Task-dependent switching of feedback controllers

Justinas Česonis[1] **and David W. Franklin**[1,2,3,*]

***For correspondence:**
david.franklin@tum.de

[1]Neuromuscular Diagnostics, Department of Sport and Health Sciences, Technical University of Munich, Germany; [2]Munich Institute of Robotics and Machine Intelligence (MIRMI), Technical University of Munich, Germany; [3]Munich Data Science Institute (MDSI), Technical University of Munich, Munich, Germany

Abstract The separation of distinct motor memories by contextual cues is a well known and well studied phenomenon of feedforward human motor control. However, there is no clear evidence of such context-induced separation in feedback control. Here we test both experimentally and computationally if context-dependent switching of feedback controllers is possible in the human motor system. Specifically, we probe visuomotor feedback responses of our human participants in two different tasks – stop and hit – and under two different schedules. The first, blocked schedule, is used to measure the behaviour of stop and hit controllers in isolation, showing that it can only be described by two independent controllers with two different sets of control gains. The second, mixed schedule, is then used to compare how such behaviour evolves when participants regularly switch from one task to the other. Our results support our hypothesis that there is contextual switching of feedback controllers, further extending the accumulating evidence of shared features between feedforward and feedback control.

Introduction

Whether it is touching a hot surface, returning a tennis serve or simply lifting an object, the human body utilises a variety of sensory inputs to produce movements of any complexity. Indeed, different feedback modalities of human motor control, such as stretch reflex (***Houk*** (***1976***); ***Pruszynski and Scott*** (***2012***); ***Dimitriou*** (***2016***)), vestibulo-ocular reflex (***Barr et al.*** (***1976***); ***Tabak and Collewijn*** (***1994***)), visuomotor (***Prablanc and Martin*** (***1992***); ***Franklin and Wolpert*** (***2008***); ***Izawa and Shadmehr*** (***2008***); ***Knill et al.*** (***2011***); ***Reynolds and Day*** (***2012***); ***Franklin et al.*** (***2012***); ***Pruszynski et al.*** (***2018***); ***Zhang et al.*** (***2018***); ***Saijo et al.*** (***2005***)), or even auditory feedback (***Baram and Miller*** (***2007***); ***Oscari et al.*** (***2012***)) have extensively been studied in prior literature. However, most studies have investigated feedback control in paradigms of either a single task (***Saunders and Knill*** (***2003***, 2005); ***Franklin et al.*** (***2017***); ***Oostwoud Wijdenes et al.*** (***2011***, 2013)), or multiple tasks presented in their own dedicated blocks (***Day and Lyon*** (***2000***); ***Diedrichsen*** (***2007***); ***Maeda et al.*** (***2018***); ***Cross et al.*** (***2019***); ***Česonis and Franklin*** (***2020***)). While such designs provide key insights into the behaviour of the feedback controller in isolation, they are not entirely reflective of human behaviour in real-life situations. For example, a realistic sequence of events could require a volleyball player to first pick up the ball from the ground by reaching for it with their hand and stopping on contact, only then to hit the same ball with the same hand a few moments later while serving. While studying both components independently has received focus in the field of motor control, any interactions between the feedback controllers in the context of rapid switching have not been broadly studied.

While feedback control in human movement is critical in correcting for random errors within movements, feedforward control corrects for movement errors that are predictable. In order to systematically predict and compensate for specific errors upcoming in a given movement, the mechanism of contextual switching via contextual cues is broadly accepted. It is now well understood that performing two opposing tasks in an alternating manner will lead to interference (***Shadmehr et al.*** (***1995***); ***Bock et al.*** (***2001***); ***Sing and Smith*** (***2010***)), resulting in behaviour that is averaged between the two tasks, failing to deal with either task. However, if the two tasks are performed in sufficiently different contexts, such as separate physical or visual workspaces (***Howard et al.*** (***2013***); ***Forano and Franklin*** (***2020***); ***Hirashima and Nozaki*** (***2012***)), or different lead-in (***Howard et al.*** (***2012***); ***Sarwary et al.*** (***2015***)) or follow-through movements (***Howard et al.*** (***2015***); ***Sheahan et al.*** (***2016***)), this interference can be reduced, allowing the formation of two separate motor memories. While it is reasonable to expect similar contextual regulation of the feedback controllers, to our knowledge the regulation of feedback control in the paradigm of "multitasking" has not yet been tested. Therefore, in this study we test whether the feedback control policies exhibit such modulation when humans are presented with different tasks in an alternating manner.

One difference between studying contextual switching in feedforward and feedback control is that it is difficult to evaluate whether the feedback control policy has changed after the intervention. Specifically, it has been shown computationally that the optimal feedback controller (OFC) with fixed parameters can produce variable responses when correcting for perturbations within the movement, for example, when the comparable perturbations are induced in different parts (e.g. early or late) of otherwise identical movements (***Liu and Todorov*** (***2007***); ***Česonis and Franklin*** (***2020***, 2021); ***Poscente et al.*** (***2021***)). Furthermore, such behaviour was also observed in experimental studies (***Franklin and Wolpert*** (***2008***); ***Oostwoud Wijdenes et al.*** (***2011***); ***Dimitriou et al.*** (***2013***); ***Franklin et al.*** (***2016***); ***Česonis and Franklin*** (***2020***); ***Poscente et al.*** (***2021***)). Hence, merely observing a difference in the feedback response is not enough to conclude a change in the control policy. However, recently we demonstrated that as long as two perturbations of the same magnitude are induced at the same time-to-target, the same feedback control policy produces the same magnitude response, independent of whether the two perturbations occurred at the same location, time from the beginning of the movement, or the movement velocity (***Česonis and Franklin*** (***2020***)). Thus, we can utilise this relationship between the magnitude (or intensity) of the feedback response to a perturbation at the same time-to-target to quantify whether the difference in the response is due to the change in the control policy or not.

In this study we test whether human participants exhibit similar switching between feedback control policies as they do for feedforward control. Specifically, we test how the feedback control policies are affected when our participants are presented with a "multitasking" scenario where they have to switch between performing two distinct tasks, i.e. reaching to and stopping at the target, or hitting through the target and stopping behind it. While the two tasks are fundamentally different, and in isolation should require different feedback control policies, here we test whether the same relationship holds true in the mixed schedule (as it would for contextual switching in feedforward control), or if the interference between two control policies results in a single policy, averaged or weighted between the two independent controllers.

Results

In this study we tested the behaviour of the human feedback controller when switching between two different tasks. Specifically, we presented our human participants with two tasks requiring different control policies – a stop task, where participants had to reach and stop at the target, and a hitting task, where participants had to punch through the target and stop behind it. In our previous work we demonstrated computationally that these two different types of movements trigger feedback responses of different magnitudes, even if the perturbations occur at the same position, time, or time-to-target (***Česonis and Franklin*** (***2020***)). However, if the two movements share the same goal (for example the goal of stopping at the same target), then these feedback responses match in magnitude if the time-to-target matches in both movements, irrespective of other movement parameters like peak velocity, movement distance, distance to the target or

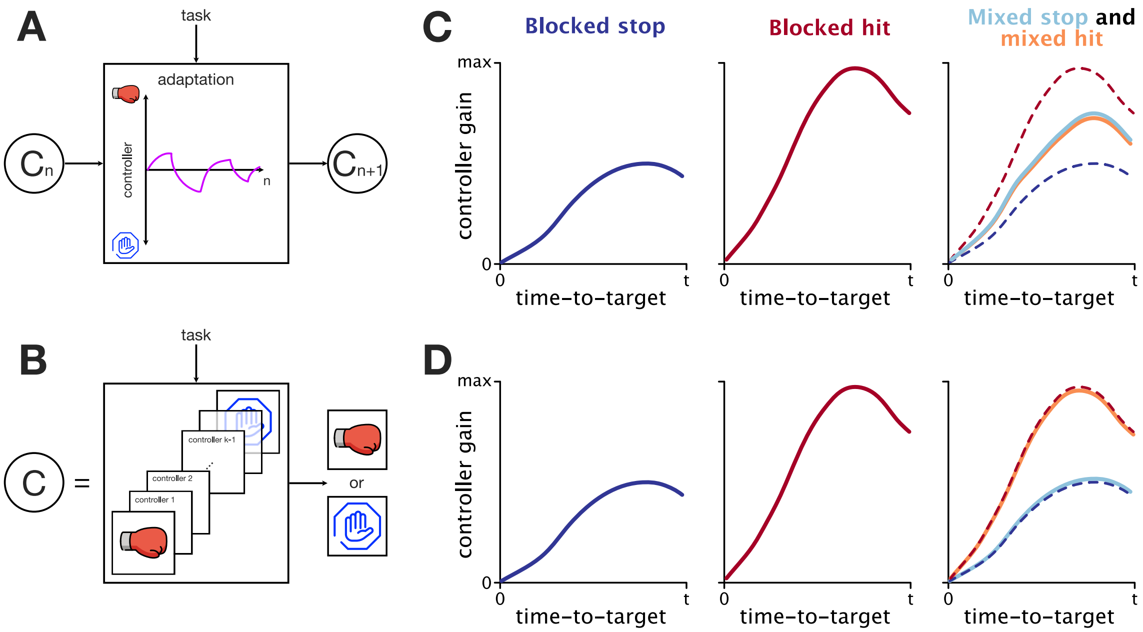

Figure 1. Theoretical predictions of two different architectures for feedback regulation. **A.** Universal feedback controller. A single feedback controller is used to produce both stop and hit movements, and is adapted to the given task over multiple trials. Such adaptive behaviour is reminiscent of the behaviour of the feedforward controller when learning two opposing force-fields without separable context. C_n indicates a feedback controller at trial n **B.** A feedback controller as a set of task-specific controllers. A task-specific controller (stop or hit) is selected based on the task-related context and is used during the given movement. Such contextual switching behaviour is reminiscent of the behaviour of the feedforward controller when learning two opposing force fields with separable context. **C.** Expected regulation of feedback responses by the universal feedback controller. When exposed to a single task for a long time (blocked schedule) the controller adapts to the given task, producing optimal responses for both stop and hit conditions. However, due to interference within the mixed schedule, such a controller would settle to the average (or weighted average) gains between the two blocked conditions. **D.** Expected regulation of feedback responses by a set of task specific controllers. Within the blocked schedule, similar regulation is expected between hit and stop as in the case of the universal controller (**C**). However, in the mixed schedule, due to the task-related context, an appropriate controller is recalled on a trial-to-trial basis, producing similar regulation as within the blocked schedule.

---

current velocity. Therefore, such a relation between time-to-target and feedback response intensity could be used to characterise the feedback control policy.

