# Time-to-target explains task-dependent modulation of temporal feedback gain evolution

**Justinas Česonis**[1] **and David W. Franklin**[1*]

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

[1]Technical University of Munich, Germany

## Abstract

Visuomotor feedback gains vary in magnitude throughout a reach, commonly explained by optimal control. However, this pattern also matches the hand velocity profile, suggesting a simpler, perhaps hardwired, controller. Here we test between these alternatives, examining whether movement velocity regulates visuomotor gains. Feedback gain profiles were modulated across five conditions, with significant effects for altered hand, but not cursor, kinematics. Using optimal control we show that visuomotor feedback gains exhibit a consistent relation with time-to-target, rather than velocity, across a range of kinematics. Although our model suggests a time beyond which feedback gains should be down-regulated to minimise cost, our participants deliberately extend their movement durations for late perturbations, where feedback gains are reduced due to the lack of time for correction. Finally, we qualitatively replicate participant results within our model, but only when accounting for this extension of time-to-target, suggesting that expected time-to-target is critical in online feedback control.

## Introduction

From intercepting a basketball pass between opponents to catching a vase accidentally knocked off the shelf – visuomotor feedback responses play a familiar role in human motor behaviour. Previous research has extensively analysed these responses in human reaching movements (***Day and Lyon*** (***2000***); ***Reichenbach et al.*** (***2014***); ***de Brouwer et al.*** (***2017***, 2018); ***Saunders and Knill*** (***2003***); ***Saunders*** (***2004***); ***Saunders and Knill*** (***2005***); ***Sarlegna et al.*** (***2003***); ***Knill et al.*** (***2011***)), and showed an interesting combination of task-dependent variability on the timescale of a single movement (***Dimitriou et al.*** (***2013***); ***Franklin et al.*** (***2014***, 2017)), as well as sub-voluntary feedback onset times (***Prablanc and Martin*** (***1992***); ***Day and Lyon*** (***2000***); ***Franklin and Wolpert*** (***2008***)). These visuomotor feedback gains have been shown to modulate throughout a movement depending on the perturbation onset location (***Dimitriou et al.*** (***2013***)). This observation was explained through optimality principles, however such control was modelled only indirectly, by replicating velocity profiles and trajectories of visually perturbed movements (***Liu and Todorov*** (***2007***); ***Rigoux and Guigon*** (***2012***)). In this study we test to what degree optimal feedback control, as opposed to other control methods, can be used to model the visuomotor feedback gains directly.

As with any model, there are limitations to optimal feedback control. For example, by itself stochastic optimal feedback control generates an infinite number of reasonable solutions to given boundary conditions (***Liu and Todorov*** (***2007***); ***Guigon et al.*** (***2008***)). In order to then select the singular model output, additional "screening" of these solutions is required by, for example, specifying movement duration as the model parameter, rather than an emerging property (***Todorov and Weiwei Li*** (***2005***). More recently, this particular issue was addressed through combining optimal

control with utility of movement (***Rigoux and Guigon*** (***2012***)) to predict the optimal duration for any movement. In addition, terminal optimal feedback control has also received attention as an alternative that addresses some classical limitations (***Guigon et al.*** (***2007***, 2008). Nevertheless, optimal control as a theory of human movement has normally been compared against other theories in terms of prediction of kinematics and dynamics (***Todorov and Jordan*** (***2002***); ***Izawa et al.*** (***2008***); ***Nagengast et al.*** (***2009***); ***Yeo et al.*** (***2016***). However, optimal feedback control has been used to motivate extensive studies investigating the control and task-dependent modulation of feedback responses (***Knill et al.*** (***2011***); ***Pruszynski and Scott*** (***2012***); ***Nashed et al.*** (***2012***, 2014)). The results of these and other studies have highlighted the flexibility of modulation of these feedback responses. While a few studies have compared the predictions of the controller feedback gains against the feedback responses in human subjects (***Knill et al.*** (***2011***)), such predictions have not been made about the temporal evolution of these feedback gains during reaching. For example, ***Dimitriou et al.*** (***2013***) show temporal evolution of feedback gains throughout a reaching movement which peak in the middle of the movement and suggest that this is similar to the feedback gain predictions of ***Liu and Todorov*** (***2007***). However a direct comparison of these feedback responses has not been made. Here we directly compare the temporal evolution of visuomotor feedback gains in human participants with the prediction of these gains in an optimal feedback control model.

Visuomotor feedback gains over a goal directed reaching movement follow a roughly bell-shaped profile, with peak gain in the middle and decay towards the beginning and the end of the movement (***Dimitriou et al.*** (***2013***)). However, this pattern of visuomotor feedback gains is qualitatively similar to the development of hand velocity profiles along the same movement. A simpler explanation might therefore be that these feedback gains vary with the velocity signals of the movement. This result posits an alternative to optimal feedback control: what if there exists a much simpler control system that does not require optimisation but instead can instantaneously modify the feedback gains? Due to this similarity, we test whether this simple, but instantaneous, control system could be using a scaling of movement velocity at the instant of the perturbation onset to generate visuomotor feedback gains. Moreover, if movement kinematics drive such control, then the question is whether haptic or visual kinematics have a stronger influence. On the other hand, in stochastic optimal feedback control theory, the movement costs are defined as time dependent variables. Therefore, we hypothesise that a relation to the time-to-target may explain visuomotor feedback gains in humans. To test our hypotheses, we devised an experimental paradigm where we offset the usual bell-shaped velocity profile in the aim to offset the visuomotor feedback gains and modify the times-to-target across conditions. Finally, we compare these results with a normative optimal feedback control model of visuomotor feedback gain in order to better understand how and whether these gains can be the result of optimality and still maintain rapid responses.

# Results

## Experimental results

In this study we tested whether a simple relationship between movement kinematics and the visuomotor feedback gains is used by the human visuomotor system. To do so, we devised an experiment consisting of five different kinematic conditions. The baseline condition required movements with a natural, bell-shaped velocity profile, while the velocity profiles were modified for the four other conditions. In these four conditions we introduced a manipulation between the hand velocity and the cursor velocity in the forward direction, such that the cursor and hand had different velocity profiles, but their positions matched at the start and end of the movement (Figure 1). Two of these four conditions (matched-cursor conditions) required different kinematics of the physical movement to successfully complete the task, but the cursor velocity profiles matched the baseline. The two other conditions (matched-hand conditions) required the same hand movement as for the baseline condition, but as a result the cursor moved with different velocity profiles (see Materials and Methods). For each condition we measured the visuomotor feedback gains at

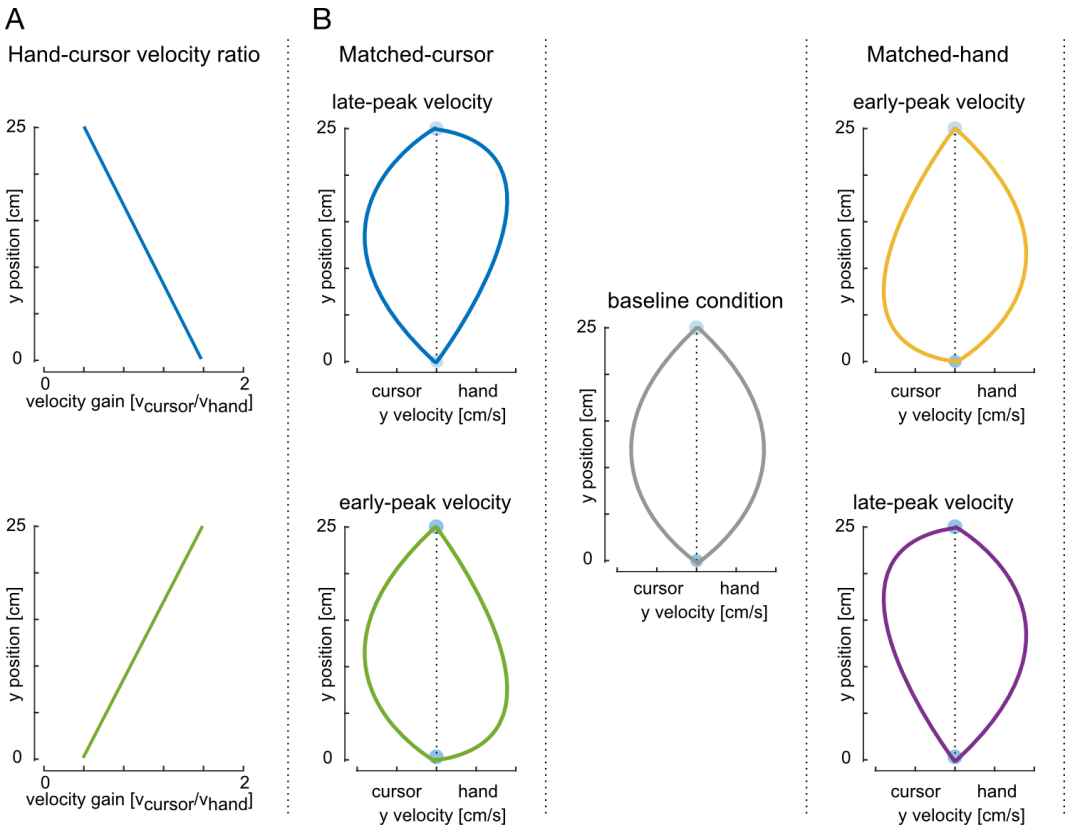


**Figure 1.** Experimental design. (**A**) Top: hand-cursor velocity scaling for conditions where the cursor position leads the hand position in y axis (matched-cursor late-peak hand velocity condition, blue, and matched-hand early-peak cursor velocity condition, yellow). Bottom: hand-cursor velocity scaling for conditions where the cursor position lags the hand position in y axis (matched-cursor early-peak hand velocity condition, green, and matched-hand late-peak cursor velocity condition, purple). (**B**) Hand and cursor velocity-position profiles required to achieve the ideal movement to the target. Left: matched-cursor velocity conditions; middle: baseline condition, where cursor position and hand position are consistent; right: matched-hand velocity conditions.

