparison metrics will be central to the validation of the facilitative value of IC in diverse contexts.

Finally, in order to apply these findings to the study of human movement, several further investigations are required. Firstly, it is necessary to confirm applicability of the current results to systems with multiple inputs and outputs. For example, an examination of the literature on postural control shows that more detailed models of human standing tend to be based on multi-segment inverted pendulums with several inputs (van der Kooij *et al.*, 1999; Alexandrov *et al.*, 2005). As such, validation in this regard will be crucial to the implementation of the direct method in human standing.

Of course, more clarity is also needed regarding the nature of the human movement control system: if human motion is continuously controlled, the facilitative value of intermittent control systems is moot. The literature review of Section 1.3 showed that, although several teams of researchers have worked diligently to solve this problem, the results remain conflicting. Nonetheless, as the volume of published literature in this area grows, it is hoped that the results of this study can serve as a lens through which to view the direct approach in a new light.

Section 6: References

Alexandrov, Av. *et al.* (2005) 'Feedback equilibrium control during human standing', *Biological Cybernetics*, 93(5), pp. 309–322. doi:10.1007/s00422-005-0004-1.

Álvarez, A. and Doublein, T. (2021) 'The inverted pendulum model', p. 3.

Antsaklis, P.J. and Michel, A.N. (2007) *A Linear Systems Primer*. Boston, MA: Birkhäuser Boston (A linear systems primer).

Åström, K.-J. and Bohlin, T. (1965) 'Numerical Identification of Linear Dynamic Systems from Normal Operating Records', *IFAC Proceedings Volumes*, 2(2), pp. 96–111. doi:10.1016/S1474-6670(17)69024-4.

Craik, K.J.W. (1947) 'Theory of the human operator in control systems; I. The operator as an engineering system.', *British Journal of Psychology*, 38, pp. 56–61.

Elias, L.A., Watanabe, R.N. and Kohn, A.F. (2014) 'Spinal Mechanisms May Provide a Combination of Intermittent and Continuous Control of Human Posture: Predictions from a Biologically Based Neuromusculoskeletal Model', *PLoS Computational Biology*. Edited by F.J. Valero-Cuevas, 10(11), p. e1003944. doi:10.1371/journal.pcbi.1003944.

Forssell, U. and Ljung, L. (1999) 'Closed-loop identification revisited', p. 27.

Gawthrop, P. *et al.* (2011) 'Intermittent control: a computational theory of human control', *Biological Cybernetics*, 104(1–2), pp. 31–51. doi:10.1007/s00422-010-0416-4.

Gawthrop, P., Gollee, H. and Loram, I. (2018) 'Intermittent Control in Man and Machine', in Miskowicz, M. (ed.) *Event-Based Control and Signal Processing*. 0 edn. CRC Press, pp. 281–350. doi:10.1201/b19013-14.

Gawthrop, P., Loram, I. and Lakie, M. (2009) 'Predictive feedback in human simulated pendulum balancing', *Biological Cybernetics*, 101(2), pp. 131–146. doi:10.1007/s00422-009-0325-6.

Gawthrop, P. and Wang, L. (2011) 'The system-matched hold and the intermittent control separation principle', *International Journal of Control*, 84(12), pp. 1965–1974. doi:10.1080/00207179.2011.630759.

Gawthrop, P.J. and Wang, L. (2009) 'Event-driven intermittent control', *International Journal of Control*, 82(12), pp. 2235–2248. doi:10.1080/00207170902978115.

Gevers, M. (2003) 'A personal view on the development of system identification 1', *IFAC Proceedings Volumes*, 36(16), pp. 747–758. doi:10.1016/S1474-6670(17)34850-4.

Gollee, H. *et al.* (2012) 'Frequency-domain identification of the human controller', *Biological Cybernetics*, 106(6–7), pp. 359–372. doi:10.1007/s00422-012-0503-9.

Gustavsson, I., Ljung, L. and Söderström, T. (1977) 'Identification of processes in closed loop—identifiability and accuracy aspects', *Automatica*, 13(1), pp. 59–75. doi:10.1016/0005-1098(77)90009-7.