We use the relationship between the time-to-target and the feedback response intensity (which serves as a proxy for feedback controller gain) as a means to analyse the controller behaviour when the task changes. Specifically, we propose two alternatives for the architecture of such control: a single universal feedback controller that exhibits adaptation to a given task (Figure 1A), or multiple task-specific controllers, gated by task context (Figure 1B). When presented with a single task in a blocked schedule (e.g. blocked stop or blocked hit), both the universal controller and task specific controllers are expected to behave similarly, as the universal controller should easily adapt its gains appropriately for the required task. However, if multiple tasks are presented in a mixed schedule (i.e. task can randomly switch from trial to trial), the different control architectures predict different responses. Particularly, a single universal controller would aim to adapt to each presented task, thus on average producing responses somewhere in between the two given tasks within the mixed schedule (Figure 1C). In contrast, a set of task-specific controllers would produce similar responses in the mixed schedule as they would in a blocked schedule, as for every trial an appropriate controller is selected from a set of controllers, rather than being adapted for the task (Figure 1D).

In order to probe the control policies of human participants within these different tasks, we

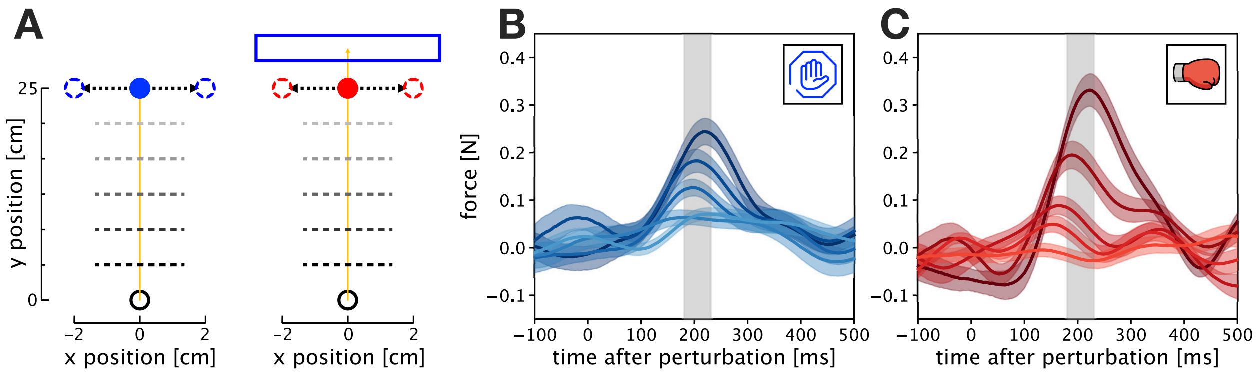

Figure 2. Experimental perturbations and responses. **A.** Perturbations in stop (left) and hit (right) conditions. Participants performed a forward reaching movement towards a target, positioned 25.0 cm in front of the start position. When the hand crossed one of five evenly spaced locations (dashed lines), a perturbation could be induced by shifting the target by 2 cm laterally for 250 ms and then returned back to the original position. Participants were instructed to either stop at the blue target (stop condition), or hit the red target and stop within the blue rectangle (hit condition). **B.** Net feedback responses to the target perturbations in the stop condition, measured via the force channel. Participants produced corrective responses to the target perturbations that varied by different perturbation onsets. Different traces represent different perturbation onsets, with darker colours indicating earlier perturbations. Shaded areas represent one standard error of the mean (SEM). The grey rectangle represents the time window of 180 – 230 ms, where the visuomotor feedback intensities are measured. **C.** Net feedback responses to the target perturbations in the hit condition.

---

occasionally perturbed participants during the movement by visually shifting the target perpendicular to movement direction and inducing a reactive visuomotor feedback response (Figure 2A). Recently it has become common practice to maintain these perturbations until the end of the movement, such that an active correction is required to successfully complete the trial (***Reichenbach et al.*** (***2013***, 2014); ***Reichenbach and Diedrichsen*** (***2015***); ***Franklin et al.*** (***2016***); ***Gallivan et al.*** (***2016***); ***Franklin et al.*** (***2017***); ***de Brouwer et al.*** (***2017***, 2018); ***Česonis and Franklin*** (***2020***)). However, we have noticed in our previous work that such maintained perturbations significantly impact the overall time-to-target, which in turn affects the visuomotor feedback gains (***Česonis and Franklin*** (***2020***)). Thus, to keep the measurements of visuomotor feedback responses consistent within time-to-target, in this study we only perturbed our participants laterally in channel trials (***Franklin and Wolpert*** (***2008***); ***Dimitriou et al.*** (***2013***); ***Scheidt et al.*** (***2000***)) and maintained these perturbations for 250 ms before switching them off, making any corrections redundant. As a result, even when producing the feedback response, participants' hands are constrained along the path of forward movement, resulting in matching movement durations independent of different perturbation onsets.

Participants produced involuntary feedback responses to the target jumps. These responses, observed as a lateral force exerted by the participants on the handle of the robotic manipulandum, were modulated by the different perturbation onsets (Figure 2BC). From these force responses we computed feedback intensities, by averaging individual responses over a time window 180 ms - 230 ms relative to the perturbation onset on each individual trial. This time window has now been used in numerous studies to quantify such responses and is associated with the involuntary, early visuomotor responses (***Cross et al.*** (***2019***); ***Franklin and Wolpert*** (***2008***); ***Dimitriou et al.*** (***2013***); ***de Brouwer et al.*** (***2017***, 2018)).

OFC model predicts differences between hit and stop conditions

We utilised the mixed-horizon OFC (***Česonis and Franklin*** (***2021***)) model, presented in our earlier work, to generate predictions of feedback control policies in our current study. Due to the experimental design of this study not requiring an extension in movement times after perturbations, the predictions of the mixed-horizon model also matched the predictions of our earlier time-to-target

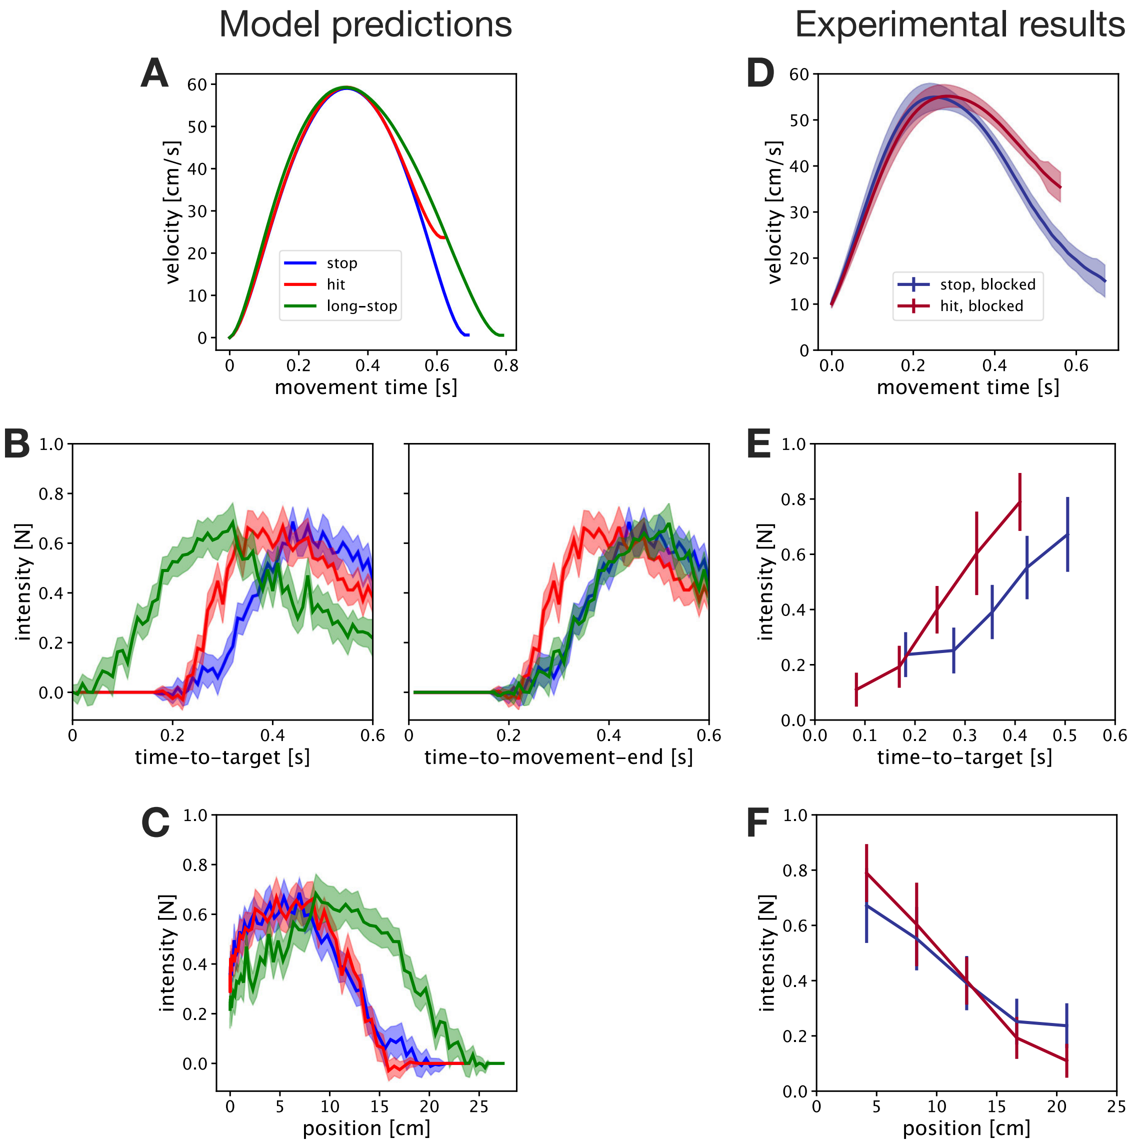

Figure 3. Comparison of OFC model predictions and experimental results. **A.** Simulated kinematics of stop, hit and long-stop conditions. Stop and hit conditions produce matching kinematics that only deviate shortly before movement end. The long-stop condition is a control simulation, that matched the kinematics of the hit condition for the duration of the hit movement, but was achieved with the same stop controller. **B.** Simulated feedback intensities as a function of time-to-target (left) and time-to-movement-end (right). Simulations predict a faster increase of response intensities for hit condition than for stop condition. As the long-stop condition is simulated via a longer (28 cm) movement, the time-to-target represents a time until the simulated movement crosses a point of 25 cm distance (the target distance). For hit and stop conditions, time-to-target and time-to-movement-end are identical. When expressed against time-to-movement-end, long-stop produces matching responses to the stop condition, as the feedback controller used for these movements is identical. With respect to the time-to-target, long-stop responses are time-shifted from the stop responses. **C.** Simulated feedback intensities as a function of the position. Stop and hit simulations with these particular kinematics produce matching feedback intensity profiles when expressed against position, even if the feedback controllers are different. In contrast, the long-stop simulation with a feedback controller matching that of the stop condition still produces different intensity profile against position. Shaded areas in simulated traces represent 95% confidence intervals for simulated results. **D.** Velocity profiles of participants in blocked stop and blocked hit conditions. The profiles match the task requirements. **E.** Feedback intensity profiles of participants in blocked stop and blocked hit conditions, expressed against time-to-target. Participants produce stronger responses at matching time-to-target in the hit condition, consistent with simulation results for hit and stop. **F.** Feedback intensity profiles of participants in blocked stop and blocked hit conditions, expressed against position. Participants produce matching responses within hit and stop conditions, supporting model simulations for stop and hit conditions, and not stop and long-stop. Error bars in experimental results represent 95% confidence intervals.

