Probabilistic Perspective on Compensatory, Pursuit and Preview Manual Control

Mulder, Max; Pool, Daan M.; van der El, Kasper; van Paassen, René M.M.

DOI
10.1016/j.ifacol.2022.10.248

Publication date
2022

Document Version
Final published version

Published in
IFAC-PapersOnline

Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnline CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 55-29 (2022) 154–159

Probabilistic Perspective on Compensatory, Pursuit and Preview Manual Control

Max Mulder, [1] **Daan M. Pool, Kasper van der El and René (M. M). van Paassen**

Control and Simulation, Faculty of Aerospace Engineering, TU Delft, 2629 HS, Delft, The Netherlands

Abstract: Mathematical human control models are widely used in tuning manual control systems and understanding human performance. Human behavior is commonly described using linear time-invariant models, averaging-out all non-linear and time-varying effects, which are gathered into the remnant. These models are limited in their capability to capture particular tracking strategies that an experienced subject may learn to use. In this paper, we consider manual control from a different perspective, namely through investigating the probability densities of the tracking error for different regions of the target signal amplitude. Results show that distinct strategies become apparent for compensatory, pursuit and preview tracking tasks. Effects of these strategies are often averaged-out by current models and can only be captured in situation-dependent models. Modeling this systematic human adaptation not captured in linear models could potentially lead to better model fits and explain/reduce part of the remnant.

Keywords: Cybernetics, manual control, system identification, probability theory

1. INTRODUCTION

Human Controllers (HC) are biological control systems capable of exhibiting an enormous variety in behaviour (McRuer and Jex, 1967). In 1960, Krendel and McRuer (1960) introduced their Successive Organisation of Perception (SOP) hierarchy for human manual control, describing the development of skill-based manual control in three stages, compensatory, pursuit/preview, and precognitive control. Depending on the task variables, mainly the type of display (the three main types of McRuer and Jex (1967) shown in Fig. 1), the controlled element (CE) dynamics and the characteristics (power, bandwidth) of the signal-to-be-followed (the target f_t), human controllers will *systematically adapt* their dynamic response to achieve high tracking performance with limited control effort.

In a recent overview, Mulder et al. (2018) argue that cybernetics theory has predominantly focused on the lowest SOP level of human control, compensatory tracking, culminating in the widely-used and universally-accepted *crossover model* (McRuer and Jex, 1967). This model, and most of those that followed, allows to describe the steady-state feedback response of operators to unpredictable quasi-random target signals. Until Drop (2016) and Van der El (2018), universal HC models for tracking with pursuit and preview displays, which allow the HC to exert a strong feedforward (and with predictive target signals even precognitive) response barely existed. Mulder et al. (2018) concluded that much more efforts should be put into understanding human control behavior in these operationally relevant tasks, and strive for describing the learning, *adaptive* human capabilities as these are the main reason humans are still included in the loop.

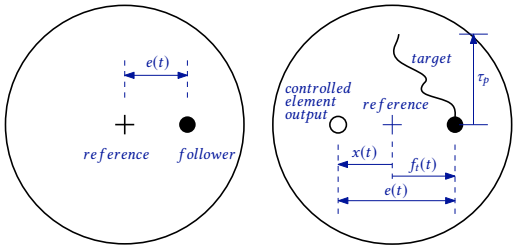

(a) Compensatory (b) Pursuit/Preview

Fig. 1. Compensatory, Pursuit and Preview displays.

Modeling adaptive behavior inevitably means moving ahead, from the current predominant linear time-invariant (LTI) models (Zaal et al., 2009) to models that allow including time-varying or even non-linear behavior. It also means that our current experimental and theoretical paradigm, which aims to fit LTI models to data from an entire experimental run, *averaging* the human response to fit LTI model and then assigning the "remaining response" to a remnant signal, needs to change. Mulder et al. (2019) argued to first study whether some strategies that follow from "common sense" heuristics actually exist in tracking data, strategies that could partially explain the observed time-varying, nonlinear behavior put into the remnant. An example strategy with the pursuit display of Fig. 1 would be to keep the CE symbol (which the HC controls) always "at the inside" of the tracking symbol, especially when the latter symbol moves closer to its extreme left/right positions (Neilson et al., 1988).