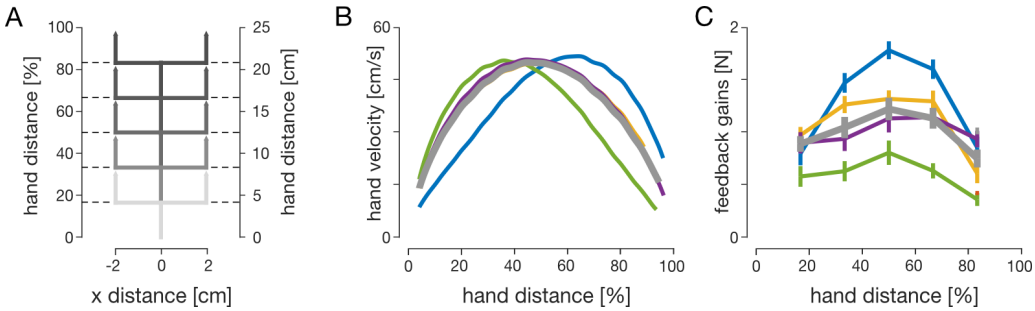


**Figure 2.** Human visuomotor feedback gains are modulated across the five experimental conditions. (**A**) Lateral perturbations of the cursor were applied in all five conditions. Perturbations were introduced as 2 cm cursor jumps perpendicular to the movement direction. The perturbation onset occurred at one of five equally spaced hand locations. (**B**) Mean velocity profiles of the hand in five experimental conditions: matched-cursor early-peak (green), matched-cursor late-peak (blue), matched-hand early-peak (yellow), matched-hand late-peak (purple) and baseline (grey). Participants successfully modulated forward movement kinematics to meet task demands – velocity profiles are skewed for matched-cursor conditions, and are similar to the baseline for matched-hand conditions. (**C**) Mean visuomotor feedback gains (mean lateral force from 180-230 ms after perturbation onset) across all participants to cursor perturbations as a function of the hand distance in the movement. Error bars represent 1 SEM. Significant regulation is observed for matched-cursor early-peak and matched-cursor late-peak conditions (blue and green), but no significant regulation is seen for matched-hand conditions (yellow and purple), relative to the baseline.

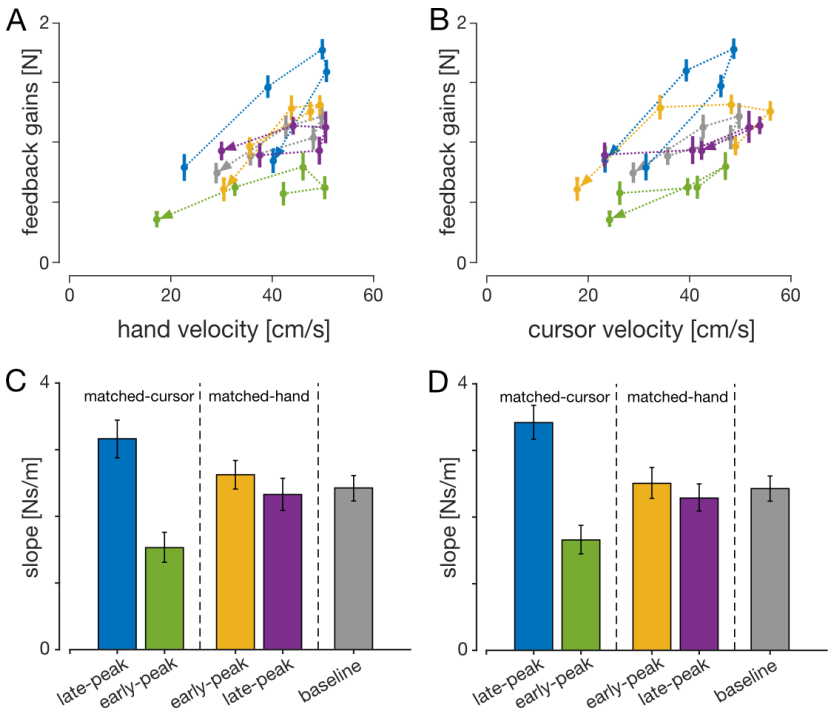


**Figure 3.** Visuomotor feedback gains as a function of (**A**) Hand velocity and (**B**) cursor velocity at the time of perturbation for all experimental conditions. Error bars represent 1 SEM, and the arrowheads represent the order of the perturbation locations. (**C**), (**D**) Regression slopes of feedback gains for each condition as a function of hand and cursor velocities respectively. Error bars represent 95% confidence intervals of the slopes. The slopes for the two matched-cursor conditions were significantly different (based on the confidence intervals) than for the baseline condition.

five different locations in the movement (Figure 2A). This paradigm allowed us to modulate the movement velocity profiles, as well as separate proprioceptive (hand) and visual (cursor) kinematics to examine their individual contribution to visuomotor feedback gains.

Different movement conditions exhibited differences in visuomotor feedback gains (Figure 2). Two-way repeated-measures ANOVA (both frequentist and Bayesian; Materials and Methods) showed significant main effects for both condition ($F_{4,36} = 10.807$, $p < 0.001$, and $BF_{10} = 9.136 \times 10^{12}$), and perturbation location ($F_{4,36} = 33.928$, $p < 0.001$, and $BF_{10} = 6.870 \times 10^{9}$). Post-hoc analysis on movement conditions revealed significant differences between baseline (grey line) and matched-cursor late-peak hand velocity condition (blue line; $t_9 = 4.262$, $p_{bonf} < 0.001$ and $BF_{10} = 247.868$), and between baseline and matched-cursor early-peak hand velocity condition (green line; $t_9 = -8.287$, $p_{bonf} < 0.001$ and $BF_{10} = 1.425 \times 10^{8}$). However, no significant differences were found between the baseline and the two matched hand velocity conditions ($t_9 = 1.342$, $p_{bonf} = 1.0$ and $BF_{10} = 0.357$ for early-peak cursor velocity, yellow; $t_9 = 0.025$, $p_{bonf} = 1.0$ and $BF_{10} = 0.154 \times 10^{8}$ for late-peak cursor velocity, purple). Our results show that different kinematics of the hand movement have a significant effect on visuomotor feedback gain regulation, but that different kinematics of the cursor movement do not.

In order to test whether a simple relationship between movement kinematics and visuomotor feedback gains exists, we mapped visuomotor feedback gains as a linear function of the hand velocity and the cursor velocity. For each experimental condition, we find a different regression slope between the velocity and the feedback gain regardless of whether this is the cursor or the hand velocity (Figure 3AB). Consistent with our previous results, this difference in slopes is significant for conditions where the hand, but not cursor, movement was different (Figure 3CD). Although feedback gains increase with increasing velocity in both cursor and hand coordinates, neither one coordinate modality could predict the changes in the feedback gain.