OFC model (***Česonis and Franklin*** (***2020***)). In order to compare differences in control throughout hit and stop movements, we first simulated two movement conditions: a 25 cm long movement with 60 cm/s peak velocity and velocity at the target distance <1 cm/s (stop condition), and a similar movement, but with velocity at the target >20 cm/s (hit condition) (Figure 3A). Both models were implemented using a linear quadratic regulator (LQR), and were identical, apart from the difference in state-dependent costs of terminal velocity and terminal force. Here we reduced these cost parameters for the hit model by a factor of 50 in order to reduce the incentive to stop at the target, and thus successfully simulate hit-like movements. In addition, we also simulated a third condition, that we term the long-stop condition, where we used the same position, velocity, force and mean activation costs as in the stop model, but applied for reaching movements of 28 cm. The concept of the long-stop model is to compare the actual hit behaviour, executed through a different controller, with "cheating" behaviour where the same stop movement is performed to an imaginary target, located beyond the actual target, resulting in non-zero velocity at the actual target, and thus appearing as a hit movement. For all three conditions we then induced virtual target perturbations by shifting a target laterally by 2 cm at every time step from movement onset to movement end. With these simulations we obtained one continuous feedback response profile per condition, showing a dependency of feedback response intensity on time-to-target (Figure 3B). This feedback response profile is characteristic of the particular movement control policy associated with the movement goal, as it is maintained even if the kinematics of the movement change (Figure 8 in ***Česonis and Franklin*** (***2020***)).

Even with similar simulated kinematics, that deviate from each other only in the last portion of the movement, the OFC model predicts striking differences in the control policies for stop and hit conditions (Figure 3B, blue and red traces) or hit and long-stop conditions (Figure 3B, red and green traces) when expressed against time-to-target. On the other hand, when expressed against position, even different controllers (hit vs stop) show no differences in feedback intensities, while identical control (stop vs long-stop) exhibit clear differences (Figure 3C). Among other things, these results point out limitations of position as a dependent variable in determining the changes of control policies, and provide yet additional support for time-to-target.

Our models make a few predictions for the behaviour of human participants. First, it challenges the classic assumption that visuomotor feedback response profiles are always bell shaped, if probed at evenly spaced locations or movement times. Instead, we propose that the bell-shaped feedback response profiles are consequential to the specific kinematic values imposed by the experiments, and other, for example monotonically decreasing intensity profiles, are also possible with faster movements (Figure 3C). Second, our simulations also make predictions on relative differences between the feedback intensity profiles in stop and hit conditions. Particularly, we expect the hit condition to produce stronger responses than the stop condition for short times-to-target, with this relationship inverting for long times-to-target if the two types of movements require different feedback controllers (Figure 3B). Note, that while in previous studies it is typical to compare such response profiles in terms of perturbation onset location, here no difference between hit and stop is predicted in position-dependent profiles (Figure 3C).

Human control policies match model predictions in hit and stop conditions

In order to compare the behaviour of our participants to the model predictions, we first analysed our results from the blocked schedule of the experiment. Here every participant has completed a block of 416 trials of hit condition and another block of 416 trials of stop condition, with the order counterbalanced across all participants. Our experimental results qualitatively match the predictions of our model. First, participants successfully differentiated between the kinematics of the hit and the stop condition, with both types of movements resulting in matching early and peak velocity, but with differences towards the end of the movement such that the velocity at the target is higher for the hit condition (Figure 3D). Specifically, in the hit condition participants produced movements with average velocity at the target of 38.5 cm/s, while successfully stopping at the target in the stop condition. In addition, similar to the model simulations, movements in the hit condition were of slightly shorter duration (630 ms vs 700 ms).

Qualitatively, the experimental feedback responses also match the model predictions (Figure 3EF). First, due to relatively fast reaching velocities in our experiment, as well as the lack of maintained perturbations, all perturbations were induced at short times-to-target (under 550 ms). For comparison, in our previous study (***Česonis and Franklin*** (***2020***)) perturbations were induced at times-to-target that ranged between 300 ms and 1000 ms, with peak feedback intensities recorded for perturbations with time-to-target at 400 ms. Second, both our data and the model produce feedback intensities at short times-to-target that are higher for the hit condition than for the stop condition, even in movement segments where the kinematics are otherwise similar. Importantly, we do not fit the model to match the data, but instead use it to qualitatively describe the relative regulation of stop and hit conditions. As such, matching features between the intensity profiles of the model (Figure 3BC) and the data (Figure 3EF) imply that similar computational mechanisms may be in action. Finally, our results also indicate that participants utilise different feedback controllers for the hit and stop conditions, as the experimental results for the blocked hit condition match the model simulations of the hit, rather than the long-stop condition.

Human participants utilise contextual switching of feedback controllers

In the previous sections we established the differences between the baseline control policies of hit and stop conditions. Here, we test how these policies change when the exposure to these conditions is no longer blocked. For example, it is natural in our daily activities to continuously switch between tasks, rather than doing a single task for many repetitions before switching to a new task. However, the question remains, how switching between different tasks affects the underlying feedback control policies. To test this, in the second half of the experiment we presented our participants with the same two types of movements (hit and stop), but now with the conditions randomly mixed across trials, instead of being presented in two separate blocks. As such, we could test for one of two possible outcomes:

1. Control policies for stop and hit movements in the mixed schedule match respectively the control policies in the stop and hit movements in the blocked schedule (Figure 1D). Such an outcome would indicate that participants are able to easily switch between different control policies (at least within consecutive trials).
2. Control policies for stop and hit movements in the mixed schedule do not match with the respective baseline policies, indicating interference when switching among multiple conditions (Figure 1C).

While both outcomes have previously been discussed from the sensorimotor adaptation perspective, to our knowledge they have not yet been demonstrated for feedback control.

Our participants successfully produced the movements required in the experiment (Figure 4A). Particularly, we observed clear distinctions in the terminal velocity between the hit and stop conditions, independent of the experimental schedule (blocked or mixed). A two-way repeated-measures ANOVA showed a significant main effect on condition (hit or stop, $F_{1,13} = 544.2$, $p \ll 0.001$), but no significant main effect on experiment schedule (blocked or mixed, $F_{1,13} = 0.710$, $p = 0.42$) or schedule/condition interactions ($F_{1,13} = 0.681$, $p = 0.42$). In addition, a complementary Bayesian repeated-measures ANOVA analysis showed similar results, with a very strong effect (***Raftery and Kass*** (***1995***)) of condition (hit or stop, $BF_{incl} = 1.6 \times 10^{25}$), and with a tendency towards no effect of schedule (blocked or mixed, $BF_{incl} = 0.379$), or condition/schedule interaction ($BF_{incl} = 0.409$). A similar analysis for peak velocities showed a significant main effect of condition (hit or stop, $F_{1,13} = 5.94$, $p = 0.03$; although $BF_{incl} = 1.12$ indicates not enough evidence to either reject or accept the null hypothesis) and condition/schedule interaction ($F_{1,13} = 19.3$, $p \ll 0.001$; $BF_{incl} = 32.6$), but not on schedule (blocked or mixed, $F_{1,13} = 1.52$, $p = 0.24$; $BF_{incl} = 0.56$ shows a weak tendency towards accepting null hypothesis). The Holm-Bonferroni corrected post-hoc analysis for the interaction

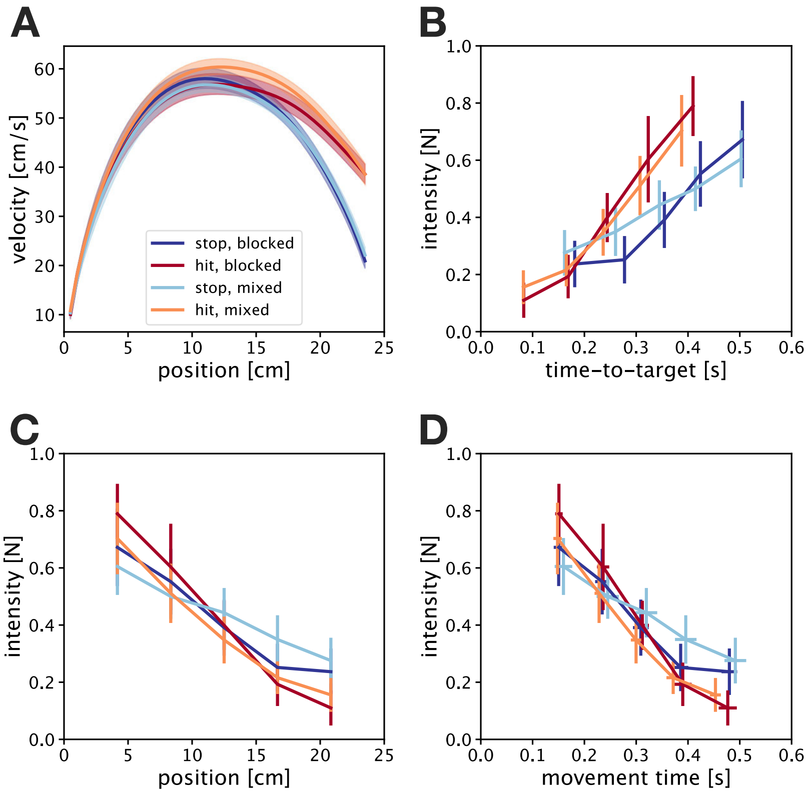

Figure 4. Experimental results of stop and hit conditions in both blocked and mixed schedules. **A.** Velocity profiles against position. Both stop conditions and both hit conditions produce respectively similar velocity profiles, showing that participants successfully performed the task in the mixed schedule. **B.** Feedback response intensities represented as a function of time-to-target. Hit and stop movements in the mixed schedule demonstrate differences when expressed against time-to-target, that match the differences between hit and stop conditions in the blocked schedule. This supports the hypothesis of contextual controller switching between multiple task-specific controllers. **C.** Feedback intensities in all four conditions show no differences when expressed against position or **D.** movement time at perturbation onset, as predicted by the OFC simulations. This questions the appropriateness of position or movement time as the reference frames in which to compare multiple feedback controllers. Error bars and shaded areas indicate 95% confidence intervals of the mean.

term revealed that participants produced slightly faster movements in the mixed-hit condition, with the peak velocities matching otherwise.