In this paper we present a probabilistic analysis of possible tracking strategies found in the experiment from Van der El et al. (2020), where the HC adaptation to target signal bandwidth was studied (and modeled) for

[1] m.mulder@tudelft.nl

10.1016/j.ifacol.2022.10.248

the compensatory, pursuit and preview displays. Section 2 discusses the main rationale of this analysis, which is based on (Mulder et al., 2019). Section 3 briefly explains the experiment of Van der El et al. (2020), and the current data analysis method. Results will be shown and discussed in Section 4; the paper ends with conclusions in Section 5.

2. BACKGROUND

2.1 Common Modeling attempts

A plethora of HC models exist and for the sake of brevity only the 'universal' model developed by Van der El (2018) will be discussed. This model has been shown to accurately describe HC behavior with *all three* tracking displays of Fig. 1 (Van der El et al., 2016, 2020).

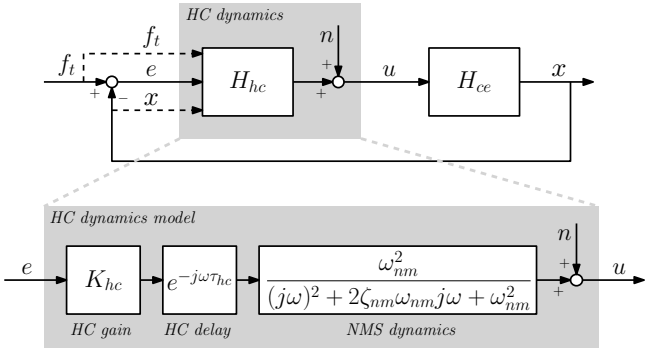

Fig. 2. Block diagram of the pursuit tracking task (top) and the HC error feedback dynamics as typically found in compensatory tracking tasks (bottom).

Fig. 2 shows the control scheme for pursuit tracking, where the HC has three inputs (Wasicko et al., 1966): the target signal f_t and CE output x, shown by the two moving symbols on the display (see Fig. 1), and the error e, i.e., the difference between both symbols $e = f_t - x$. Paraphrasing Mulder et al. (2019), because x is available the HC can 'see what she is doing', can have proper eye-hand coordination and can explore the CE dynamics. Because f_t is available, the HC can characterize its (stochastic) properties, in case of simple target signals perhaps even remember (parts of) the signal, and can try to anticipate its movements. Note that all these elements, especially the latter ones, form a problem for cyberneticists to capture with their LTI models, as will be discussed in the next subsection.

In compensatory tracking only the error e is shown, which makes the task markedly simpler to analyze and model. The HC has just one input and does not have all the benefits a pursuit display offers. With quasi-random unpredictable target signals the HC can only feed back the error, successfully captured in McRuer and Jex (1967)'s crossover model. For tracking with the preview display, the benefits of the pursuit display can be exploited even better, as then also the *future* τ_p seconds of the target signal is shown, see Fig. 1 (Van der El et al., 2016).

Van der El (2018) proposes a general model to describe HC behaviour with all three displays; the 'pursuit' version is illustrated in Fig. 3. In short, he proposes a compensatory tracking control scheme for all three displays, extended

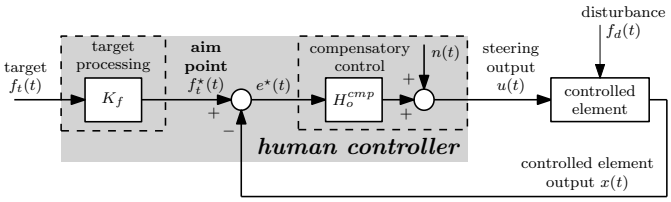

Fig. 3. Block diagram of the HC, modeled as a compensatory controller with a reference "aim point" obtained by pre-filtering the target signal f_t.

with a "target processing" block that describes the way in which the HC processes the perceived movement of the target symbol (i.e., signal f_t), to self-generate an "error" signal $e^\star$ minimized by a compensatory feedback mechanism. For the compensatory display this pre-processing block is not existent, and the self-generated error will be the same as the presented error e. With a pursuit display, the pre-processing of f_t is simply a scaling with gain K_f, see Fig. 3. With the preview display, the pre-processing block includes this same gain, a low-pass filter to filter out high-frequency movements of the target signal f_t that the HC is likely to ignore, and a negative delay to account for the preview of the target signal. This elementary model has been extensively and successfully validated, as reported in (Van der El et al., 2016, 2018b,a, 2020). Other recent models to describe HC behaviour are described in (Zhang et al., 2017; Sheffler et al., 2019).