Hinrichsen, D. and Pritchard, A.J. (2005) 'Introduction to State Space Theory', in Hinrichsen, D. and Pritchard, A. J., *Mathematical Systems Theory I*. Berlin, Heidelberg: Springer Berlin Heidelberg (Texts in Applied Mathematics), pp. 73–192. doi:10.1007/3-540-26410-8_2.

Ho, B.L. and Kalman, R.E. (1966) 'Editorial: Effective construction of linear state-variable models from input/output functions: Die Konstruktion von linearen Modeilen in der Darstellung durch Zustandsvariable aus den Beziehungen für Ein- und Ausgangsgrößen', *auto*, 14(1–12), pp. 545–548. doi:10.1524/auto.1966.14.112.545.

Hwang, S. *et al.* (2016) 'Identification of the Unstable Human Postural Control System', *Frontiers in Systems Neuroscience*, 10. doi:10.3389/fnsys.2016.00022.

Iqbal, K. (2017) *A first course in control system design*. Available at: https://go.openathens.net/redirector/umoncton.ca?url=https%3A%2F%2Fsearch.ebscohost.com%2Flogin.aspx%3Fdirect%3Dtrue%26scope%3Dsite%26db%3Dnlebk%26AN%3D1800582 (Accessed: 10 April 2022).

Isermann, R. and Münchhof, M. (2011) *Identification of Dynamic Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/978-3-540-78879-9.

Johansson, R., Magnusson, M. and Akesson, M. (1988) 'Identification of human postural dynamics', *IEEE Transactions on Biomedical Engineering*, 35(10), pp. 858–869. doi:10.1109/10.7293.

van der Kooij, H. *et al.* (1999) 'A multisensory integration model of human stance control', *Biological Cybernetics*, 80(5), pp. 299–308. doi:10.1007/s004220050527.

van der Kooij, H., van Asseldonk, E. and van der Helm, F.C.T. (2005) 'Comparison of different methods to identify and quantify balance control', *Journal of Neuroscience Methods*, 145(1–2), pp. 175–203. doi:10.1016/j.jneumeth.2005.01.003.

van der Kooij, H. and de Vlugt, E. (2007) 'Postural Responses Evoked by Platform Pertubations Are Dominated by Continuous Feedback', *Journal of Neurophysiology*, 98(2), pp. 730–743. doi:10.1152/jn.00457.2006.

Krinitsyn, N.S. *et al.* (2016) 'Plant Identification in the Closed-loop Control System', *Hong Kong*, p. 5.

Ljung, L. (2003) 'Version 6 of the system identification toolbox', *IFAC Proceedings Volumes*, 36(16), pp. 957–962. doi:10.1016/S1474-6670(17)34884-X.

Ljung, L. (2006) 'Frequency Domain Versus Time Domain Methods in System Identification – Revisited', in Francis, B.A., Smith, M.C., and Willems, J.C. (eds) *Control of Uncertain Systems: Modelling, Approximation, and Design*. Berlin/Heidelberg: Springer-Verlag (Lecture Notes in Control and Information Science), pp. 277–291. doi:10.1007/11664550_15.

Ljung, L. (2010) 'Perspectives on system identification', *Annual Reviews in Control*, 34(1), pp. 1–12. doi:10.1016/j.arcontrol.2009.12.001.

Loram, I.D. *et al.* (2014) 'Does the Motor System Need Intermittent Control?', *Exercise and Sport Sciences Reviews*, 42(3), pp. 117–125. doi:10.1249/JES.0000000000000018.

MathWorks United Kingdom (2022a) *MATLAB tfest*, *MathWorks*. Available at: https://uk.mathworks.com/help/ident/ref/tfest.html (Accessed: 1 February 2022).

MathWorks United Kingdom (2022b) *tfestOptions - Option set for tfest*, *Option set for tfest*. Available at: https://uk.mathworks.com/help/ident/ref/tfestoptions.html (Accessed: 17 April 2022).

Morud, J. and Skogestad, S. (1996) 'Dynamic behaviour of integrated plants', *Journal of Process Control*, 6(2–3), pp. 145–156. doi:10.1016/0959-1524(95)00045-3.

Narasimhan, S. and Bombois, X. (2012) 'Plant friendly input design for system identification in closed loop', *IFAC Proceedings Volumes*, 45(16), pp. 1335–1340. doi:10.3182/20120711-3-BE-2027.00124.