We examined the evolution of the experimental visuomotor responses as a function of perturbation onset position or onset time across the four different conditions (Figure 4CD). When expressed against either position or time, the visuomotor intensity profiles do not show the classical bell-shaped profile where strongest responses occur in the middle of the movement and are reduced towards the beginning and end. Instead, our participants produced the strongest responses for the earliest perturbations, induced at 1/6 of the total forward movement, with further responses decaying in intensity as perturbations occurred closer to the target. Moreover, we observed no significant differences in visuomotor responses across the different conditions and schedules. Three-way repeated-measures ANOVA with condition (stop or hit), schedule (blocked or mixed) and perturbation location (5 levels) as main factors showed no effect of condition ($F_{1,13} = 0.486$, $p = 0.50$; $BF_{incl} = 0.238$ shows substantial evidence towards no effect), schedule ($F_{1,13} = 0.096$, $p = 0.76$; $BF_{incl} = 0.142$ shows substantial evidence towards no effect) or condition/schedule interaction ($F_{1,13} = 0.657$, $p = 0.43$; $BF_{incl} = 0.305$ shows substantial evidence towards no effect). While we observed a significant main effect of the perturbation location ($F_{2.9,37.7} = 61.2$, $p \ll 0.001$ after Greenhouse-Geisser sphericity correction; $BF_{incl} = 9.3 \times 10^{36}$), such an effect was expected due to the temporal evolution of feedback responses. In addition, we observed a significant interaction between perturbation onset location and the condition ($F_{2.1,27.0} = 6.26$, $p = 0.005$ after the sphericity correction; $BF_{incl} = 6.86$), however a Holm-Bonferroni corrected post-hoc analysis on the interaction term did not indicate any meaningful interaction effects, with none of the significant interactions appearing at the same perturbation onset location. Finally, the remaining interactions of schedule/perturbation ($F_{2.6,33.9} = 2.67$, $p = 0.07$ after Greenhouse-Geisser sphericity correction; $BF_{incl} = 0.289$) and condition/schedule/perturbation ($F_{2.8,36.8} = 0.233$, $p = 0.86$ after Greenhouse-Geisser sphericity correction; $BF_{incl} = 0.075$) showed no significant effects. Thus, as a whole our analysis indicates that the feedback controllers could not be differentiated when expressed as a function of the position within the movement.

When expressed against time-to-target, the visuomotor feedback responses show decreasing feedback intensities with decreasing time-to-target, with responses virtually vanishing when the time-to-target approaches zero (Figure 4B). This behaviour is consistent with our previous models describing the time-gain relationship (***Česonis and Franklin*** (***2020***)). In addition, we observe stronger increases in visuomotor feedback intensity with increasing time-to-target for the hit condition compared to the stop condition, in both blocked and mixed schedules. Such regulation was previously predicted by our time-to-target OFC model (see Fig 9C in ***Česonis and Franklin*** (***2020***)) for short times-to-target. Finally, we also observe a qualitative match between the two stop conditions (mixed and blocked) as well as between the two hit conditions (mixed and blocked), suggesting first evidence of rapid feedback controller switching in the mixed schedule.

Qualitatively the increase of visuomotor response intensities with time-to-target for our specific results could be well described by a line function for each of the four combinations of condition and schedule. In order to get a quantitative estimate of the differences between the conditions we performed a Two-way ANCOVA analysis of visuomotor response intensity, with schedule and condition as the two factors, and time-to-target as the covariate. The results showed a significant main effect of condition (hit or stop, $F_{1,275} = 24.8$, $p \ll 0.001$; $BF_{incl} = 9.46 \times 10^3$), and time-to-target ($F_{1,275} = 222.8$, $p \ll 0.001$; $BF_{incl} = 1.04 \times 10^{33}$), but no effect of the experimental schedule (blocked or mixed, $F_{1,275} = 0.098$, $p = 0.75$; $BF_{incl} = 0.138$) or of schedule/condition interaction ($F_{1,275} = 1.06$, $p = 0.30$; $BF_{incl} = 0.304$ shows tendency towards no effect). Such results indicate that we can successfully separate the two different controllers when expressing their feedback response intensities (or their gains) against the time-to-target. Furthermore, we also show that such differences are only present when comparing the controllers for different tasks, and are not dependent on the presentation schedule of these tasks. Thus, we demonstrate that our participants successfully selected an appropriate controller for a hit or a stop task, even in a schedule where the task could change on consecutive trials.

Discussion

In this study we have demonstrated that humans are capable of rapid switching between appropriate feedback controllers in the presence of different contextual cues. Specifically, our participants show systematic differences in feedback responses when performing hitting movements, compared to reach-and-stop movements. Moreover, the same systematic differences are present, both when the two tasks are performed in isolation (blocked schedule), or when rapidly switching from one task to the other (mixed schedule), showing that these differences are evoked within a single trial, and not gradually adapted. Finally, these feedback responses are also well matched with the optimal feedback control predictions for these responses in hit and stop tasks, further reinforcing accumulating evidence of optimality principals in the feedback control of human movements.

In order to gain insight into computational mechanisms that are employed when humans switch between hit and stop conditions, in this study we formulate our hypothesis through normative modelling (***Harris and Wolpert*** (***1998***); ***Todorov and Jordan*** (***2002***); ***Liu and Todorov*** (***2007***); ***Verstynen and Sabes*** (***2011***); ***Rigoux and Guigon*** (***2012***); ***Berniker and Penny*** (***2019***)). Such an approach compares the behavioural experimental data to the results simulated computationally through a known bottom-up design. In turn, any mismatch between the data and the model rules out the mechanism, while matching behaviour provides support for the likelihood of such a mechanism. Specifically, here we simulate three different types of control movements: stop movement, where a point mass is stopped at a target 25 cm away from the start position; hit movement, where the point mass is instead brought to the same target with nonzero terminal velocity; and a long-stop movement, with similar kinematics to the hit movement within the 25 cm segment, generated by a stop movement to a secondary virtual target at 28 cm distance. The hit and stop simulations differed in the implementation of the feedback controller, with the state dependent costs for the terminal velocity and terminal acceleration reduced by a factor of 50 for the hit condition. As a result, the two models inherently simulate the behaviour that is achieved via different controllers. On the other hand, the long-stop condition was simulated by using the same controller as the stop condition, but to a target at 28 cm instead of 25 cm. Consequently, such a movement still maintained a non-zero velocity at 25 cm, virtually simulating a hit-like movement. Notably, in order to better match the kinematics of a long-stop movement to the kinematics of the hit and stop movements, we temporally modulated the activation cost R of the long-stop controller, which we have previously shown does not affect the overall feedback responses in terms of time-to-target (***Česonis and Franklin*** (***2020***)). In general, while kinematics of hit and long-stop models matched well, the two simulations predicted very different feedback response profiles when expressed both against time-to-target and against position. Finally, the responses of our participants in the hit condition matched better with the model simulation of the hit condition, rather than the simulation of the long-stop, providing evidence that humans use different feedback controllers for different tasks.

Principles of contextual switching have been extensively studied in the context of feedforward adaptation (***Wada et al.*** (***2003***); ***Lee and Schweighofer*** (***2009***); ***Howard et al.*** (***2013***); ***Sarwary et al.*** (***2015***); ***Howard et al.*** (***2017***); ***Forano and Franklin*** (***2020***); ***Forano et al.*** (***2021***)). While these cues vary in effectiveness (***Wada et al.*** (***2003***); ***Howard et al.*** (***2013***)) and are typically considered as relative weightings of multiple feedforward models (***Wolpert and Kawato*** (***1998***)), strong dynamic cues such as differences in follow-through (***Howard et al.*** (***2015***); ***Sheahan et al.*** (***2016***)), lead-in (***Howard et al.*** (***2012***, 2017)), or visual workspace (***Forano and Franklin*** (***2020***); ***Hirashima and Nozaki*** (***2012***)) can effectively separate the feedforward models. As multiple recent papers have demonstrated that voluntary (feedforward) and feedback control likely share neural circuits (***Ahmadi-Pajouh et al.*** (***2012***); ***Hayashi et al.*** (***2016***); ***Maeda et al.*** (***2018***, 2020); ***Voudouris and Fiehler*** (***2021***); ***Poscente et al.*** (***2021***)), it is reasonable to believe that similar contextual regulation would also be present in feedback control. However, studies that have shown this parallel changes in the feedback responses to the learning of the feedforward dynamics, either examined before and after adaptation to novel dynamics ***Wagner and Smith*** (***2008***); ***Ahmadi-Pajouh et al.*** (***2012***); ***Maeda et al.*** (***2018***, 2021), or during the process of adaptation ***Cluff and Scott*** (***2013***); ***Franklin et al.*** (***2017***); ***Franklin and Franklin*** (***2021***); ***Coltman and Gribble*** (***2020***), meaning that the they could not distinguish between the slow adaptation of the feedback controller to each condition or the rapid switching between two controllers. Moreover, other studies have suggested that feedforward and feedback controllers are learned separately and may even compete with one another (***Kasuga et al.*** (***2015***)), suggesting that these share different neural circuits and may have different properties. In this study we showed that in the mixed schedule, where the task goal unpredictably switched between hit or stop tasks on consecutive trials, participants evoked different control policies for each task. Furthermore, these policies, evoked within mixed schedule, well matched with the respective policies in the blocked schedule, suggesting that they were not only different from one another, but also appropriate for each task, showing the strong separation of the two contexts. While this is not unexpected, as the two hit and stop tasks are significantly different in their dynamics and thus should act as a strong contextual cue, one important result is that we demonstrated that the context regulates the feedback, and not only feedforward control. Finally, our results are also consistent with the accumulating evidence of the shared relationship between feedforward and feedback control in human reaching.

One reason why contextual effects on feedback control have not been broadly studied, is that it is difficult to quantify what really constitutes a change in feedback control policy. For example, we can trigger responses of different magnitudes by changing the size of the perturbation (***Franklin et al.*** (***2016***); ***Nashed et al.*** (***2012***)), inducing perturbations at different positions (***Dimitriou et al.*** (***2013***); ***Česonis and Franklin*** (***2020***)) or at different times (***Liu and Todorov*** (***2007***); ***Oostwoud Wijdenes et al.*** (***2011***, 2019)). However, computationally such differences in response intensity can be achieved within the same optimal feedback controller without ever changing control parameters. In a recent study ***de Comite et al.*** (***2021***) demonstrated that the control signal is modified during reaching when the task goal changes. While the study provided important insights into online modification of control in tasks conceptually similar to the change in context (i.e. obstacles appearing in the path of reaching), the change in the control signal alone does not directly imply the change in the feedback controller. That is, from the computational point of view, the control signal u is modulated not only by the control gain L, but also the observed state vector $\hat{x}$, which conventionally also includes the target representation. Thus, it is equally possible that the change in the control signal is achieved through an update of the target state representation, as it is through the recomputation of the control policies. Even though the results, described in the article, are indeed likely to be achieved through the update of the controller gains and not only through the update of the target, they also highlight the need of more direct methods to evaluate the change in the control. In our study we separate the two variables (controller gain L and state estimate $\hat{x}$) through normative modelling, where we make predictions on the human behaviour. Specifically, we simulate the behaviour either by recomputing the controller L (Figure 3BC, hit and stop), or by updating the state estimate $\hat{x}$ and using the same controller L (Figure 3BC, stop and long-stop), to compare with the experimental results (Figure 3EF). These results show that humans indeed change their control policies when the task goal (e.g. hit or stop) changes. Thus, by combining behavioural results with normative control models we can clearly identify that it is specifically the change in control, and not other mechanisms, that is responsible for the regulation observed in the experimental data.