2.2 Nonlinear Control Strategies

Mulder et al. (2019) discussed situation-dependent control strategies with the pursuit display that are a direct consequence of the way our tracking experiments are set-up. Basically, one always scales the target signal f_t such that – with the pursuit and preview displays – the target (=reference) symbol never leaves the screen. In addition, the target signal is often a sum-of-sinusoids signal, defined such that it has a quasi-random appearance and is normally distributed. Whereas the quasi-randomness prevents the HC to predict or even remember (parts of) f_t and exert precognitive control, the statistical property of being normally distributed also has an important consequence. It means that the target symbol will be mostly moving around the display center, and will only for short times move to the left or right extremes. A smart HC will quickly learn to use this property, to improve performance, but especially to reduce workload (Neilson et al., 1993).

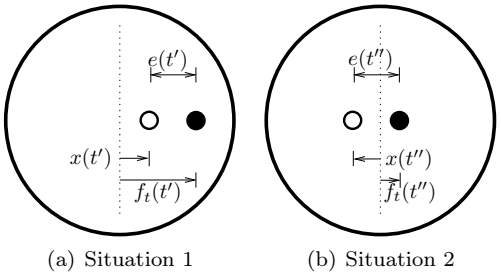

(a) Situation 1 (b) Situation 2

Fig. 4. Two situations with the same error $e = e(t') = e(t'')$, but which may lead to different HC responses.

Consider Fig. 4, which shows a pursuit display for two situations, both with the *same* error e. In Situation 1 the

target symbol has an extreme position and is close to the display border, hence, given the way f_t is constructed it is very likely that *it will move back to the screen center*. A skilled HC will learn to use this property and *not* try to force the CE symbol more to the right. In Situation 2, however, the target symbol is close to the display center, and the HC cannot anticipate well in what direction it may be moving, leaving her with no alternative to control the CE symbol to the right, closer to the target.

Quoting Mulder et al. (2019), "... the level of predictability of f_t depends on the value of f_t itself, and it is likely that a trained HC will adapt her behavior *depending on where the target symbol is located on the screen*." None of the existing LTI HC models are able to capture this behavior, as they include no mechanism to account for this dependency in a time-varying (that is, situation-dependent, here "f_t amplitude"-dependent) manner. As a result, if the strategy would indeed exist, any fitted linear model will *average-out* the phenomenon and will put all effects which cannot be captured with the linear model in the remnant signal.

This potential workload-reduction strategy with the pursuit display follows from common sense, and can indeed become an important heuristic for a trained HC. Whether this strategy can indeed be found in the data, and whether the same strategy exists for the compensatory and preview displays, is unknown, however; this forms the topic of this paper. The next section describes the experiment from which the data will be analyzed within the current context.

3. METHOD

3.1 Experiment

To validate his empirical HC model Van der El (2018) performed a number of experiments studying human manual control behavior with compensatory, pursuit and preview displays. Effects of CE dynamics (Van der El et al., 2018b), preview time (Van der El et al., 2018a) and target signal bandwidth (Van der El et al., 2020) have been studied and the changes in human control behavior captured in one single model. In this paper only the data from the target signal bandwidth experiment are studied.

In short, Van der El et al. (2020) describe an experiment with two independent variables: (i) display, and (ii) target signal bandwidth. Three displays were used: Compensatory (C), Pursuit (P) and Preview (PR). Three bandwidths ω_i were applied: 1.5, 2.5 and 4.0 rad/s, corresponding with the target signals from McRuer and Jex (1967). The controlled element dynamics were a double integrator; the preview time in the PR display was fixed at 2 seconds. A factorial design led to nine conditions.

Nine subjects were instructed to closely follow the target signal f_t: minimize error e between the target f_t and CE output x. The target was constructed using a sum of sines, yielding a quasi-random (that is: unpredictable for humans), approximately normally-distributed signal. To avoid subjects to recognize or remember the target signal, five different phase sets were used. The variance of the target signal was fixed at 1.61 cm^2 for all conditions. In addition to the target signal a disturbance signal f_d was added to the CE, of which the phases and variance (0.26 cm^2) were fixed for all conditions. In this paper, the HC response to the disturbance signal will be neglected.