Skogestad, S., Havre, K. and Larsson, T. (2002) 'CONTROL LIMITATIONS FOR UNSTABLE PLANTS', *IFAC Proceedings Volumes*, 35(1), pp. 485–490. doi:10.3182/20020721-6-ES-1901.00330.

Tanabe, H. *et al.* (2016) 'Effect of intermittent feedback control on robustness of human-like postural control system', *Scientific Reports*, 6(1), p. 22446. doi:10.1038/srep22446.

Appendix A: Default Values, Constants, and Software Information

Table 10: A tabulation of the default values that were implemented in simulations with each controller configuration

	Open-Loop Interval (Δ_{OL}) (s)	Threshold (rad)	Time Delay (s)
Clock-Driven	VoI	0	0
Event-Driven	0.10	VoI	0
Continuous	0	N/A	VoI

For **all** simulations, a ladder of step inputs was used for the setpoint, as shown in Figure 17. The magnitudes were as follows: +0.25 at t = 5, +0.25 at t = 10, +0.25 at t = 15, -0.25 at t = 20. All simulations were allowed to run for thirty seconds.

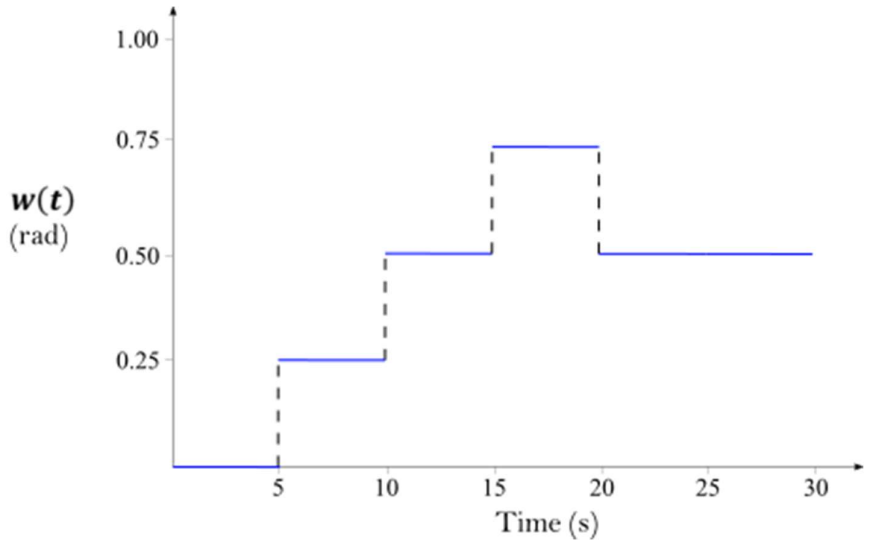

Figure 17: A plot showing the variation of the setpoint input $w(t)$ over time

Table 11: A tabulation of the constant values implemented for simulations of the Álvarez Doublein model

Constant	Value	Constant	Value
J	77 kgm^2	c	0.85 (no units)
V	2.9 $Nm(rad)^{-1}$	h	0.92 m
m	70 kg	g	9.81 ms^{-2}

Important Note:
All simulations were completed using MATLAB R2021b, as published by The MathWorks, Nantick, Massachusetts, United States. Additionally, the following MATLAB add-on packages were used: the Signal Processing Toolbox Version 8.7, the Statistics and Machine Learning Toolbox Version 12.2, the System Identification Toolbox Version 9.15, and the Control System Toolbox Version 10.11. All add-on packages are published by The MathWorks, Nantick, Massachusetts, United States.