Previous studies have demonstrated that visuomotor feedback intensity profiles are roughly bell-shaped along the movement – low at the beginning and the end, and peaking in the middle (***Dimitriou et al.*** (***2013***); ***Česonis and Franklin*** (***2020***)) – leading to assumptions that these gains might parallel the velocity (***Voudouris and Fiehler*** (***2021***); ***Poscente et al.*** (***2021***)). Our simulations and experimental results (Figure 3CF) demonstrate that this bell-shape profile is not fixed, and that other profiles are possible. In our previous work, we established a robust relationship between the visuomotor feedback intensities and time-to-target, demonstrating that time-to-target is the fundamental variable that modulates the responses, given that the task goal (and thus the feedback controller) remains the same (***Česonis and Franklin*** (***2020***)). This means that the bell-shaped profile is simply a by-product of a specific timing of perturbations, and is not regulated by their onset location. As a consequence, the shape of these feedback intensity profiles can be modulated away from the bell-shaped profile by changing movement speed, target distance or acceleration profile. Such results illustrate possible caveats in the experimental paradigms of motor control: historically, some of the task requirements have been largely consistent, particularly in terms of reaching distance, reaching speed or duration. This may result in some measured behavioural outcomes being specific to these kinematics or conditions rather than representing the general features of the motor control system. Thus, while we do not advocate for routinely altering the standard experimental and analytical methods, it is worth considering the specific biases that such methods may contribute to a given study.

One popular way of looking at the visuomotor responses in humans is how they vary with position in a movement. Indeed, numerous studies either analyse the evolution of responses against position (***Dimitriou et al.*** (***2013***); ***Česonis and Franklin*** (***2020***); ***Poscente et al.*** (***2021***)), or induce perturbations based on a fixed position (***Franklin and Wolpert*** (***2008***); ***Knill et al.*** (***2011***); ***Franklin et al.*** (***2016***); ***Gallivan et al.*** (***2016***); ***de Brouwer et al.*** (***2017***, 2018); ***Cross et al.*** (***2019***)), with the expectation that these perturbations induce similar responses unless the control changes. For example, **Knill et al.** (**2011**) demonstrated different feedback responses, induced at a matched position in movements towards different targets. While we believe that these different target properties indeed suggest different feedback controllers, such a distinction cannot be reliably tested with only one perturbation, matched by position. Our results clearly demonstrate the limitations of position as the main variable to probe such control. On one hand, even with similar kinematics for the majority of the movement, simulations of stop and long-stop movements predict radically different responses at matching positions (Figure 3C), despite the fact that these are generated with identical controllers. On the other hand, different controllers for hit and stop conditions still produced roughly matching feedback responses at the same position, consistent with the experimental data (Figure 3CF). In contrast to position as the main variable, OFC simulations in both this study and our previous work (**Česonis and Franklin** (**2020**)) show that the same controller, when expressed against time-to-target, produces matching response profiles, independent of other kinematic factors such as movement velocity or position of the perturbation onset (Figure 3B). Furthermore, different controllers, such as hit and stop, produce feedback responses with systematic differences when expressed against time-to-target, exactly as demonstrated by our participants. Thus, we propose that time-to-target is the better reference frame for comparing feedback responses.

In this study we have raised two alternative hypotheses about the regulation of feedback controllers within the mixed schedule. The first possibility is that the feedback control gradually adapts to a given task over a few consecutive trials, similar to the feedforward control during learning of a force field or visuomotor rotation. If such adaptation was true, we expect different feedback intensities between the hit and stop conditions in the blocked schedule as the controller has enough trials to reach steady-state behaviour. However, in the mixed schedule the controller would drift between the equilibrium of hit and stop conditions, producing similar responses for mixed hit and mixed stop conditions. Note that even in such a case where only a single controller is performing both hit and stop trials, we would not necessarily expect any effects on the kinematics or the participant's ability to complete the task. Instead, due to the feedback nature of the control, a sub-optimal controller would still complete the movement, but produce sub-optimal (e.g. more costly) responses in the presence of external disturbances. The second possibility is that an appropriate controller is selected before each movement based on the provided context, allowing immediate switching between tasks. In this case, the feedback intensity profiles would match for the same task, regardless of the schedule of their presentation. That is, we expect to see similarities between both hit conditions, as well as between both stop conditions, but differences between any two hit and stop conditions. Our experimental results strongly support the latter option, as we not only observe differences between mixed hit and mixed stop conditions, but also observe their respective match with the blocked conditions. While our results do not rule out the adaptation of feedback controllers in general, we do demonstrate that different optimal controllers can be rapidly selected and switched between for familiar tasks.

One important aspect of the relationship between feedforward and feedback control is that modulating one of them should affect the behaviour of the other. Indeed, previous work has demonstrated that human participants changed their feedback gains after adapting their feedforward models to novel dynamics (**Wagner and Smith** (**2008**); **Ahmadi-Pajouh et al.** (**2012**); **Franklin et al.** (**2012**); **Cluff and Scott** (**2013**); **Franklin et al.** (**2017**); **Maeda et al.** (**2018**); **Coltman and Gribble** (**2020**); **Maeda et al.** (**2021**); **Franklin and Franklin** (**2021**)). However, an adapted movement in the force field typically produces kinematics that are similar to those in baseline movements, suggesting that such change of gains is achieved at matching times-to-target, and with the same task goal. Thus, our proposed framework that the relation between feedback intensities and time-to-target is unique for a unique controller would predict that the feedback gains would remain unchanged. As a result, we can not directly explain this change of control gains, unless the feedback controller somehow changes during adaptation. One factor driving such a change is that adapted movements in the force field are more effortful than baseline movements, due to additional muscle activity required to compensate for the force. An increased effort in the context of OFC simulations would thus increase the model activation cost R, resulting in a change of optimal feedback gains and intensities at matching times-to-target. In addition, the presence of a force field likely influences the biomechanics of the movement (particularly the muscle viscosity b), changes the state transition due to the external dynamics (via state transition matrix A), and updates the state uncertainty (***Izawa et al.*** (***2008***)), resulting in the same controller being applied to a different control plant, and thus producing different responses. Moreover, if the controller is optimised to to this new control plant, adaptation will inevitably require a new feedback controller. Therefore, such changes in feedback control are expected, even though conventionally it appears that the task goal remains the same after adaptation to the novel dynamics.

Even though many recent studies use force channel trials (***Scheidt et al.*** (***2000***)) to accurately measure the visuomotor feedback responses (***Franklin and Wolpert*** (***2008***); ***Franklin et al.*** (***2012***)), often these brief perturbation trials are complemented with maintained perturbation trials (***Reichenbach et al.*** (***2013***, 2014); ***Reichenbach and Diedrichsen*** (***2015***); ***Franklin et al.*** (***2016***, 2017); ***de Brouwer et al.*** (***2017***, 2018); ***Česonis and Franklin*** (***2020***)). This is because brief perturbations within a channel trial are task-irrelevant, and can be ignored without compromising the task, whereas maintained perturbations strengthen these responses as they require an active correction for the participant to reach the target. However, we have recently shown that these maintained perturbations also force a non-trivial extension of the movement duration compared to the non-perturbed movement, and thus complicate the relationship between the perturbation onset location and the time-to-target. Hence, in order to consistently evaluate the control behaviour and its relation to the time-to-target, here we deliberately chose to only induce perturbations within the force channels and not to include the maintained perturbations. Although this generally decreases overall feedback intensities, our participants produced clear responses that exhibited the temporal evolution as predicted by the OFC model simulations.

In summary, here we again demonstrate that time-to-target (elsewhere referred to as urgency) (***Česonis and Franklin*** (***2020***); ***Crevecoeur et al.*** (***2013***); ***Poscente et al.*** (***2021***); ***Oostwoud Wijdenes et al.*** (***2011***, 2019)), and not position or velocity, act as a primary predictor for the feedback response intensity when the task goal is fixed. Moreover, when comparing multiple tasks, the time-to-target reference frame consistently separates the feedback control policies for these tasks – an outcome that fails when comparing two different controller gains within the position reference frame. While position within the movement, and velocity at the time of a perturbation, definitely influence the controller responses, our results clearly demonstrate that the effect of these variables on overall control may be somewhat exaggerated in the previous literature. For example, our participants produced temporal evolution of the responses to visual perturbations that neither paralleled the velocity, nor showed the typical variation with position (with peak responses achieved mid-movement), but could be explained by the time-to-target dependency that was predicted by OFC. In addition, participants were able to switch their feedback controller from one trial to another, demonstrating the principle of contextual switching for feedback control. Such switching, well known in feedforward control, further reinforces accumulating evidence of the shared connections between feedforward and feedback control. Most importantly, our results demonstrate that the visuomotor feedback control in humans not only follows the principles of optimal control for a singular task, but also selects an appropriate controller for that task upon presenting the relevant context.

Methods

Participants

Fourteen right-handed (***Oldfield*** (***1971***)) human participants (age 21-29 years, 5 females) with no known neurological diseases took part in the experiment. All participants were naïve to the purpose