To successfully complete each trial, participants were required to reach the target. However, the distance to reach the target is affected by the perturbation onset – later perturbation locations lead to larger correction angles (Figure 4A) and thus longer movement distances (Figure 4B). This effect is clearly seen where the extension of movement distance is enhanced for the perturbations closest to the target, with movement distance extended by almost half a centimetre compared to

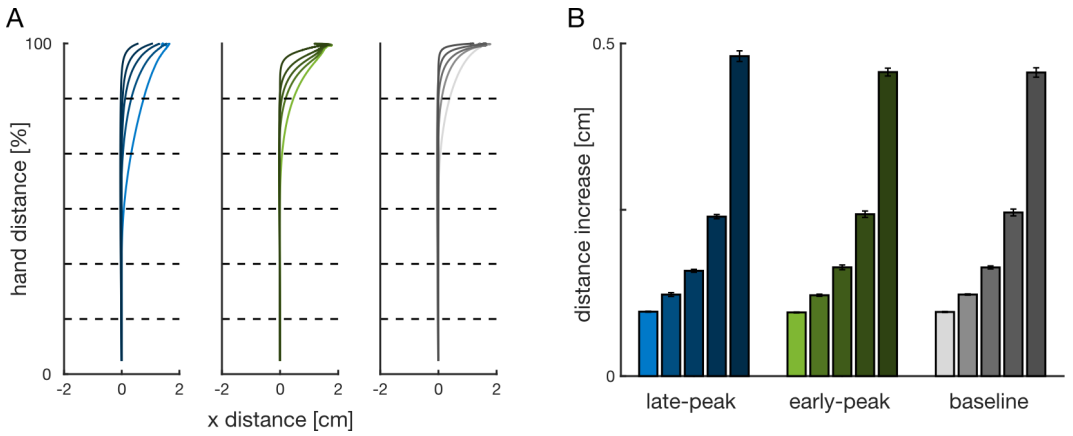


**Figure 4.** (**A**) Mean hand movement trajectories for matched-cursor late-peak (left), matched-cursor early-peak (middle) and baseline (right) conditions recorded in our participants, with perturbation onset at five locations (colour light to dark: 4.2 cm (16.7%), 8.3 cm (33.3%), 12.5 cm (50%), 16.7 cm (66.7%) and 20.8 cm (83.4%) from the start position; dashed lines). Corrections to rightward perturbations were flipped and combined with leftward corrections. (**B**) Distance increase for each perturbation location recorded in our participants. Perturbation locations closest to the target required the largest increases in movement distance. Error bars represent 1SEM.

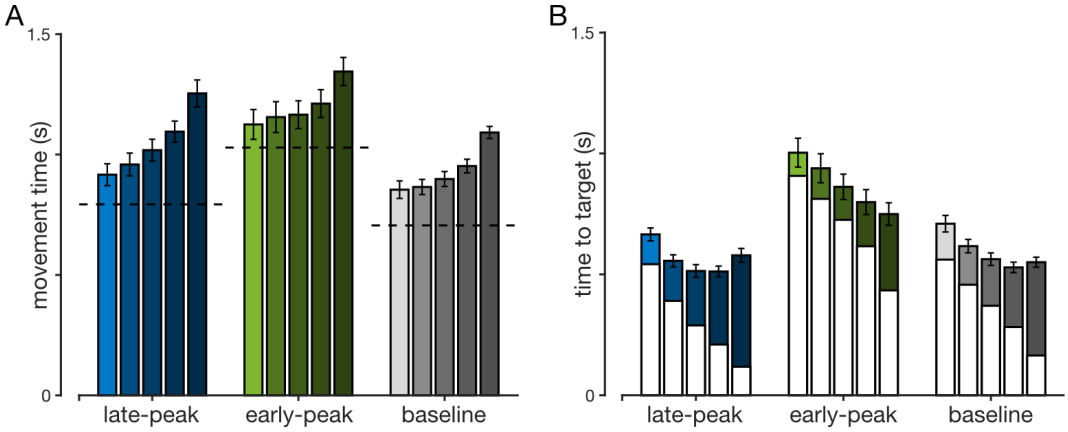


**Figure 5.** (**A**) Movement durations in maintained perturbation trials recorded by our participants in late-peak, early-peak and baseline conditions. Separate bars within the same colour block represent different perturbation onset locations (left to right: 4.2 cm, 8.3 cm, 12.5 cm, 16.7 cm and 20.8 cm from the start position). Error bars represent 1SEM while the horizontal dashed lines represent movement durations in the same movement condition for non-perturbed movements. (**B**) Full bars represent times-to-target in maintained perturbation trials in our participants for late-peak, early-peak and baseline conditions, for perturbations occurring at each of the test locations. White bars represent the time-to-target for each movement if the movement was not perturbed, calculated as the difference between the non-perturbed movement duration and the average time of the perturbation onset. The coloured part of the bars shows the extension in times-to-target due to the perturbation in a non-constrained movement. Error bars represent 1SEM

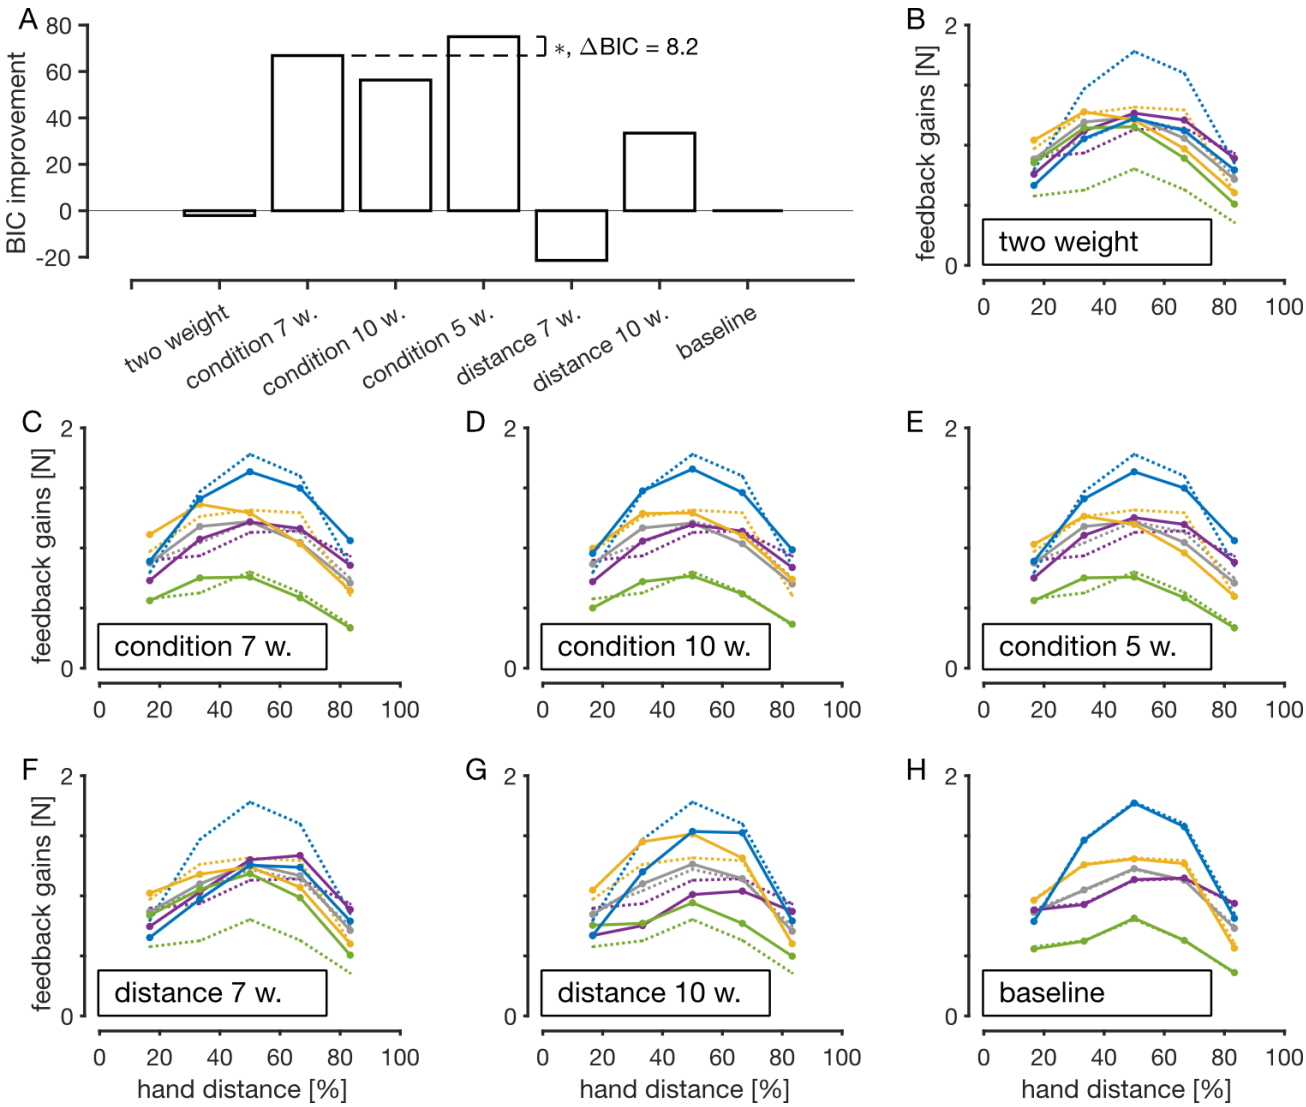


**Figure 6.** Linear model comparison. (**A**) BIC improvement with respect to the baseline model for all models. Condition dependent 5 weight model provides the best fit to the data, with ΔBIC = 8.2 compared to the condition dependent 7 weight model, and a further ΔBIC=10.5 to the condition dependent 10 weight model. ΔBIC > 6 is considered a strong effect, and ΔBIC > 10 a very strong effect (***Raftery and Kass*** (***1995***)). (**B**)-(**H**) Linear models (solid lines) overlaid with mean experimentally recorded visuomotor feedback gains (dashed).

less than one millimetre for the closest perturbations. Any extension of the movement distance requires an appropriate increase in movement duration. Consequently, participants extended their movement time, with longest durations for perturbations close to the target (Figure 5A). This increase in movement duration increases the time to target for these late perturbations (Figure 5B), and now allows sufficient time for the controller to issue any corrective commands.