Participants were extensively trained; for each condition five measurement runs were recorded. Conditions were ordered using a latin square. Each run lasted 128 seconds, of which the final 120 seconds were used. Data were sampled at 100 Hz; each run led to 12,000 data points. With the nine participants, nine conditions and five runs per condition the data of a total of 405 runs are available; for more details, see (Van der El et al., 2020).

3.2 Data analysis

In this paper we focus on just two signals from the experiment of Van der El et al. (2020): the measured error signal $e(t)$ and the target signal $f_t(t)$. Based on the discussion of possible nonlinear control strategies in Section 2 the rationale of our analysis was simple. Consider Fig. 5(a) which shows the left/right movement of the P/PR display target symbol on the screen in cm (horizontal axis) as a function of time (vertical axis), for $\omega_i = 2.5$ rad/s.

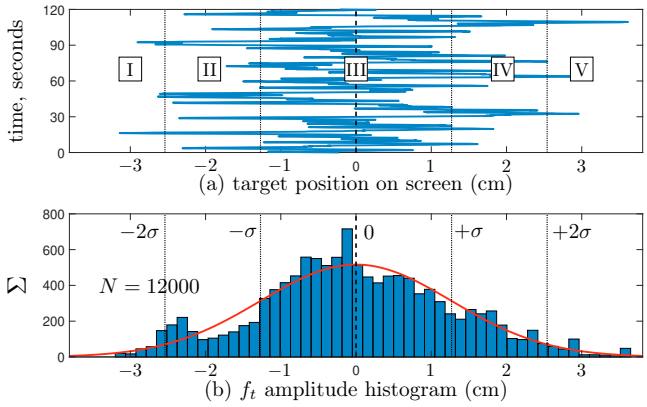

Fig. 5. Position of the reference symbol as a function of time (f_t, top); density (bottom, with the Normal density estimate in red) (bandwidth 2.5 rad/s)

Mulder et al. (2019) argued that human controllers, in due time, will recognize the approximately normal distribution of the target symbol position, Fig. 5(b). In our analysis we will study the probability density of the error signal $e(t)$ for five different regions of the target signal amplitude: Regions I, II, III, IV and V in Fig. 5(a). Region III is centered around the (zero) average of f_t and corresponds to the $\pm 1\sigma$ area of the density of f_t. Regions II and IV correspond to the amplitudes of f_t that are, respectively, in-between -2σ and -1σ and $+1\sigma$ and $+2\sigma$. Regions I and V belong to the amplitudes of f_t that are, respectively, smaller than -2σ and larger than $+2\sigma$. For a normal distribution, Region III has about 68.26% of all f_t data, Regions II and IV together 27.18% and Regions I and V together only 4.56% of the data.

Now, for *each* sample $f_t(t')$ belonging to one of the five regions we will take the sample of the error signal $e(t')$ at that moment t'. Repeating this procedure for all conditions, averaging out the effects of participants, we obtain an estimate of the probability density functions of the error signal in each of the five regions: $\hat{f}_{\bar{e}_i}(e)$ (with i = I, II, III, IV, V). When our assumption about HC

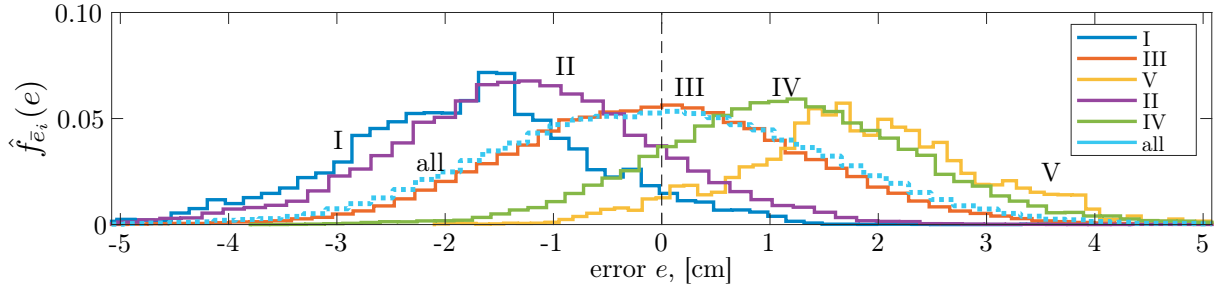

Fig. 6. Densities of the error signal e when the target signal f_t amplitude belongs to Regions I, II, III, IV, V and "all".

heuristics, discussed in Section 2, is correct, these densities will differ and depend on the type of display (C, P or PR), likely also on the target f_t bandwidth.