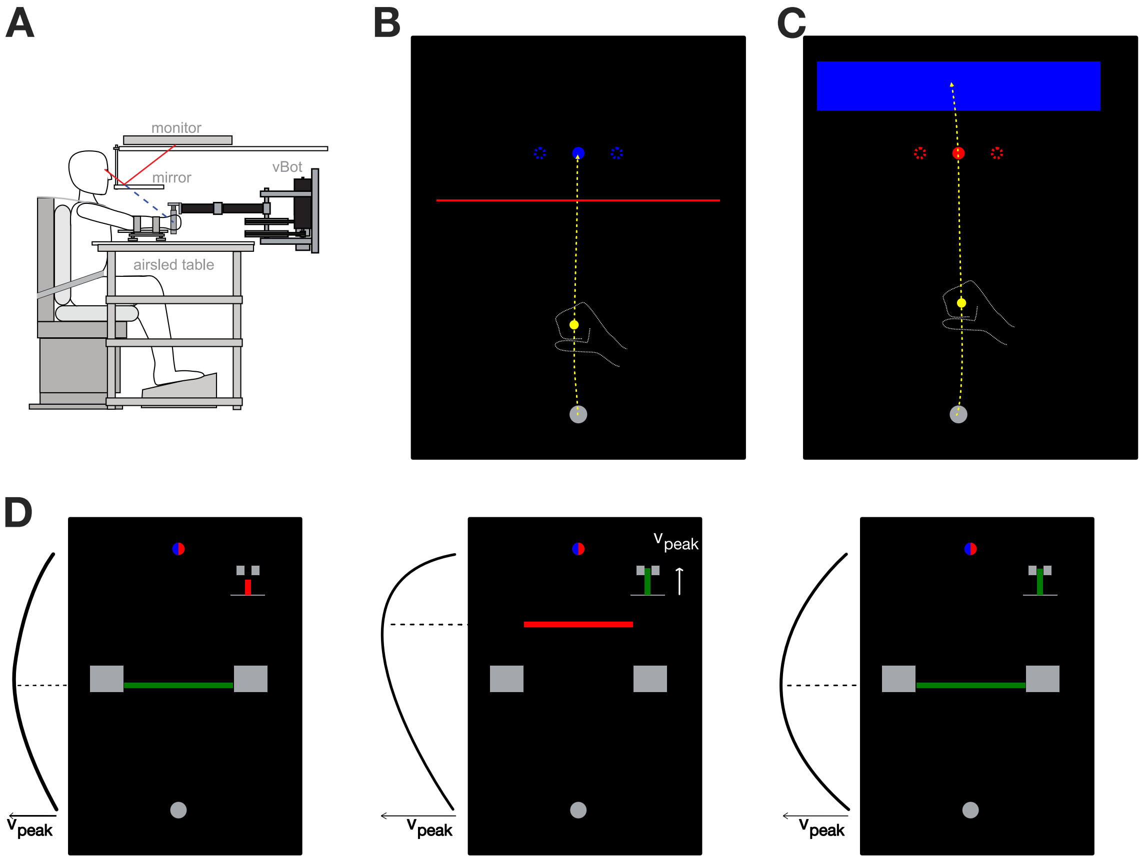

Figure 5. Experimental setup. **A.** Participants controlled a yellow cursor by moving a robotic handle. The cursor was projected via a screen-mirror system directly into the plane of the participant's hand. **B.** Stop condition. Participants were instructed to reach with the cursor through a red line and stop within the blue target. Target perturbations were occasionally induced via target jumps of 2 cm laterally. **C.** Hit condition. Participants were instructed to reach through the red target and stop within the blue area. Target perturbations (2 cm laterally) were again induced on random trials. **D.** Visual feedback was presented after each trial. Participants were shown the workspace with the start position and the target still present. In addition, two indicators were displayed. A bar chart at the top-right part of the workspace scaled proportionally with the absolute peak velocity, and was green if the velocity was within the required range as indicated by two grey brackets. A horizontal bar indicating the actual forward location where this peak velocity was achieved was displayed between the start and target positions. This bar was green if the peak location matched the experimental requirements, indicated by two large rectangular blocks. If both location and peak amplitude criteria were successfully fulfilled, participants were rewarded with one point. If at least one of the two criteria was not fulfilled, the respective indicator turned red instead of green, and no point was provided. In both hit and stop experiments participants were instructed to move through the red workspace element and stop at the blue, and were rewarded with one point if they both intercepted the target and fulfilled both velocity requirements.

of the study, and provided a written informed consent before participating. The study was approved by the Ethics Committee of the Medical Faculty of the Technical University of Munich.

Experimental setup

Participants performed forward reaching movements either to a target (stop condition) or through the target (hit condition) while grasping the handle of a robotic manipulandum (vBOT, ***Howard et al.*** (***2009***)) with their right hand, with their right arm supported on an air sled. Participants were seated in an adjustable chair and restrained using a four-point harness in order to limit the movement of the shoulder. A six-axis force transducer (ATI Nano 25; ATI Industrial Automation) measured the end-point forces applied by the participant on the handle. Position and force data were sampled at 1 kHz, while velocity information was obtained by differentiating the position over time. Visual feedback was provided via a computer monitor and mirror system, such that this system prevented direct vision of the hand and arm, and the virtual workspace appeared in the horizontal plane of the hand (Figure 5A). The exact timing of any visual stimulus presented to the participant was determined from the graphics card refresh signal.

Participants controlled a yellow cursor (circle of 1.0 cm diameter) by moving the robotic handle. The centre position of this cursor in the virtual workspace always matched the physical position of the handle. Every experimental trial was initiated when the cursor was brought into the start position (grey circle of 1.6 cm diameter), which was located 20 cm in front of participants' chest and centred with the body. When the cursor was within this start position, the circle changed from grey to white and the type of experimental trial was indicated by the presentation of a target. After a random delay, sampled from an exponential distribution with $\lambda = 0.7$ and truncated outside 1.0 s - 2.0 s interval, a tone was played to indicate the start of the movement. If participants failed to leave the start position within 1000 ms after this tone, the procedure of the current trial was aborted and restarted.

Over the course of the experiment the participants were tasked to complete two types of movements: stop movements, where they were required to stop within the target (a circle of 1.2 cm diameter, located 25.0 cm in front of the start position) (Figure 5B), or hit movements, where they had to intercept the target without stopping, and instead stop in a designated stopping area (a blue rectangle, [width, height] = [15 cm, 4 cm], centred 5 cm beyond the target) (Figure 5C). The reaching movement was considered complete once the centre of the cursor was maintained for 600 ms either within the area of the target in stop trials, or within the stopping area in the hit trials. In addition, if the movement duration was longer than 4.0 s, the trial was timed-out and had to be repeated. After each trial, the participant's hand was passively moved back to the start position by the vBOT, while the feedback of the previous trial was provided on screen (Figure 5D). All movements were self-paced, with short breaks provided every 208 trials, and a longer break (5-10 minutes) provided at the half-way point of the experiment.

Experimental paradigm

Participants performed reaching movements in four conditions – blocked stop, blocked hit, mixed stop and mixed hit – that were part of a single experiment. Across these conditions, participants were required to either reach to the target and stop there (the stop conditions), or to reach through the target and stop in the designated stopping area (hit conditions). In order to easily cue the distinction between the hit and stop conditions, the two types of trials had small visual differences. For the hit condition participants were presented with a red target (a red circle of 1.2 cm diameter) and a rectangular blue stopping area of dimensions 15 cm by 4 cm, centred 5 cm beyond the target (Figure 5C). For the stop condition participants were presented with a target that was otherwise identical to the target in hit condition, but was blue in colour, and with a horizontal, 15 cm wide red line, that was placed 3 cm before the target (Figure 5B). While this line had no functional interaction with the experiment, it allowed us to consistently instruct the participants to always perform reaching movements so that they intercept the red element in the workspace, and stop within the blue element.

In order to probe the visuomotor feedback responses of participants, during some reaching movements we briefly perturbed the target by shifting it 2.0 cm laterally for 250 ms before returning back to the original position (Figure 2A). These perturbed trials were always performed within the virtual mechanical channel, where participants were free to move along the line between the start position and the target, but were laterally constrained by a virtual viscoelastic wall with stiffness of 2 N/m and damping 4000 Ns/m (***Scheidt et al.*** (***2000***); ***Franklin and Wolpert*** (***2008***); ***Dimitriou et al.*** (***2013***)). As the perturbations were always task-irrelevant, this channel therefore did not obstruct participants to complete the trial. However, as participants still produced involuntary feedback responses due to the target shift, the virtual channel allowed us to record the forces that participants produced due to the perturbations and measure the intensities of the visuomotor feedback responses.

For each type of movement (i.e. hit or stop) there was a total of 11 different perturbations. Ten of these perturbations were cued during the reaching movement as participants crossed one of the five perturbation onset locations, equally spaced between the start position and the target position (4.2, 8.3, 12.5, 16.7, and 20.8 cm from the centre of the start position). At all of these five locations the target could either shift to the left or to the right. In addition, one zero-magnitude perturbation was also included, where the movement was simply performed within the channel without any target shift in order to probe the force profile of the natural movement. Finally, in addition to the perturbation trials we also included non-perturbed trials where participants simply reached towards the target without any target perturbation and without the virtual channel constraining the hand.

In order to present the different perturbations in a balanced manner, we combined different types of trials in blocks of 16 trials. One block of 16 trials contained 11 perturbed trials (5 perturbation onset locations x 2 directions, plus one neutral movement in the force channel), and 5 non-perturbed movements without the force channel. Each of the four experimental conditions consisted of 26 such blocks, with the order of trials fully randomised within each block, resulting in 416 trials per condition and 1664 trials overall.

In the first half of the experiment, participants were always presented with the two blocked-design conditions (blocked hit and blocked stop), with the order of the conditions balanced across the population of participants. That is, each participant started with 416 trials of stop trials, followed by 416 hit trials or vice-versa. In the second half of the experiment, the two final conditions – mixed hit and mixed stop – were presented in a pseudo-random order within the same blocks. While individual trials within mixed conditions were identical to the individual trials within the respective blocked conditions, they were now presented in a pseudo-randomised order. Specifically, the remaining 832 trials were divided into 26 blocks of 32 trials, with each block consisting of 16 hit and 16 stop trials fully randomised within this block.

Feedback regarding movement kinematics

In theory, the movements in hit condition could be interpreted as the movements, where the goal is to go through the via-point (the red target) and stop at the blue stopping area. As a result, such movements could simply be treated by participants as the stop movements with longer movement distance and a less restrictive target. Typically for such reaching movements, humans would produce a velocity profile that is bell-shaped, with peak velocity near the middle of the movement, and therefore further along the movement than in the stop condition. In order to avoid such differences and keep the velocity profiles comparable between the two conditions, we provided the task-relevant feedback on the velocity profiles, specifically the peak velocity and peak velocity location, to our participants (Figure 5D).

Independent of the experimental condition, participants were required to produce the movements with the peak velocity of 60 cm/s ± 8 cm/s, and the peak velocity location within 11.25 cm - 13.75 cm movement distance (or 45%-55% of the distance between the start location and the target). The peak velocity was indicated as the small bar chart at the top-right of the screen, with the required velocity range indicated by two grey brackets. If the velocity target was matched, the bar chart turned green, otherwise it was red. Similarly, the peak velocity location was shown as a horizontal bar, centred around the movement distance where the peak velocity was reached. If this location was within the target range (also indicated by grey brackets), it was displayed as green, otherwise it was red. Participants were rewarded one point if both velocity requirements were successfully met, and the cursor intercepted the target during the movement.

Data Analysis and Code Availability

All data was pre-processed for the analysis in MATLAB 2017b: force and kinematic time series were low-pass filtered with a tenth-order zero-phase-lag Butterworth filter with 15 Hz cutoff and resampled at 1 kHz to account for an occasional missed sample during the signal recording. All subsequent analysis was performed in Python 3.9.4 and JASP v0.14.1 (***JASP Team*** (***2021***)). First, raw visuomotor feedback intensities were calculated from the force responses, recorded after the induction of a target perturbation. Specifically, for every perturbation trial we averaged the lateral force response over a time window of 180 ms - 230 ms after the onset of the perturbation, and subtracted a neutral force profile over the matching time window. This method and the particular time window has now been used in numerous studies to calculate the intensity of the early involuntary visuomotor feedback response (***Franklin and Wolpert*** (***2008***); ***Dimitriou et al.*** (***2013***); ***Franklin et al.*** (***2016***); ***Česonis and Franklin*** (***2020***); ***Cross et al.*** (***2019***); ***de Brouwer et al.*** (***2017***, 2018)). As the direction of the response differed based on the perturbation direction, we reversed the direction of the intensities of responses to the leftward perturbations and grouped all intensities by the perturbation onset location. Second, we normalised mean feedback responses between 0 and 1 for each participant in order to avoid the group effect being biased towards participants with stronger responses. Finally, in our analysis the start of all movements was defined as the last time sample where the cursor is still within the area of the start circle, and the end of the movement was defined as the last time sample before the cursor enters the target circle. Time-to-target values were extracted from the data for every perturbation trial by subtracting the perturbation onset time from the movement end time.