## Linear models

Can there be a simple relation between movement velocity and feedback gains? In order to examine the relative contributions of visual and haptic modalities, we dissociated the visual and haptic kinematics in our experimental conditions. According to the theory of multisensory integration these two inputs would be weighted and linearly combined (***Ernst and Banks*** (***2002***)), suggesting that the visuomotor feedback gains may depend on a linear combination of the visual and haptic components of movement velocity. To investigate the relative contributions of these separate modalities, we modelled feedback gains as a weighted sum of the hand and cursor velocities (Figure

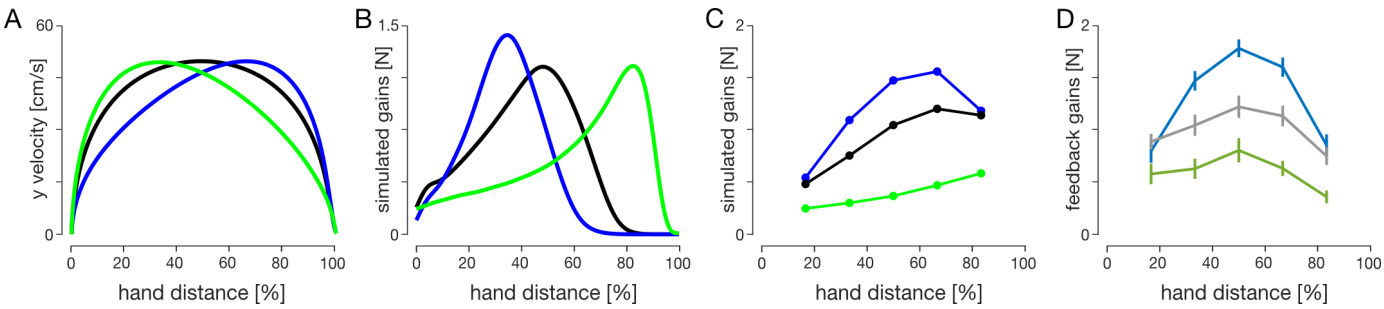


**Figure 7.** Comparison of feedback gains between the two OFC models and experimental data. (**A**) Simulated velocity profiles, and (**B**) Simulated feedback gain profiles of baseline (black), early-peak (green) and late-peak (blue) velocity condition simulations for the classical OFC model. Velocity profiles were obtained by constraining the velocity peak location and magnitude and optimising for movement duration and activation cost function. Simulated feedback gain profiles were obtained by applying virtual target jumps perpendicular to the movement direction during these movements and calculating the force exerted by the controller in the direction of the target jumps. (**C**) Simulated feedback gains obtained via the time-to-target OFC model. Pre-perturbation movements were simulated as if no perturbation would occur, in order to keep the controller naive to an upcoming perturbation. At the perturbation onset the remaining movement duration is adjusted to match the mean time-to-target for a similar perturbation onset in human participants (Figure 5B). The velocity profiles for the time-to-target model match the velocity profiles of the classical model, shown in (**A**). (**D**) Visuomotor feedback gains recorded in human participants.

6), with model weights varying with movement condition or perturbation location (see Materials and Methods), and evaluated our models using the Bayesian Information Criterion improvement over the baseline model (ΔBIC). Our movement condition dependent models outperformed all other models (condition dependent 7 weights: ΔBIC = 66.9; condition dependent 10 weights: ΔBIC = 56.3, condition dependent 5 weights: ΔBIC = 75.0). Importantly, the condition dependent 5 weight model showed a significant BIC improvement (ΔBIC = 8.2) compared to the second best model (condition dependent 7 weights). These results provide additional evidence that visual kinematics have a minor effect on visuomotor feedback gains, compared to hand kinematics. As the best fit model includes a movement condition-dependent scaling, these results argue against the simple relationship between movement velocity and feedback gains.

## Optimal Feedback Control

As optimal control has been suggested to predict the temporal evolution of feedback gains ***Dimitriou et al.*** (***2013***); ***Liu and Todorov*** (***2007***), we built two Optimal Feedback Control (OFC) models: the classical model (***Liu and Todorov*** (***2007***)), and a time-to-target model. For the classical model we implemented an OFC (***Todorov*** (***2005***)) to simulate movements with different velocity profiles, similar to the experiments performed by our participants. We extended this classical model to the time-to-target model, by increasing the movement duration after each perturbation onset according to experimental results (Figure 5). For both models we only simulated different hand kinematics for computational ease and as our participants showed little effect of cursor kinematics on their feedback gains.

For both models we controlled the activation cost R to simulate three conditions in which the location of the peak velocity was shifted to match the experimental hand kinematics (Figure 7A). Specifically we solved for the activation cost R and movement duration N by optimising the log-likelihood of our model's peak velocity location and magnitude using Bayesian Adaptive Direct Search (BADS, ***Acerbi and Ma*** (***2017***)). The optimised movement durations (mean ± SEM) were N = 930 ± 0 ms for the baseline condition, N = 1050 ± 10 ms for the late-peak condition and N = 1130 ± 20 ms for the early-peak condition (10 optimisation runs per condition). In comparison, experimental movement durations were N = 932 ± 30 ms for the baseline condition, N = 1048 ± 47 ms for the late-peak condition and 1201 ± 59 ms for the early-peak condition, matching well with the OFC predictions. Overall this shows that specific constraints on the magnitude and location of peak velocity that we imposed on our participants resulted in a modulation of reaching times that matched OFC predictions under the same constraints.

For the classical model we estimate simulated feedback gains by shifting the movement target at each timepoint in the movement and measuring the magnitude of the force response in the direction of this shift. The simulated feedback gain profiles follow the same general shape as in human participants – response strength increases from the beginning of the movement and then falls off at the end (Figure 7B). However, the overall profile of these simulated feedback gains is very different for each of the kinematic conditions. For the early velocity peak condition the simulated feedback gain peaks towards the end of the movement (green line), whereas for the late peak velocity condition the simulated feedback gain profile peaks early in the movement (blue line). These simulated feedback gains do not appropriately capture the modulation of visuomotor feedback gains in our experimental results. While the simulated feedback gains are qualitatively similar to the experimental results within each condition, overall this model cannot appropriately capture the modulation of visuomotor feedback gains across the conditions.

For the time-to-target OFC model, we extend the classical model to account for the different movement durations for each perturbation location (and movement condition) that is seen in the experimental results. After a perturbation, the remaining time-to-target was adjusted to match the experimentally recorded times-to-target for this specific movement, while before the perturbation both the classical model and the time-to-target model were identical. After adjusting for the individual durations of each perturbation condition we are now able to qualitatively replicate the general regulation of feedback gain profiles for different kinematics using OFC (Figure 7). In the late-velocity peak condition we predict a general increase in the feedback gains throughout the movement compared to the baseline condition, whereas in the early velocity peak condition we predict a general decrease in these feedback gains compared to the baseline condition. Thus we show that within the OFC the time-to-target is critical for the regulation of feedback gains, and when we take this into account we are able to replicate the feedback gain modulation of our participants.

While in our experiment, we manipulated the time-to-target through skewing the velocity profiles, time-to-target is naturally modified through changing the peak velocity. Therefore, we can further analyse the effect of the time-to-target by calculating the feedback gains for movements with different peak velocities (Figure 8A). The simulated feedback gains vary widely across peak velocities, with a shift of peak feedback gains towards the earlier locations for faster movements (Figure 8B). However, when these distinct simulated feedback gain profiles are re-mapped as a function of time-to-target, the simulated feedback gains follow a consistent, albeit non-monotonic, relationship (Figure 8C). This relationship is also consistent over a range of peak velocities across all three kinematic conditions and is well described by a combination of a square-hyperbolic and logistic function (Figure 8D). The squared-hyperbolic arises from the physics of the system: the lateral force necessary to bring a point mass to a target is proportional to $1/t^2$ (Materials and Methods). The logistic function simply provides a good fit to the data. Overall our models show that the feedback gain profiles under OFC are independent of the peak velocity or movement duration. Instead, our simulations suggest that time-to-target is a key variable in regulating visuomotor feedback gains.