Note that the quality of these probability density function estimates will differ between regions. For Region III we can use approximately 68.26% of all data; for Regions I and V only a meagre 2.28% of our data. With 405 available runs of 12,000 data points each, 1% of the data corresponds to 48,600 samples; for an analysis of each of the 9 experiment conditions separately this number reduces to 5,400. So whereas the estimate for $\hat{f}_{\bar{e}_{\mathrm{III}}}(e)$ can be based on about 386,604 samples for each condition, the estimate for $\hat{f}_{\bar{e}_{\mathrm{I}}}(e)$ is based on only about 12,312 samples. This is an inevitable consequence of the approach taken.

3.3 Hypotheses

Our main hypothesis is that with the pursuit display the f_t-dependency heuristic discussed in Section 2 will be the most prevalent, especially when the task gets more difficult (higher target signal bandwidth). Evidence for this strategy will be when the data for the different regions will be distributed in a different way, with non-zero means for Regions II and IV and especially Regions I and V.

We further expect that this strategy will be less strong with the preview display, as here the need for the HC to predict where the target symbol will move vanishes, and she can aim to maximize performance even in cases where the target symbol moves to the extreme left/right of the display. With the compensatory display the heuristic will not work, the HC only sees "error", hence we expect the data to be zero-mean normally-distributed.

4. RESULTS AND DISCUSSION

4.1 Example densities: Pursuit, 4.0 rad/s

Fig. 6 shows the six probability density function estimates of the error signal corresponding to the target signal being in one of the five regions of interest (I-V), including the 'all regions' density function. Data are shown for the pursuit display (bandwidth 4.0 rad/s), to bediscussed first.

The figure clearly illustrates that the density is different for each Region. For Region III, where the target symbol is positioned within $\pm 1\sigma$ around the center of the pursuit display, HCs cannot easily predict "where the symbol will go", so they steer the CE close to the target symbol at all times, yielding an almost perfect normal density of the error signal around zero. In other words, the CE symbol will be oscillating around the (moving) target symbol when the target symbol is located within $\pm 1\sigma$.

This dramatically changes when the target symbol moves farther away from the center display positions, i.e., where $f_t(t)$ belongs to Regions II and IV, and especially when it moves to its extreme positions, corresponding to Regions I and V. When considering the estimated error density functions in the latter two extreme regions, we clearly see that the average error will be non-zero. The negative average error for Region I (where $f_t(t)$ is negative) and positive average error for Region V ($f_t(t)$ is positive) are convincing evidence for our main hypothesis, namely that with pursuit displays operators will learn that it is beneficial to "keep the CE symbol at the inner side of the target symbol". Already in Regions II and IV we can see that this strategy is activated. This means that – on hindsight – also Region III could have been split-up in a left/right region to verify our assumption that the heuristic is only activated for larger f_t amplitudes. With the high-bandwidth, rapidly-moving target signal (here $\omega_i = 4.0$ rad/s) this strategy is likely to be strong. To investigate effects of display and bandwidth, we need to study all other conditions as well, which is done in the next subsection.

Fig. 6 shows that the "all regions" density is similar to the 'Region III' density, an almost perfect zero-mean normal density. This follows from the fact that the majority of the samples belong to Region III, the target symbol moving close to the display center; any model fit will therefore aim to approximate these data the best. In addition, the skews for Regions I and V (and, to a lesser extent Regions II and IV) are symmetric, adding these data cancels the skews. The skew in the estimated probability density functions for data for Regions I and V could be the result from having only a small number of samples (2.28 %).