In this article we provide two types of statistical analysis: the conventional frequentist statistics, as well as complementary Bayesian analysis that is presented as Bayesian factors (***Raftery and Kass*** (***1995***)), which instead of a simple hypothesis testing provides evidence for or against the null hypothesis. As a result, among other things, Bayesian analysis allows us to distinguish between accumulating evidence for the null hypothesis, and simply lacking evidence in either direction due to low power or small sample size.

All the Jupyter notebooks for the data analysis, pre-processed experimental data and statistical analysis conducted in this article are available at https://figshare.com/s/ddc74e11dc25d04457ae.

Computational modelling

In this study we formulated our initial hypothesis about the feedback control mechanisms in humans by first simulating the behaviour of the optimal feedback controller (OFC). Specifically, we used a finite-horizon linear-quadratic regulator framework – a relatively simple OFC that assumes perfect sensory input, as well as no control-dependent noise, while still being able to capture a significant part of the variance of human reaching movements (***Kuo*** (***1995***); ***Česonis and Franklin*** (***2021***)). In order to model the feedback behaviour of our human participants, we first simulated virtual movements of a point mass with $m = 1$ kg, and an intrinsic muscle damping $b = 0.1$ Ns/m. This point mass was controlled in two dimensions by two orthogonal force actuators that simulated muscles, and regulated by a control signal u_t via a first-order low-pass filter with a time constant $\tau = 0.06$ s. At time t within the movement, such system could be described by the state transition model:

$$x_{t+1} = Ax_t + B(u_t + \xi_t), \tag{1}$$

where A is a state transition matrix, B is a control matrix, and ξ_t is additive control noise. For one spacial dimension A and B are defined in discrete time as:

$$A = \begin{bmatrix} 1 & \delta t & 0 \\ 0 & 1 - b\delta t/m & \delta t/m \\ 0 & 0 & 1 - \delta t/\tau \end{bmatrix}$$

,

$$B = \begin{bmatrix} 0 \\ \delta t/\tau \\ 0 \end{bmatrix}$$

Finally, to simulate our model in discrete time we used the sampling rate $\delta t = 0.01$ s

State x_t exists in the Cartesian plane and consists of position $\mathbf{p}$, velocity $\mathbf{v}$ and force $\mathbf{f}$ (two dimensions each). The control signal u_t is produced via the feedback control law:

$$u_t = -Lx_t \tag{2}$$

where L is a matrix of optimal feedback control gains, obtained by optimising the performance index (also known as the cost function):

$$J = \sum_{t=0}^{N} x_t^T Q_t x_t + u_t^T R_t u_t = \sum_{t=0}^{N} \omega_{p,t}(\mathbf{p}_t - \mathbf{p}^*)^2 + \omega_{v,t}||\mathbf{v}_t||^2 + \omega_{f,t}||\mathbf{f}_t||^2 + \omega_{r,t}||u_t||^2. \tag{3}$$

Here $x_t^T Q x_t$ and $u_t^T R u_t$ are two components of the total cost, known as state-cost and a control-cost respectively. In addition, ω_p, ω_v and ω_f are position, velocity and force state cost parameters, $\mathbf{p}^*$ is a target position, ω_r is the activation cost parameter and N is the duration of the movement, here required as a model input. Within the finite-horizon formulation, the cost parameters can be non-stationary and thus be different for every time-point. However, in our simulations we set $Q = 0$ for $t \neq N$, consistent with (***Todorov*** (***2005***); ***Liu and Todorov*** (***2007***)).

In this study we simulate three different controllers that we call stop, hit and long-stop. While the stop and long-stop controllers are derived from the identical set of costs state-costs Q, they are used for slightly different movements (25 cm and 700 ms for stop, 28 cm and 800 ms for long-stop). We used ω_p = 1.5, ω_v = 1 and ω_f = 0.1 as the values for the state cost parameters in this model, and the activation cost R = 3 $\times 10^{-6}$. Furthermore, in order to better match the forward velocity profiles, we also introduced a non-stationarity in the activation cost R of the long stop movement, where the total integral of the activation cost over the movement is not changed, but this cost develops over time during the movement. Specifically, at a time t in the trial, the activation cost for the long-stop movement was computed by:

$$R_{long-stop}(t) = RC(t), \tag{4}$$

where

$$C(t) \propto \exp\left(p\frac{t+q}{r}\right), \tag{5}$$

and the mean of $C(t)$ equals 1 for the duration of the trial, so that $R_{long-stop}$ produces the same amount of activation as R over the duration of the trial. Here p = 1, q = -1000 and r = 65 are constants, fit via trial and error in order to produce the forward velocity profile of long-stop condition that matches the velocity of stop and hit conditions. We have previously shown that such modulation only affects the kinematics of the movement, but does not change the feedback responses when expressed against the time-to-target (***Česonis and Franklin*** (***2021***)). On the other hand, in order to incentivise the hit controller to produce faster movements at the target, we reduced the cost parameters for terminal velocity and terminal force by a factor of 50. As a result, such controller produced hit-like movements that were aimed directly at a target, positioned at 25 cm distance, over 620 ms, which matched the kinematics of the long-stop controller over this movement segment.

Finally, for each controller we simulated feedback response intensity profiles along the movement, which we then used to compare the control policies predicted by each controller. To do so, we induced lateral target perturbations of 2 cm magnitude during the simulated movement to the target and recorded the corrective force, produced by each controller as a result of these perturbations. While in the experimental study we only induced such perturbations at five different onsets due to practical reasons, in our simulations we could perturb the movements at every point in time and fully map the response intensity profiles over the movement. Thus, for each model we simulated different movements with perturbations at each movement time-step (i.e. every 10 ms), with one perturbation only happening once per movement. In addition, to simulate the visuomotor delay that is present in humans, we delayed the onset of each perturbation by 150 ms, so that for the perturbation triggered at time t, the target is shifted at time $t + 150$ ms. We then averaged the force, produced by our model over a time interval 10 ms – 60 ms after the target was shifted (160 ms – 210 ms after the perturbation was triggered), representing the visuomotor response window of 180 ms – 230 ms in human subjects. Note that we used an earlier window for the model simulations than for the human subjects as the responses in the simulations ramp up fast due to muscles simplified to a single low pass filter.

Acknowledgements

We thank Hanna Hoogen, Isabelle Hoxha and Oliver Gerke for contributions to preliminary projects related to this manuscript. We thank Clara Günter, Jing Zhang, Sae Franklin, and Marion Forano for their feedback on this manuscript.

References

Ahmadi-Pajouh, M. A., Towhidkhah, F., and Shadmehr, R. (2012). Preparing to reach: Selecting an adaptive long-latency feedback controller. *Journal of Neuroscience*, 32(28):9537–9545.

Baram, Y. and Miller, A. (2007). Auditory feedback control for improvement of gait in patients with Multiple Sclerosis. *Journal of the Neurological Sciences*, 254(1-2):90–94.

Barr, C. C., Schultheis, L. W., and Robinson, D. A. (1976). Voluntary, non-visual control of the human vestibulo-ocular reflex. *Acta Oto-Laryngologica*, 81(5-6):365–375.

Berniker, M. and Penny, S. (2019). A normative approach to neuromotor control. *Biological Cybernetics*, 113(1-2):83–92.

Bock, O., Schneider, S., and Bloomberg, J. (2001). Conditions for interference versus facilitation during sequential sensorimotor adaptation. *Experimental Brain Research*, 138(3):359–365.

Česonis, J. and Franklin, D. W. (2020). Time-to-target simplifies optimal control of visuomotor feedback responses. *eNeuro*, 7(2):1–17.

Česonis, J. and Franklin, D. W. (2021). Mixed-horizon optimal feedback control as a model of human movement. *Neurons, Behavior, Data analysis, and Theory*, pages 1–36.

Cluff, T. and Scott, S. H. (2013). Rapid Feedback Responses Correlate with Reach Adaptation and Properties of Novel Upper Limb Loads. *Journal of Neuroscience*, 33(40):15903–15914.

Coltman, S. K. and Gribble, P. L. (2020). Time course of changes in the long-latency feedback response parallels the fast process of short-term motor adaptation. *Journal of Neurophysiology*, 124(2):388–399. PMID: 32639925.

Crevecoeur, F., Kurtzer, I., Bourke, T., and Scott, S. H. (2013). Feedback responses rapidly scale with the urgency to correct for external perturbations. *Journal of Neurophysiology*, 110(6):1323–1332.

Cross, K. P., Cluff, T., Takei, T., and Scott, S. H. (2019). Visual Feedback Processing of the Limb Involves Two Distinct Phases. *Journal of Neuroscience*, 39(34):6751–6765.

Day, B. L. and Lyon, I. N. (2000). Voluntary modification of automatic arm movements evoked by motion of a visual target. *Experimental Brain Research*, 130(2):159–168.

de Brouwer, A. J., Gallivan, J. P., and Flanagan, J. R. (2018). Visuomotor feedback gains are modulated by gaze position. *Journal of Neurophysiology*, 120(5):2522–2531.

de Brouwer, A. J., Jarvis, T., Gallivan, J. P., and Flanagan, J. R. (2017). Parallel Specification of Visuomotor Feedback Gains during Bimanual Reaching to Independent Goals. *Eneuro*, 4(2):ENEURO.0026–17.2017.

de Comite, A., Crevecoeur, F., and Lefèvre, P. (2021). Online modification of goal-directed control in human reaching movements. *Journal of neurophysiology*, 125(5):1883–1898.

Diedrichsen, J. (2007). Optimal Task-Dependent Changes of Bimanual Feedback Control and Adaptation. *Current Biology*, 17(19):1675–1679.

Dimitriou, M. (2016). Enhanced Muscle Afferent Signals during Motor Learning in Humans. *Current Biology*, 26(8):1062–1068.

Dimitriou, M., Wolpert, D. M., and Franklin, D. W. (2013). The Temporal Evolution of Feedback Gains Rapidly Update to Task Demands. *Journal of Neuroscience*, 33(26):10898–10909.

Forano, M. and Franklin, D. W. (2020). Timescales of motor memory formation in dual-adaptation. *PLoS Computational Biology*, 16(10):1–33.