## Discussion

Here we examined how movement kinematics regulate visuomotor feedback gains. Participants extended their movement duration after perturbations to successfully reach the target. In addition visuomotor feedback gains were modulated when the hand followed different kinematics, but not when the cursor followed different kinematics. In order to better understand this modulation we built two normative models using OFC. While the classical model failed to predict the experimental visuomotor feedback gains, the time-to-target model qualitatively replicated the visuomotor feedback gain profile of our participants. Overall, optimal feedback control models suggested that feedback gains for each perturbation location depended on the time-to-target rather than distance or velocity. Simulated feedback gains under all movements followed the same profile with respect to time-to-target, suggesting a critical role in the regulation of visuomotor feedback gains.

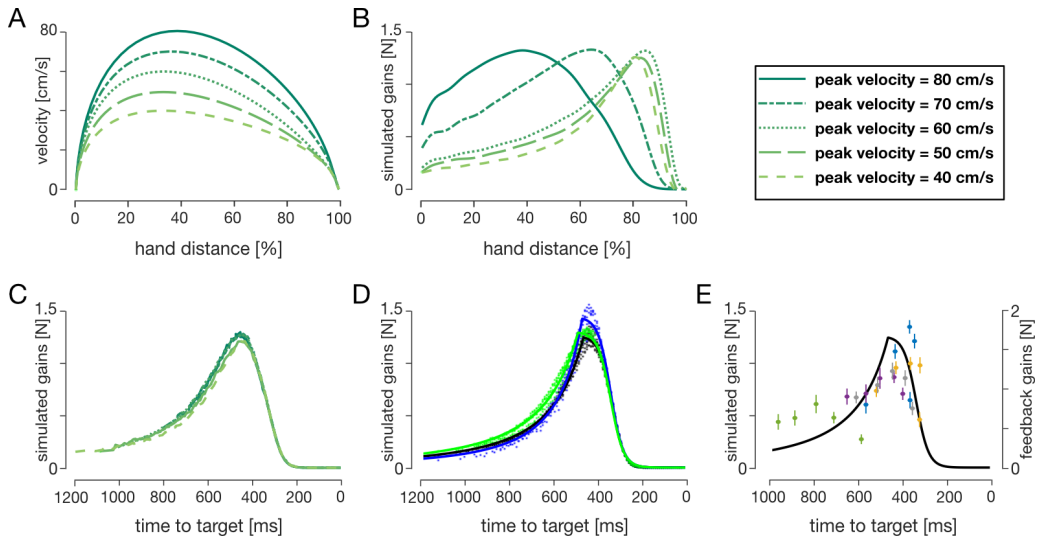


**Figure 8.** OFC simulations of (**A**) velocity profiles and (**B**) simulated feedback gain profiles for different desired peak velocities (in order from light to dark line colours: 40 cm/s, 50 cm/s, 60 cm/s, 70 cm/s, 80 cm/s). (**C**) simulated feedback gains of (**B**) re-mapped as a function of time-to-target at the time of target perturbation. (**D**) Simulated feedback gains vs time-to-target for the three kinematic conditions over the five peak velocities simulated by OFC (coloured dots). Solid lines represent the tuning curves (Equation 8) fit to the data. Both the tuning curves and the simulated feedback gain profiles are similar across a variety of different kinematics when expressed as a function of time-to-target. (**E**) Experimental visuomotor feedback gains for all five experimental conditions (scatter plot) overlaid with the OFC model for baseline condition, as a function of time-to-target.

Experimentally, our participants exhibited a temporal evolution of visuomotor feedback gains for each condition, confirming the findings of ***Dimitriou et al.*** (***2013***). In addition, we also showed the regulation of visuomotor feedback gains across conditions, allowing us to investigate the underlying mechanism of this temporal evolution. Specifically, our experimental results demonstrated strong regulation of visuomotor feedback gain profiles with different hand kinematics, but not with different cursor kinematics (Figure 2C). Compared to the baseline condition, in the matched-cursor early-peak velocity condition participants produced longer times-to-target at each perturbation location (Figure 5B), resulting in lower feedback gains based on the relationship between time-to-target and visuomotor feedback gains (Figure 8E). The opposite is true for the matched-cursor late-peak velocity condition. As the two matched-hand conditions produced similar times-to-target as the baseline due to similar hand kinematics, we did not observe a different regulation in feedback gains. Therefore, the condition dependent visuomotor feedback gain modulation exhibited by our participants meshes nicely with a control policy whereby the time-to-target regulates the feedback gains.

We examined whether multiple linear regression models can describe visuomotor feedback gains as a function of movement kinematics. We found that within the framework of linear models, our data is best described by a condition dependent model with consistent multisensory integration over the movement. Moreover, this model was sensitive to changes in the hand kinematics, but not visual kinematics. In hindsight, this result is also consistent with our time-to-target model: visual kinematics have no effect on visuomotor feedback gains because the time-to-target is only influenced by the change in hand, but not visual, kinematics. Furthermore, the relationship of time-to-target and visuomotor feedback gains also provides an explanation to the up-and-down regulation of visuomotor feedback gain profiles under different hand kinematics. Thus, overall results of linear modeling are consistent with the results provided by the optimal feedback control.

It has long been suggested that we select movements that minimize the noise or endpoint variability (***Harris and Wolpert*** (***1998***)). Within the framework of optimal control, this idea has been expanded to the corrective movements – that is optimality in reaching movements is achieved in

part by minimizing the noise during any corrective response (***Todorov and Jordan*** (***2002***)). As motor noise scales proportionally to muscle activation (***Jones et al.*** (***2002***); ***Hamilton et al.*** (***2004***)), one way of minimising such noise is reducing the peak levels of muscle activation during the correction. Mathematically, the optimal solution to correct any perturbation approximates a constant activation, resulting in a constant force for the whole duration between perturbation onset and target interception. Such a solution assumes that the brain is capable of estimating the remaining duration of the movement and that the force follows the squared-hyperbolic relationship to this duration (Eqn. 9). Our results show that time-to-target strongly modulates visuomotor feedback gains, consistent with the idea that human participants aim to behave optimally. More specifically, we suggest that, among different optimality variables, the temporal evolution of visuomotor feedback gains serves to reduce effects of system noise.

Optimal feedback control predicts a time beyond which feedback gains are suppressed. Beyond this critical time, a logistic function well describes the relation between time-to-target and feedback gains, with gains reducing as the time-to-target decreases. OFC suggests that this logistic function describes the controller's unwillingness to react to the perturbation. That is, increasing accuracy of the movement at this point costs more than missing the target. However, as participants were required to intercept the target, they extended their movement durations past this point of cost equilibrium, effectively producing a movement of a longer duration. Our time-to-target OFC model accounts for such extensions, resulting in the controller moving out of the logistic function section (into the squared-hyperbolic function) and producing visuomotor feedback gain profiles similar to those in human data.

A limitation of our model is that it takes time-to-target as an input in order to generate feedback gain predictions, rather than obtain the time-to-target as a model output. As a result, our models do not describe exactly how the change in movement geometry after the perturbation influences this time-to-target, which in turn regulates the visuomotor feedback gains. However, utility of movement has recently been used within optimal control to characterise reaching movements (***Rigoux and Guigon*** (***2012***); ***Shadmehr et al.*** (***2016***)) in which optimal movement time falls out automatically from a trade-off between reward and effort. With respect to our models, this adds additional complexities to capturing the different movement conditions. Future approaches could attempt to model these results within the utility of movement framework.

Rapid feedback responses scale with the urgency to correct for mechanical perturbations (***Crevecoeur et al.*** (***2013***)). Here we have shown that visuomotor feedback gains also follow a similar regulation, suggesting that these two systems share the same underlying control policy. Our work further extends this finding of Crevecoeur by not just showing that urgency affects feedback gains, but explaining the manner in which these gains are regulated with respect to urgency. That is, here we have shown that for visual perturbations the feedback gains scale with a squared-hyperbolic of the time-to-target, which is a direct measure of urgency. Moreover, the feedback gains were rapidly adjusted due to the change in urgency as the task changed. Specifically, when the cursor jumps close to the target, the expected time-to-target is prolonged, and therefore the optimal visuomotor feedback gain needs to be adjusted appropriately to this increase in time. Our results show that participants produce a visuomotor response consistent with the actual, post-perturbation, time-to-target, as opposed to the expected time-to-target prior to the perturbation. Therefore, our results not only suggest that similar computations might occur for both stretch and visuomotor feedback response regulation, but also that this regulation originates from task-related optimal feedback control.