4.2 All Density Estimates

Fig. 7 shows the error density function estimates for Regions I (extreme left), III (center) and V (extreme right), for the three displays (C, P, PR, the columns) and three three f_t bandwidths (1.5, 2.5 and 4.0 rad/s, the rows). The density functions for the intermediate Regions II and IV are omitted for reasons of clarity. The results for the pursuit, 4 rad/s bandwidth case have already been discussed above. From Fig. 7 one can directly see the increasing tracking performance (low error) when moving from the compensatory (left) to preview (right) displays, with more and more data closer to the zero average and higher values of the density functions around zero. This effect becomes stronger when the target bandwidth increases (from top to

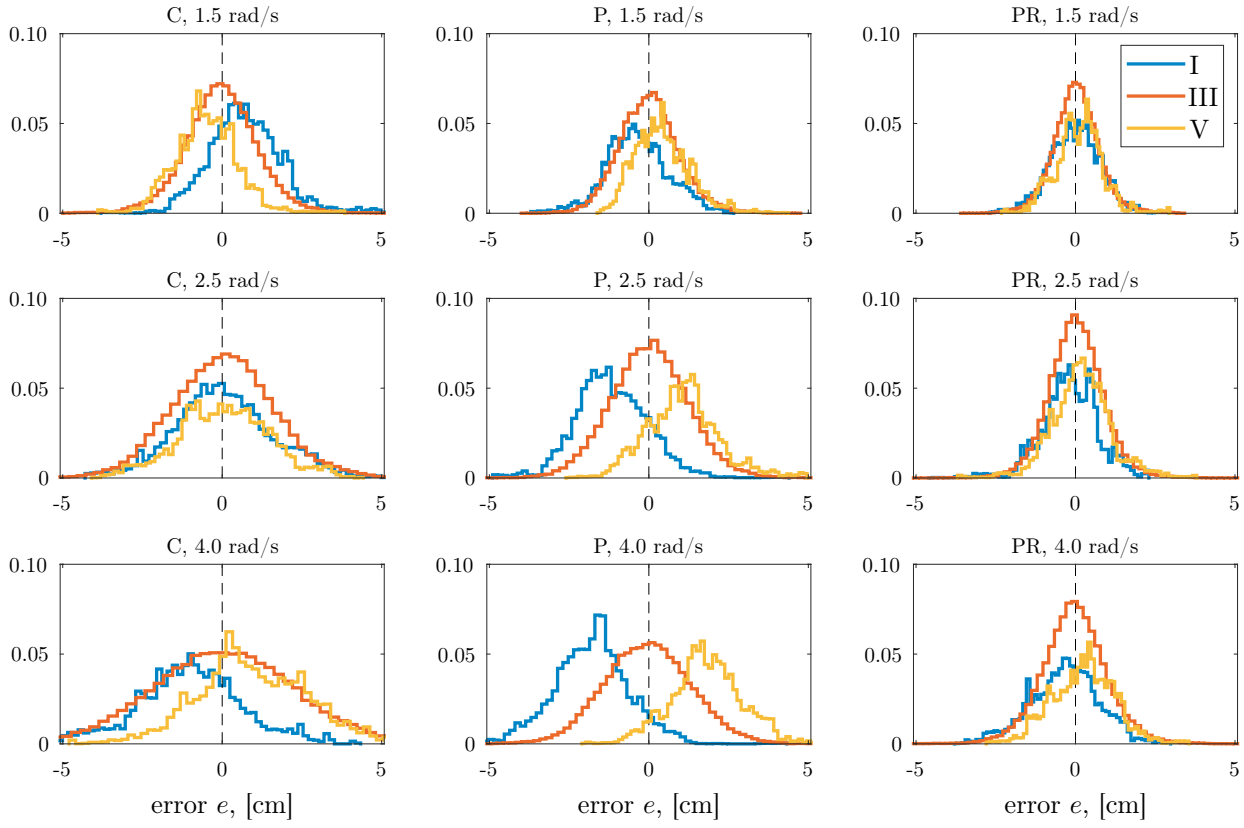

Fig. 7. Error signal densities when the target f_t amplitude belongs to Regions I, III and IV.

bottom). These results are (obviously) in line with those presented by Van der El et al. (2020). In the following we will discuss the main results for the pursuit, preview and compensatory displays, respectively.

Pursuit Display Data Here the error density functions for Regions I and V are markedly shifted to the left and right, respectively, an effect that becomes stronger when the target signal bandwidth increases. Clearly, with a faster-moving target signal the HC will keep the CE symbol away from the screen edges; it makes no sense to try and follow the rapid motions of the target symbol when it is close to an extreme, also because it is very likely that the symbol will be moving back to the center anyway.