Appendix B – Unstable System Results for Section 3.3

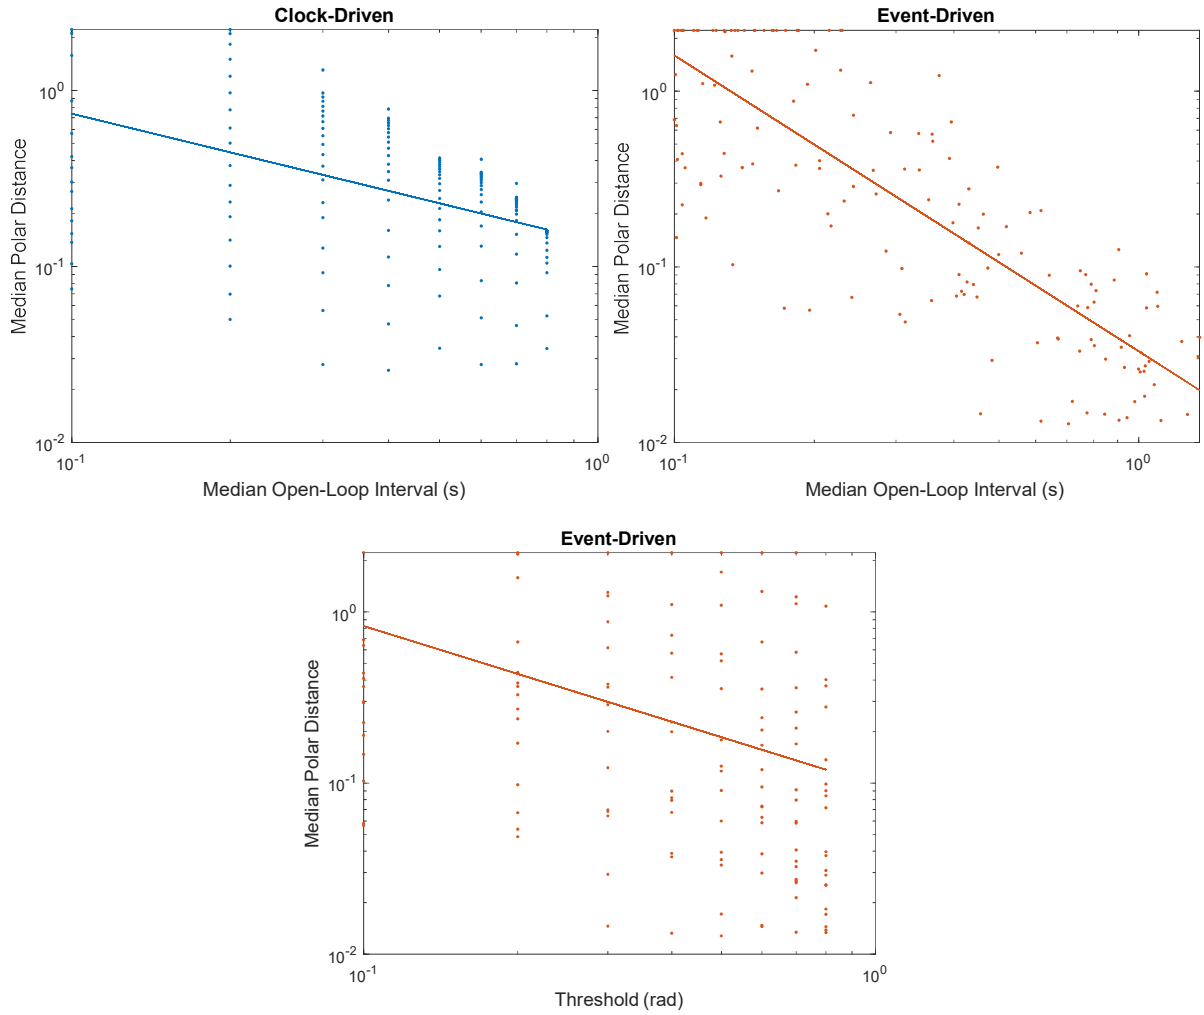

Figure 18a (Top Left), 18b (Top Right), and 18c (Bottom): Figures 18a and 18b are logarithmic plots showing the impact of median open-loop interval on median polar distance for the unstable pendulum. Note that the open-loop intervals in clock-driven control are evenly spaced since t_i for any single controller is constant. Figure 18c is a logarithmic plot displaying the impact of threshold on median polar distance.

Table 12: A tabulation of the gradients of the least-square regression lines for each of the datasets in Figures 18a, 18b, and 18c (above), with confidence intervals

	x^1 coefficient	Lower Bound x^1 CI	Upper Bound x^1 CI
Clock-Driven (t_i)	-0.728	-0.944	-0.5124
Event-Driven (t_i)	-1.684	-1.864	-1.5045
Event-Driven (Threshold)	-0.928	-1.3017	-0.5551