Our work has shown that simulated feedback gains from OFC exhibit the same underlying pattern as a function of time-to-target over a wide range of movement kinematics, matching well the feedback gains of our human participants (Figure 7). These results suggest that, for known dynamics, we do not need to recalculate the feedback gains prior to each movement, but instead can access the appropriate pattern as a function of the estimated time-to-target in each movement. Therefore gain computation in reaching movements may not be a computationally

expensive process, but instead could be part of an evolutionary control strategy that allows for rapid estimation of the appropriate feedback gains. Moreover, the fact that both stretch reflex and visuomotor feedback systems exhibit similar control policies despite different sensory inputs, perhaps only sharing the final output pathway, suggests that this simple feedback pathway may be an evolutionary old system. Indeed, several studies have suggested that visuomotor feedback is controlled via a pathway through the colliculus (***Reynolds and Day*** (***2012***); ***Gu et al.*** (***2018***); ***Corneil et al.*** (***2004***)). Such a system would then only need to be adapted as the dynamics change, allowing for fine tuning of the feedback gains according to changes in the environment (***Franklin et al.*** (***2017***)).

Our results have shown the connection between the visuomotor feedback gain regulation and the time left to complete the movement. Specifically, in our human participants we recorded the increase in the time-to-target after the perturbation onset, which consequently increased the movement durations (Figure 5). This increase was also longer for later perturbations, consistent with previous studies (***Liu and Todorov*** (***2007***)). According to our normative OFC model, the time-to-target alone is enough to successfully regulate visuomotor feedback gains as observed in humans. This result was independent of the kinematics of the movement or the onset times of the perturbations. This suggests that there is no recalculation of a control scheme for the rest of the movement after the perturbation, but rather a shift to a different state within the same control scheme. Such findings are consistent with the idea that visuomotor feedback gains are pre-computed before the movement, allowing for faster than voluntary reaction times (***Franklin*** (***2016***)). Moreover, through our results, we gain a deeper insight into how optimal feedback control governs these feedback gains – through a straightforward relationship to the estimated time-to-target, based on physics.

## Materials and Methods

### Participants

Eleven right-handed (***Oldfield*** (***1971***)) human participants (5 females; 27.3 $\pm$ 4.5 years of age) with no known neurological diseases took part in the experiment. All participants provided written informed consent before participating. All participants except one were naïve to the purpose of the study. Each participant took part in five separate experimental sessions, each of which took approximately 3 hours. One participant was removed from analysis as their kinematic profiles under the five experimental sessions overlapped. 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 to a target while grasping the handle of a robotic manipulandum with their right hand. Participants were seated in an adjustable chair and restrained using a four-point harness. The right arm of participants was supported on an air sled while grasping the handle of a planar robotic interface (vBOT, ***Howard et al.*** (***2009***)). 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 1kHz. Visual feedback was provided in the plane of the hand via a computer monitor and a mirror system, such that this system prevented direct visual feedback of the hand and arm. The exact onset time of any visual stimulus presented to the participant was determined from the graphics card refresh signal.

Participants initiated each trial by moving the cursor (yellow circle of 1.0 cm diameter) into the start position (grey circle of 1.6 cm diameter) located approximately 25 cm in front of the participant, centred with their body. This start position turned from grey to white once the cursor was within the start position. Once the hand was within the start position for a random delay drawn from a truncated exponential distribution (1.0-2.0 s, mean 1.43 s), a go cue (short beep) was provided signalling participants to initiate a straight reaching movement to the target (red circle of 1.2 cm diameter, located 25.0 cm directly in front of the start position). If participants failed to initiate the

movement within 1000 ms the trial was aborted and restarted. Once the cursor was within 0.6 cm of the centre of the target, participants were notified by the target changing colour to white. The movement was considered complete when the participants maintained the cursor within this 0.6 cm region for 600 ms. After each trial, the participant's hand was passively returned by the robot to the start position while visual feedback regarding the success of the previous trial was provided (Figure 9). Movements were self-paced, and short breaks were enforced after every 100 trials.

## Experimental paradigm

Participants performed the experiment under five different conditions, each performed in a separate session. In the baseline condition the cursor matched the forward movement of the hand, with a peak velocity in the middle of the movement. In the other four conditions, the cursor location was scaled relative to the hand location in the forward direction, such that the cursor and the hand location matched only at the start and end of the movements (Figure 1). In two of the conditions (matched-hand velocity), the hand velocity matched the baseline condition throughout the movement (with the peak in the middle of the movement) but the cursor velocity peaked either earlier (33% of movement distance) or later (66% of movement distance). In the other two conditions (matched-cursor velocity), the cursor velocity was matched to the baseline condition throughout the movement (with the peak in the middle of the movement) but the hand velocity peaked either earlier (33% of movement distance) or later (66% of movement distance). The difference between the cursor velocity and the hand velocity was produced through a linear scaling of the cursor velocity as a function of the forward position (Figure 1A). Specifically, for the two conditions where the position of the peak cursor velocity is earlier than the position of the peak hand velocity (Figure 1 top), this scaling was implemented as:

$$\frac{v_c}{v_h} = -0.012d + 1.6, \tag{1}$$

where $v_c$ and $v_h$ are cursor and hand velocities respectively, and $d$ is the distance along the movement direction in %. The cursor velocity was therefore manipulated by a linear scaling function such that its velocity is 160% of the hand velocity at the beginning of the movement, linearly decreasing to 40% at the target location (Figure 1 top). For the two conditions where the position of the peak cursor velocity is later than the position of the peak hand velocity (Figure 1 bottom), this scaling was implemented as:

$$\frac{v_c}{v_h} = 0.012d + 0.4 \tag{2}$$

such that the velocity gain function linearly increased from 40% hand velocity at the start of the movement to 160% at the end of the movement (Figure 1, bottom). Desired velocity profiles of both the hand and the cursor are shown in Figure 1B for each condition.

## Feedback regarding movement kinematics

In all conditions, one of the velocity modalities (cursor or hand) was required to be similar to the baseline velocity profile. Feedback was always provided about this specific velocity modality. Ideal trials were defined as trials in which this peak velocity was between 42 cm/s and 58 cm/s with the peak location between 45% and 55% of the movement distance with no target overshoot. After each trial, visual feedback about the peak velocity and the location at which this peak occurred was provided to the participants graphically (Fig 9). The peak velocity was indicated on the right hand side of the screen with the length of a bar and the velocity target. This bar changed colour from red to green if the velocity was within the ideal range. The location of the peak velocity was indicated as a horizontal line between home and target positions at the exact location it was achieved, along with the ideal range. This line was green when the location of the peak velocity was within the ideal range, and red otherwise. Overshooting the target was defined as the position of the cursor exceeding the centre of the target in the y-coordinate by more than 0.9 cm.

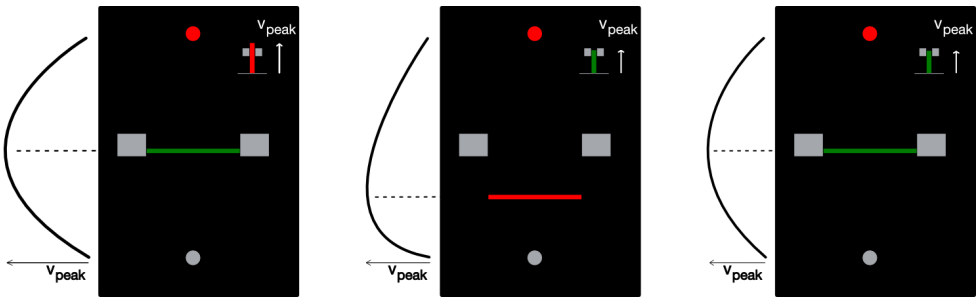


**Figure 9.** Examples of feedback presented to the participants. Feedback regarding the peak velocity and the timing of the peak velocity was provided after each trial. Large grey blocks indicate the velocity peak location target, while the bar chart at the top-right corner indicates peak y-velocity magnitude. Feedback was provided on the modality (cursor or hand) that matched the baseline. Left: velocity peak location is within the target, but the movement was too fast (unsuccessful trial); middle: velocity peak location is too early, but the movement speed is within the target (unsuccessful trial); right: successful trial.

## Probe trials

During each session, probe trials were used to measure the visuomotor feedback gains – the corrective motor responses to a change in the visual feedback of hand position. To measure the feedback gains, visual perturbations were initiated laterally ($\pm 2.0$ cm) at five different hand distances (4.2, 8.3, 12.5, 16.7, and 20.8 cm) from the start (Figure 2A). In addition a zero amplitude perturbation (cursor matched to the lateral position of the hand) was included, resulting in eleven different probe trials. On these trials the visual perturbations lasted 250 ms, after which the cursor was returned to the lateral location of the hand. The lateral hand position was constrained in a simulated mechanical channel throughout the movement, thereby requiring no correction to reach the target. The simulated mechanical channel was implemented with a stiffness of 4000 N/m and damping of 2 Ns/m acting perpendicularly to the line connecting the start position and the target (***Scheidt et al.*** (***2000***); ***Milner and Franklin*** (***2005***)), allowing measurement of any lateral forces in response to a visual perturbation.