Preview Display Data The densities are markedly different from those of the pursuit display. Although the densities for Regions I and V shift in the same direction (left and right, respectively) as for the pursuit display, the size of the shift is much smaller, and the average of all densities remains close to zero. Showing the *future* of the target symbol (here, 2 seconds ahead) releases the HC from the effort of predicting where the target symbol will move, and also allows her to much more closely follow that symbol, even when it moves to the extremes of the display. The certainty of knowing where the symbol will be in the immediate future, together with the instruction and motivation to track as accurately as possible, persuade the HC to also try to achieve zero error in extreme cases, leading to the superior performance with the preview display.

Compensatory Display Data The densities shift in the same direction for the high-bandwidth condition as for the pursuit display, shift in the *opposite* direction for the low-bandwidth condition, and remain approximately zero-centered for the intermediate bandwidth condition. Apparently, the HC matches the CE output and target signal amplitudes in the intermediate condition, *undershoots* the target in the high-bandwidth condition and *overshoots* the target in the low-bandwidth condition. Keep in mind that with the compensatory display the HC *cannot see* either the target or CE symbols (respectively, f_t and x) but *only* acts on the presented difference between the two, the error e. This means that the 'over' and 'under'-shooting of the target signal happens *without* the operator knowing it. An explanation of this finding could be the 'crossover regression' effect, the reduction of the HC gain when the target signal bandwidth increases (McRuer and Jex, 1967).

4.3 Discussion

Fig. 8 shows the tracking error e variance and remnant contribution to error e_n (both normalized with the power of the f_t and f_d forcing functions), the estimated gains K_f, and the Variance Accounted For (VAF) values of the HC model predictions, for all conditions of the experiment by Van der El et al. (2020). The contribution of the remnant to the error variance is lowest with the preview display, but almost the same for the pursuit display. The estimated K_f gains are close to 1 for the preview display, and decrease from 0.75 to 0.55 for the pursuit display for higher target bandwidths; results which correspond with our analysis. With the preview display, the HCs maintain a strategy to fully correct for any visible error, independent of where the target symbol moves. Apparently, lowering the K_f gain for the pursuit display sufficiently captures the

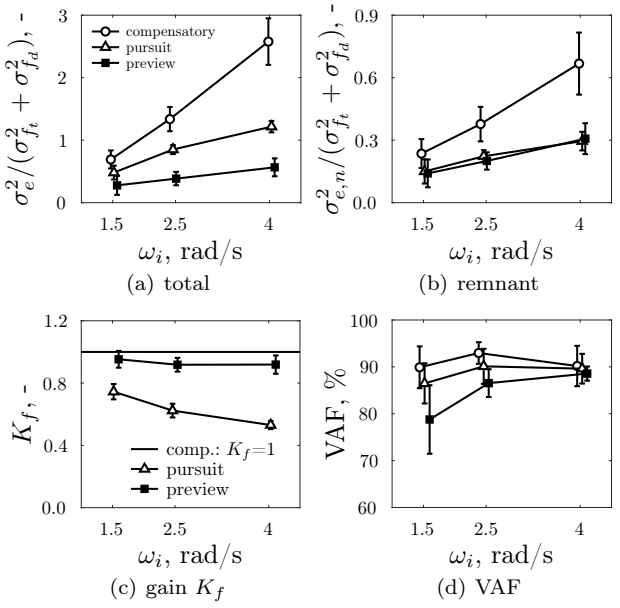

Fig. 8. Variance of tracking error (a); remnant contributions (b); estimated values of the gain K_f (c), and VAF (d); (average over nine participants, 95% confidence intervals).

HC strategy of maintaining the CE symbol closer to the display center than the target symbol, also shown by the VAF values which are high. In other words, the lower K_f gain allows the LTI model fit to account on average quite well for the (from the present analysis) evidently time-varying HC behavior. Recall that because f_t is normally-distributed, the majority of the f_t data (and with that, the HC behavior) remain close to its average.

Future work, with a possible time-varying K_f (e.g., modelling the HC with a gain K_f which depends on the f_t magnitude), should be conducted to study whether any improvements in HC modeling for the pursuit display are worth the effort of increasing model complexity. Evidence could be found in lowering the model remnant (as the HC behavioral variation can be better accounted for), or a better reproduction of the operator control signal u (using VAF-like metrics). In the time-domain, one could consider the measured and modeled HC control signal u and check whether indeed it will be less for the phases where the f_t amplitude is large (Regions I, V).