Forano, M., Schween, R., Taylor, J. A., Hegele, M., and Franklin, D. W. (2021). Direct and indirect cues can enable dual adaptation, but through different learning processes. *Journal of Neurophysiology*, 126(5):1490–1506. PMID: 34550024.

Franklin, D. W., Reichenbach, A., Franklin, S., and Diedrichsen, J. (2016). Temporal Evolution of Spatial Computations for Visuomotor Control. *Journal of Neuroscience*, 36(8):2329–2341.

Franklin, D. W. and Wolpert, D. M. (2008). Specificity of Reflex Adaptation for Task-Relevant Variability. *Journal of Neuroscience*, 28(52):14165–14175.

Franklin, S. and Franklin, D. W. (2021). Feedback Gains modulate with Motor Memory Uncertainty. *Neurons, Behavior, Data analysis, and Theory*, 5(2):1–28.

Franklin, S., Wolpert, D. M., and Franklin, D. W. (2012). Visuomotor feedback gains upregulate during the learning of novel dynamics. *Journal of Neurophysiology*, 108(2):467–478.

Franklin, S., Wolpert, D. M., and Franklin, D. W. (2017). Rapid visuomotor feedback gains are tuned to the task dynamics. *Journal of Neurophysiology*, page jn.00748.2016.

Gallivan, J. P., Logan, L., Wolpert, D. M., and Flanagan, J. R. (2016). Parallel specification of competing sensorimotor control policies for alternative action options. *Nature Neuroscience*, 19(2):320–326.

Harris, C. M. and Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394:780.

Hayashi, T., Yokoi, A., Hirashima, M., and Nozaki, D. (2016). Visuomotor Map Determines How Visually Guided Reaching Movements are Corrected Within and Across Trials. *eNeuro*, 3(3):1–13.

Hirashima, M. and Nozaki, D. (2012). Distinct motor plans form and retrieve distinct motor memories for physically identical movements. *Current Biology*, 22(5):432–436.

Houk, J. C. (1976). An Assessment of Stretch Reflex Function. *Progress in Brain Research*, 44(C):303–314.

Howard, I. S., Ford, C., Cangelosi, A., and Franklin, D. W. (2017). Active lead-in variability affects motor memory formation and slows motor learning. *Scientific Reports*, 7(1):1–12.

Howard, I. S., Ingram, J. N., Franklin, D. W., and Wolpert, D. M. (2012). Gone in 0.6 seconds: The encoding of motor memories depends on recent sensorimotor states. *Journal of Neuroscience*, 32(37):12756–12768.

Howard, I. S., Ingram, J. N., and Wolpert, D. M. (2009). A modular planar robotic manipulandum with end-point torque control. *Journal of Neuroscience Methods*, 181(2):199–211.

Howard, I. S., Wolpert, D. M., and Franklin, D. W. (2013). The effect of contextual cues on the encoding of motor memories. *Journal of Neurophysiology*, 109(10):2632–2644.

Howard, I. S., Wolpert, D. M., and Franklin, D. W. (2015). The value of the follow-through derives from motor learning depending on future actions. *Current Biology*, 25(3):397–401.

Izawa, J., Rane, T., Donchin, O., and Shadmehr, R. (2008). Motor Adaptation as a Process of Reoptimization. *Journal of Neuroscience*, 28(11):2883–2891.

Izawa, J. and Shadmehr, R. (2008). On-Line Processing of Uncertain Information in Visuomotor Control. *Journal of Neuroscience*, 28(44):11360–11368.

JASP Team (2021). JASP (Version 0.14.1)[Computer software].

Kasuga, S., Telgen, S., Ushiba, J., Nozaki, D., and Diedrichsen, J. (2015). Learning feedback and feedforward control in a mirror-reversed visual environment. *Journal of Neurophysiology*, 114(4):2187–2193.

Knill, D. C., Bondada, A., and Chhabra, M. (2011). Flexible, Task-Dependent Use of Sensory Feedback to Control Hand Movements. *Journal of Neuroscience*, 31(4):1219–1237.

Kuo, A. (1995). An optimal control model for analyzing human postural balance. *IEEE Transactions on Biomedical Engineering*, 42(1):87–101.

Lee, J.-Y. and Schweighofer, N. (2009). Dual Adaptation Supports a Parallel Architecture of Motor Memory. *Journal of Neuroscience*, 29(33):10396–10404.

Liu, D. and Todorov, E. (2007). Evidence for the Flexible Sensorimotor Strategies Predicted by Optimal Feedback Control. *Journal of Neuroscience*, 27(35):9354–9368.

Maeda, R. S., Cluff, T., Gribble, P. L., and Pruszynski, J. A. (2018). Feedforward and feedback control share an internal model of the arm's dynamics. *Journal of Neuroscience*, 38(49):10505–10514.

Maeda, R. S., Gribble, P. L., and Pruszynski, J. A. (2020). Learning New Feedforward Motor Commands Based on Feedback Responses. *Current Biology*, 30(10):1941–1948.e3.

Maeda, R. S., Kersten, R., and Pruszynski, J. A. (2021). Shared internal models for feedforward and feedback control of arm dynamics in non-human primates. *European Journal of Neuroscience*, 53(5):1605–1620.

Nashed, J. Y., Crevecoeur, F., and Scott, S. H. (2012). Influence of the behavioral goal and environmental obstacles on rapid feedback responses. *Journal of Neurophysiology*, 108(4):999–1009.

Oldfield, R. (1971). The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia*, 9(1):97–113.

Oostwoud Wijdenes, L., Brenner, E., and Smeets, J. B. (2013). Comparing online adjustments to distance and direction in fast pointing movements. *Journal of Motor Behavior*, 45(5):395–404.

Oostwoud Wijdenes, L., Brenner, E., and Smeets, J. B. J. (2011). Fast and fine-tuned corrections when the target of a hand movement is displaced. *Experimental Brain Research*, 214(3):453–462.

Oostwoud Wijdenes, L., Van Beers, R. J., and Medendorp, W. P. (2019). Vestibular modulation of visuomotor feedback gains in reaching. *Journal of Neurophysiology*, 122(3):947–957.

Oscari, F., Secoli, R., Avanzini, F., Rosati, G., and Reinkensmeyer, D. J. (2012). Substituting auditory for visual feedback to adapt to altered dynamic and kinematic environments during reaching. *Experimental Brain Research*, 221(1):33–41.

Poscente, S. V., Peters, R. M., Cashaback, J. G., and Cluff, T. (2021). Rapid Feedback Responses Parallel the Urgency of Voluntary Reaching Movements. *Neuroscience*.

Prablanc, C. and Martin, O. (1992). Automatic control during hand reaching at undetected two-dimensional target displacements. *Journal of Neurophysiology*, 67(2):455–469.

Pruszynski, J. A., Gribble, P. L., and Corneil, B. D. (2018). A rapid visuomotor response on the human upper limb is selectively influenced by implicit motor learning. *Journal of Neurophysiology*, 121(1):85–95.

Pruszynski, J. A. and Scott, S. H. (2012). Optimal feedback control and the long-latency stretch response. *Experimental Brain Research*, 218(3):341–359.

Raftery, A. E. and Kass, R. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795.

Reichenbach, A., Costello, A., Zatka-Haas, P., and Diedrichsen, J. (2013). Mechanisms of responsibility assignment during redundant reaching movements. *Journal of Neurophysiology*, 109(8):2021–2028.

Reichenbach, A. and Diedrichsen, J. (2015). Processing reafferent and exafferent visual information for action and perception. *Journal of Vision*, 15(8):1–12.

Reichenbach, A., Franklin, D. W., Zatka-Haas, P., and Diedrichsen, J. (2014). A dedicated binding mechanism for the visual control of movement. *Current Biology*, 24(7):780–785.

Reynolds, R. F. and Day, B. L. (2012). Direct visuomotor mapping for fast visually-evoked arm movements. *Neuropsychologia*, 50(14):3169–3173.

Rigoux, L. and Guigon, E. (2012). A Model of Reward- and Effort-Based Optimal Decision Making and Motor Control. *PLoS Computational Biology*, 8(10).

Saijo, N., Murakami, I., Nishida, S., and Gomi, H. (2005). Large-Field Visual Motion Directly Induces an Involuntary Rapid Manual Following Response. *Journal of Neuroscience*, 25(20):4941–4951.

Sarwary, A. M. E., Stegeman, D. F., Selen, L. P. J., and Medendorp, W. P. (2015). Generalization and transfer of contextual cues in motor learning. *Journal of Neurophysiology*, 114(3):1565–1576.

Saunders, J. A. and Knill, D. C. (2003). Humans use continuous visual feedback from the hand to control fast reaching movements. *Experimental Brain Research*, 152(3):341–352.

Saunders, J. A. and Knill, D. C. (2005). Humans use continuous visual feedback from the hand to control both the direction and distance of pointing movements. *Experimental Brain Research*, 162(4):458–473.

Scheidt, R. A., Reinkensmeyer, D. J., Conditt, M. A., Rymer, W. Z., and Mussa-Ivaldi, F. A. (2000). Persistence of motor adaptation during constrained, multi-joint, arm movements. *Journal of neurophysiology*, 84(2):853–862.

Shadmehr, R., Brashers-Krug, T., and Mussa-Ivaldi, F. A. (1995). Interference in Learning Internal Models of Inverse Dynamics in Humans. *Advances in Neural Information Processing Systems 7*, 7:1117–1124.

Sheahan, H. R., Franklin, D. W., and Wolpert, D. M. (2016). Motor Planning, Not Execution, Separates Motor Memories. *Neuron*, 92(4):773–779.

Sing, G. C. and Smith, M. A. (2010). Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation. *PLoS Computational Biology*, 6(8).

Tabak, S. and Collewijn, H. (1994). Human vestibulo-ocular responses to rapid, helmet-driven head movements. *Experimental Brain Research*, 102(2):367–378.

Todorov, E. (2005). Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation*, 17(5):1084–1108.

Todorov, E. and Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235.

Verstynen, T. and Sabes, P. N. (2011). How Each Movement Changes the Next: An Experimental and Theoretical Study of Fast Adaptive Priors in Reaching. *Journal of Neuroscience*, 31(27):10050–10059.

Voudouris, D. and Fiehler, K. (2021). Dynamic temporal modulation of somatosensory processing during reaching. *Scientific Reports*, 11(1):1–12.

Wada, Y., Kawabata, Y., Kotosaka, S., Yamamoto, K., Kitazawa, S., and Kawato, M. (2003). Acquisition and contextual switching of multiple internal models for different viscous force fields. *Neuroscience Research*, 46(3):319–331.

Wagner, M. J. and Smith, M. A. (2008). Shared internal models for feedforward and feedback control. *Journal of Neuroscience*, 28(42):10663–10673.

Wolpert, D. M. and Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, 11(7-8):1317–1329.

Zhang, Y., Brenner, E., Duysens, J., Verschueren, S., and Smeets, J. B. (2018). Postural responses to target jumps and background motion in a fast pointing task. *Experimental Brain Research*, 236(6):1573–1581.