In previous experiments, feedback responses gradually decreased during the course of the experiment (***Franklin and Wolpert*** (***2008***); ***Franklin et al.*** (***2012***)). However, it has been shown that including perturbation trials where the perturbations were maintained until the end of the movement, and where participants had to actively correct for the perturbation to reach the target, prevents this decrease in the feedback responses (***Franklin et al.*** (***2016***)). Therefore half of the trials contained the same range of perturbations as the probe trials but where these perturbations were maintained throughout the rest of the trial and participants had to correct for this perturbation. These maintained perturbations have now been used in several studies ***Franklin et al.*** (***2016***, 2017); ***de Brouwer et al.*** (***2017***).

## Session design

Prior to each session, participants performed 100 to 300 training trials in order to learn the specific velocity profiles of the reaching movements. All training trials contained no visual perturbations and were performed in the null force field. The training trials were stopped early once participants achieved an accuracy of 75% over the last 20 trials, and were not used for the analysis.

Each session consisted of 40 blocks, where each block consisted of 22 trials performed in a randomized order. Eleven of these 22 trials were probe trials (5 perturbation locations $\times$ 2 perturbation directions + zero perturbation condition) performed in the mechanical channel. The other eleven trials consisted of the same perturbations but maintained throughout the trial and performed in the null field. Therefore in each of the five sessions participants performed a total 880 trials (440 probe trials). The order of the five different conditions (sessions) was pseudo-randomized and counterbalanced across participants.

## Data analysis

Data was analyzed in MATLAB R2017b and JASP 0.8.2. Force and kinematic time series were low-pass filtered with a tenth-order zero-phase-lag Butterworth filter (40 Hz cutoff). The cursor velocity was calculated by multiplying the hand velocity by the appropriate scaling function. The visuomotor feedback response was measured for each perturbation location as the difference between the force responses to the leftward and rightward perturbations within a block. To measure the visuomotor feedback gain (change in force to a fixed-size visual perturbation) the force difference was averaged over a time window of 180-230 ms, a commonly used time interval for the involuntary visuomotor feedback response (***Franklin and Wolpert*** (***2008***); ***Dimitriou et al.*** (***2013***); ***Franklin et al.*** (***2012***, 2016)). In order to compare any differences across the conditions a two-way repeated-measures ANOVA was performed with main effects of condition (5 levels) and perturbation location (5 levels). As a secondary method to frequentist analysis we also used the Bayesian factor analysis (***Raftery and Kass*** (***1995***)) to verify our statistical results. Bayesian factor analysis is a method that in addition to the conventional hypothesis testing (evaluating evidence in favour of the alternative hypothesis) allows us to evaluate evidence in favour of the null hypothesis, therefore distinguishing between the rejection of the alternative hypothesis and not enough evidence to accept the alternative hypothesis.

Although we used the time window of 180-230 ms to estimate visuomotor feedback gains, we also verified whether the onset of the visuomotor feedback response in our data is consistent with previously reported values. To estimate this onset time, we first estimated individual onset times for each participant at each perturbation location and movement condition. To do so, we used the Receiver Operator Characteristic (ROC) to estimate where the force reaction to leftwards cursor perturbations deviated from the reaction to rightwards cursor perturbations (***Pruszynski et al.*** (***2008***)). For each type of trials we built the ROC curve for the two signals at 1 ms intervals, starting from 50 ms before the perturbation, and calculated the area under this curve (aROC) for each of these points until the aROC exceeded 0.75 for ten consecutive milliseconds. In order to find where the force traces start deviating from each other we then fit a function of the form $max(0.5, k \times (t - \tau)$ to the aROC curve. The time point where the linear component of this function first overtakes the constant component was taken as the threshold value. Overall, the mean onset times across all conditions and perturbation locations were $138 \pm 7$ ms (mean + SD), with onset times consistent among movement conditions ($F_{4,36} = 1.410$, $p = 0.25$, and $BF_{10} = 0.105$), perturbation locations ($F_{4,36} = 1.582$, $p = 0.20$, $BF_{10} = 0.252$), and their interactions ($F_{16,144} = 1.350$, $p = 0.176$, and $BF_{10} = 0.005$)

## Modelling

### Linear modelling

Here we tested whether the pattern of feedback gains in all five experimental conditions could be explained by any linear function of hand and/or cursor velocity. Specifically, we performed multiple linear regression models of the general equation:

$$G = w_{h,mc,pl} v_h + w_{c,mc,pl} v_c, \tag{3}$$

where $w_{h,mc,pl}$ and $w_{c,mc,pl}$ are regression weights for hand velocity ($v_h$) and cursor velocity ($v_c$) respectively that depend on movement condition ($mc$) and perturbation location ($pl$). These regression weights were constrained to be non-negative. These models were fit to the scaled mean of each participant's data (250 data points: 10 participants $\times$ 5 movement conditions $\times$ 5 perturbation locations). Each participant's data was scaled so that the mean responses across all trials were equal to one, to ensure equal contributions from each subject.

We fit seven models with the number of parameters ranging from 2 to 25 (as we fit 25 data points per participant). For each model, we determined the best fit model parameters using a non-linear least-squares solver (`lsqnonlin`, MATLAB 2017b). We then evaluated model fits using

Bayesian Information Criterion (BIC), which allows us to compare models and account for potential overfitting by penalising additional parameters. The seven models for comparison were:

*Two weight model:* $G = w_h v_h + w_c v_c$, where $G$ is the feedback gain, $v_h$ and $v_c$ are hand and cursor velocities respectively at the time of the perturbation, and $w_h$ and $w_c$ are the corresponding model weights. The two weight model has one weight for the hand velocity and one for the cursor velocity. These weights are the same for all five movement conditions and perturbation locations. This model therefore assumes there is a simple relation between movement velocity and visuomotor feedback gains.

*Condition dependent 7 weight model:* $G = k_{mc}(w_h v_h + w_c v_c)$. This is an extension of the two weight model, including an additional multiplier for each movement condition $k_{mc}$. This coefficient allows for the modelling of condition dependent regulation of feedback gains. Here the ratio between $w_h$ and $w_c$ is consistent for all movement conditions. This model therefore assumes a consistent integration of visual and proprioceptive information, but allows for condition dependent regulation.

*Condition dependent 10 weight model:* $G = w_{h,mc} v_h + w_{c,mc} v_c$. For each of the five movement conditions, there are individual weights for the hand and cursor velocities. This model is an extension of the condition dependent 7 weight model, with the relative integration of visual and proprioceptive information not fixed across movement conditions. Therefore this allows for variations in the relative contributions of vision and proprioception from condition to condition.

*Condition dependent 5 weight model:* $G = k_{mc\ hand}(w_h v_h + w_c v_c)$. This is a more constrained version of the condition dependent 7 weight model, with condition-dependent scaling varying only for conditions with different hand kinematics. Therefore, $k_{mc\ hand}$ in this model is identical for the baseline condition and the two matched-hand movement conditions. This model assumes that visual kinematics have no effect on the visuomotor feedback gains.

*Distance dependent 7 weight model):* $G = k_{pl}(w_h v_h + w_c v_c)$. This is an extension of the two weight model, including an additional multiplier for each perturbation location $k_{pl}$. Here the ratio between $w_h$ and $w_c$ is consistent for all perturbation locations. This model treats all movement conditions the same, however it allows for modulation of visuomotor feedback gain by perturbation location. This is equivalent to feedback modulation with movement distance.

*Distance dependent 10 weight model:* $G = w_{h,pl} v_h + w_{c,pl} v_c$. For each of the five perturbation locations, there are individual weights for the hand and cursor velocities. This model is an extension of the distance dependent 7 weight model, with the relative integration of visual and proprioceptive information not fixed across perturbation locations. Therefore, this allows for variations in the relative contributions of vision and proprioception along the movement.

*Baseline model:* $G = w_{mc,pl}(v_h + v_c)$. This model has 25 parameters that are used to fit the 25 data points. This model has a separate weight for each movement condition and perturbation location, and therefore should be able to fit the mean participant data.

## Optimal feedback control

In addition to our linear models we implemented two different Optimal Feedback Control (OFC) models: the classical model (***Liu and Todorov*** (***2007***)) and the time-to-target model. In both models we modelled the hand as a point mass of $m = 1.1$ kg and the intrinsic muscle damping as a viscosity $b = 7$ Ns/m. This point mass was controlled in a horizontal plane by two orthogonal force actuators to simulate muscles. These actuators were controlled by the control signal $u_t$ via a first order low-pass filter with a time constant $\tau = 0.05$ s. The state-space representation of the dynamic system used to simulate the reaching movements can be expressed as

$$x_{t+1} = Ax_t + B(I + C)u_t + \xi_t, \tag{4}$$

where $A$ is a state transition matrix, $B$ is a control matrix, and C is a $2 \times 2$ matrix whose each element is a zero-mean normal distribution representing control-dependent noise. Variables $x_t$ and $u_t$ are state and control at time t respectively. State $x_t$ exists in the Cartesian plane and consists of position

$\mathbf{p}$ (2 dimensions), velocity $\mathbf{v}$ (2), force $\mathbf{f}$ (2) and target position $\mathbf{p}^*$ (2). For our simulation purposes we treat the control-independent noise $\xi_t$ as zero.