5. CONCLUSIONS

Experimental tracking data were analyzed to investigate possible heuristics human controllers use to maximize their performance-to-effort ratio. Results show that these strategies indeed exist, especially for the pursuit display, an effect that becomes stronger for higher target signal bandwidths. The heuristic implies that the HC control strategy depends on the amplitude of the target signal. Current models account for this effect through a target signal pre-processing gain smaller than one, yet average-out this effect over the entire run. Time-varying models, where the gain changes as a function of time (or better: scale with the target amplitude) could better describe this strategy, increasing accuracy and reducing the (unexplained) 'remnant'-part of our HC models.

REFERENCES

Drop, F.M. (2016). *Control-Theoretic Models of Feedforward in Manual Control*. Ph.D. dissertation, TU Delft.

Krendel, E.S. and McRuer, D.T. (1960). A Servomechanics Approach to Skill Development. *Journal of the Franklin Institute*, 269(1), 24–42.

McRuer, D.T. and Jex, H.R. (1967). A Review of Quasi-Linear Pilot Models. *IEEE Trans. on Human Factors in Electronics*, 8(3), 231–249.

Mulder, M., Pool, D.M., Abbink, D.A., Boer, E.R., Zaal, P.M.T., Drop, F.M., Van der El, K., and Van Paassen, M.M. (2018). Manual Control Cybernetics: State-of-the-Art and Current Trends. *IEEE Trans. on Human-Machine Systems*, 48(5), 468–485.

Mulder, M., Pool, D.M., Van der El, K., Drop, F.M., and Van Paassen, M.M. (2019). Manual Control with Pursuit Displays: New Insights, New Models, New Issues. In *14th IFAC Human-Machine Systems symposium, Tallinn, Estonia, Sept. 16-19*, 139–144. IFAC.

Neilson, P.D., Neilson, M.D., and O'Dwyer, N.J. (1993). What Limits High Speed Tracking Performance? *Human Movement Science*, 12, 85–109.

Neilson, P.D., O'Dwyer, N.J., and Neilson, M.D. (1988). Stochastic Prediction in Pursuit Tracking: An Experimental Test of Adaptive Model Theory. *Biological Cybernetics*, 58, 113–122.

Sheffler, A.J.S., Mousavi, S.A.S., Hellström, E., Jankovic, M., Santillo, M.A., Seigler, T.M., and Hoagg, J.B. (2019). Effects of Reference-Command Preview as Humans Learn to Control Dynamic Systems. In *IEEE SMC Conference, Bari, Italy, Oct. 6-8*, 4200–4205.

Van der El, K. (2018). *How Humans use Preview Information in Manual Control*. Ph.D. dissertation, TU Delft.

Van der El, K., Padmos, S., Pool, D.M., Van Paassen, M.M., and Mulder, M. (2018a). Effects of Preview Time in Manual Control. *IEEE T-HMS*, 48(5), 486–485.

Van der El, K., Pool, D.M., Damveld, H.J., Van Paassen, M.M., and Mulder, M. (2016). An Empirical Human Controller Model for Preview Tracking Tasks. *IEEE Trans. on Cybernetics*, 46(11), 2609–2621.

Van der El, K., Pool, D.M., Van Paassen, M.M., and Mulder, M. (2020). Effects of Target Trajectory Bandwidth on Manual Control Behavior in Pursuit and Preview Tracking. *IEEE Trans. on HMS*, 50(1), 68–78.

Van der El, K., Pool, D.M., Van Paassen, M.M., and Mulder, M. (2018b). Effects of Preview on Human Control Behavior in Tracking Tasks with Various Controlled Elements. *IEEE Trans. on Cyb.*, 48(4), 1242–1252.

Wasicko, R.J., McRuer, D.T., and Magdaleno, R.E. (1966). Human Pilot Dynamic Response in Single-loop Systems with Compensatory and Pursuit Displays. Technical Report AFFDL-TR-66-137.

Zaal, P.M.T., Pool, D.M., Mulder, M., and Van Paassen, M.M. (2009). Multimodal Pilot Control Behaviour in Combined Target-Following Disturbance-Rejection Tasks. *J. of Guidance, Cont., Dyn.*, 32(5), 1418–1428.

Zhang, X., Wang, S., Hoagg, J.B., and Seigler, T.M. (2017). The Roles of Feedback and Feedforward as Humans Learn to Control Unknown Dynamic Systems. *IEEE Trans. on Cybernetics*, 48(2), 543–555.