The state of the plant is not directly observable, but has to be estimated from noisy sensory information. We model the observer as

$$y_t = Hx_t + D_t, \tag{5}$$

where $H = \text{diag}[1, 1, 1, 1, 1, 1, 0, 0]$ is the observation matrix, and $D_t$ is a diagonal matrix of zero-mean normal distributions representing state-independent observation noise. Therefore, our observer can infer the state information of position, velocity and applied force of the plant, consistent with human participants.

The simulated movements were guided by the LQG controller with state dependent cost Q, activation cost R, a reaching time N, and a time step t = 0.01 s. However, due to the presence of the control-dependent noise, the estimation and control processes are not anymore separable as in the classic LQG theory. In order to obtain optimal control and Kalman gain matrices we utilised the algorithm proposed by ***Todorov and Weiwei Li*** (***2005***) where control and Kalman gain matrices are iteratively updated until convergence.

For both classical and time-to-target model we simulated three different movement kinematics representing three different conditions in our experiment – the baseline and the two matched-cursor conditions. The state dependent cost Q was identical for all three kinematics:

$$Q(t) = \begin{cases} 0, & \text{for } t \neq N \\ (\omega_{\mathbf{p}}(\mathbf{p}(t) - \mathbf{p}^*(t)))^2 + \omega_v ||\mathbf{v}(t)||^2 + \omega_f ||\mathbf{f}(t)||^2, & \text{for } t = N \end{cases} \tag{6}$$

where $\omega_{\mathbf{p}} = [0.5, 1]$, $\omega_v = 0.02$, and $\omega_f = 2$. The activation cost R($t$) = 0.00001 was constant throughout the movement for the baseline condition, but was modulated for the two matched-cursor conditions by multiplying it elementwise by a scaling function:

$$R'(t) = \frac{\exp(p\frac{t+q}{r})}{||\exp(p\frac{t+q}{r})||}, \tag{7}$$

where $p$, $q$ and $r$ are constants. All other parameters were kept constant across the three conditions.

Although LQG is a fixed time horizon problem, we did not pre-define the movement duration N. Instead, we obtained the N, and constants $p$, $q$ and $r$ using Bayesian Adaptive Direct Search (BADS, ***Acerbi and Ma*** (***2017***)) to maximise the log-likelihood of the desired movement kinematics given the location and magnitude of the peak velocity.

The classical and the time-to-target models only differed in the way the perturbations were handled. For the classical model, we simulated perturbation trials at every time step $t_p$ by shifting the target x-coordinate by 2 cm at the time $t_p + 120$ ms. This 120 ms delay was used in order to mimic the visuomotor delay in human participants, and was taken from ***Liu and Todorov*** (***2007***). We then averaged the force response of the controller over the time window $t_p + 130$ - $t_p + 180$ as an estimate of the simulated feedback gains, equivalent of visuomotor feedback gains in our participants. For perturbations occurring at times where the movement is over before the end of this time window, this simulated feedback gain is set to zero.

For the time-to-target model we introduced an extension in the time-to-target after the onset of any perturbation similar to that observed in our participants. Simulated feedback gains were modelled at five locations, matching the perturbation locations in our experiment to obtain the appropriate increase in time-to-target after each perturbation. In order to simulate the response to perturbations we first extracted the perturbation onset times from movement kinematics by performing an unperturbed movement and recording the timepoint $t_p$ at which this movement passed the perturbation onset location. We then simulated the post-perturbation portion of the movement as a new LQG movement with an initial state matching the state at $t_p + 120$ ms of the

unperturbed movement, and movement duration matching the time-to-target recorded in our participants for the particular perturbation. Together this keeps our simulated reaches “naive" to the perturbation prior to its onset and allows the time-to-target of the simulated reaches to match the respective time-to-target of our human participants. Finally, we calculated the simulated feedback gains as described previously, using a time window 10 ms - 60 ms of the post-perturbation movement.

### Time-to-target tuning function

In order to understand the mechanisms that might underlie the consistent relationship between the simulated feedback gains and the time-to-target, we fit a mathematical expression to the simulated feedback gains. We modelled the relationship as the minimum of a squared-hyperbolic function and a logistic function:

$$G(t) = \min\left(\frac{\beta}{\left(t - t_1\right)^2};\ \frac{\alpha}{1 + \exp\left(-\frac{t - t_0}{\tau}\right)}\right) \tag{8}$$

and used BADS to optimise the log-likelihood of this model.

The logistic function was chosen simply as it provided a good fit to the data. The squared-hyperbolic arises from the physics of the system. A point mass ($m$) travelling a distance ($d$) under the influence of force $F$ can be expressed as:

$$d = \frac{Ft^2}{2m} + v_0 t \leftrightarrow F = \frac{2md}{t^2}, \tag{9}$$

where $v_0 = 0$ is the lateral velocity at the start of perturbation correction. Hence the lateral force necessary to bring a point mass to the target is proportional to $1/t^2$.

## Acknowledgments

We thank Matthew Millard, Michael Dimitriou, Sae Franklin and Raz Leib for their comments on an earlier version of this manuscript.

## References

Acerbi, L. and Ma, W. J. (2017). Practical Bayesian Optimization for Model Fitting with Bayesian Adaptive Direct Search. *Advances in Neural Information Processing Systems*, (30):1836–1846.

Corneil, B. D., Olivier, E., and Munoz, D. P. (2004). Visual responses on neck muscles reveal selective gating that prevents express saccades. *Neuron*, 42(5):831–841.

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.

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.

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.

Ernst, M. O. and Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415:429–433.

Franklin, D. W. (2016). Rapid Feedback Responses Arise From Precomputed Gains. *Motor Control*, 20(2):171–176.

Franklin, D. W., Franklin, S., and Wolpert, D. M. (2014). Fractionation of the visuomotor feedback response to directions of movement and perturbation. *Journal of Neurophysiology*, 112(9):2218–2233.

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., 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*, 118(5):2711–2726.

Gu, C., 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.

Guigon, E., Baraduc, P., and Desmurget, M. (2007). Computational Motor Control: Redundancy and Invariance. *Journal of Neurophysiology*, 97(1):331–347.

Guigon, E., Baraduc, P., and Desmurget, M. (2008). Optimality, stochasticity, and variability in motor behavior. *Journal of Computational Neuroscience*, 24(1):57–68.

Hamilton, A. F., Jones, K. E., and Wolpert, D. M. (2004). The scaling of motor noise with muscle strength and motor unit number in humans. *Experimental Brain Research*, 157(4):417–430.

Harris, C. M. and Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394:780–784.

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.

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

Jones, K. E., Hamilton, A. F., and Wolpert, D. M. (2002). Sources of signal-dependent noise during isometric force production. *Journal of Neurophysiology*, 88(3):1533–1544.

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.

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

Milner, T. E. and Franklin, D. W. (2005). Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *Journal of Physiology*, 567(2):651–664.

Nagengast, A. J., Braun, D. A., and Wolpert, D. M. (2009). Optimal control predicts human performance on objects with internal degrees of freedom. *PLoS Computational Biology*, 5(6).

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.

Nashed, J. Y., Crevecoeur, F., and Scott, S. H. (2014). Rapid Online Selection between Multiple Motor Plans. *Journal of Neuroscience*, 34(5):1769–1780.

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

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., Kurtzer, I., and Scott, S. H. (2008). Rapid Motor Responses Are Appropriately Tuned to the Metrics of a Visuospatial Task. *Journal of Neurophysiology*, 100(1):224–238.

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., 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).

Sarlegna, F., Blouin, J., Bresciani, J. P., Bourdin, C., Vercher, J. L., and Gauthier, G. M. (2003). Target and hand position information in the online control of goal-directed arm movements. *Experimental Brain Research*, 151(4):524–535.

Saunders, J. A. (2004). Visual Feedback Control of Hand Movements. *Journal of Neuroscience*, 24(13):3223–3234.

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., Huang, H. J., and Ahmed, A. A. (2016). A Representation of Effort in Decision-Making and Motor Control. *Current Biology*, 26(14):1929–1934.

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.

Todorov, E. and Weiwei Li (2005). A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. *Proceedings of the 2005, American Control Conference, 2005.*, pages 300–306.

Yeo, S.-H., Franklin, D. W., and Wolpert, D. M. (2016). When Optimal Feedback Control Is Not Enough: Feedforward Strategies Are Required for Optimal Control with Active Sensing. *PLOS Computational Biology*, 12(12):e1005190.