

Published in final edited form as:

Nat Neurosci. 2016 February ; 19(2): 320–326. doi:10.1038/nn.4214.

Parallel Specification of Competing Sensorimotor Control Policies for Alternative Action Options

Jason P. Gallivan^{*,1,2,3}, Lindsey Logan⁴, Daniel M. Wolpert⁵, and J. Randall Flanagan^{*,1,3}

¹Departments of Psychology, Queen's University, Kingston, Ontario, Canada

²Biomedical and Molecular Sciences, Queen's University, Kingston, Ontario, Canada

³Centre for Neuroscience Studies, Queen's University, Kingston, Ontario, Canada

⁴Life Sciences Program, Queen's University, Kingston, Ontario, Canada

⁵Department of Engineering, University of Cambridge, Cambridge, U.K.

Abstract

Recent theory proposes that the brain, when confronted with several action possibilities, prepares multiple competing movements in advance of deciding between them. To date, psychophysical supporting evidence for this idea comes from rapid reaching tasks in which individuals, under cases of target uncertainty, launch initial movements in between competing targets, consistent with multiple prepared reaches being executed simultaneously. Reach planning, however, involves far more than specifying movement direction(s); it requires the specification of sensorimotor control policies that set feedback gains shaping how the motor system responds to errors induced by noise or perturbations. Here we show that the feedback gains for reaches toward multiple potential targets correspond to an average of the gains on trials in which the targets are known in advance. These findings provide evidence that the brain, when presented with multiple action options, computes multiple competing sensorimotor control policies in parallel prior to instituting one of them.

Keywords

Action; decision-making; motor planning; parallel encoding; reflexes; reaching; feedback gains; online control

*Correspondence should be addressed to: Jason P. Gallivan, Depts. of Psychology and Biomedical and Molecular Sciences, Centre for Neuroscience Studies, Queen's University, Kingston, Ontario, Canada K7L 3N6, gallivant@queensu.ca; J. Randall Flanagan, Dept. of Psychology, Centre for Neuroscience Studies, Queen's University, Kingston, Ontario, Canada K7L 3N6, flanagan@queensu.ca.

A supplementary methods checklist is available.

Authors Contributions

J.P.G., L.L., D.M.W., and J.R.F. designed the experiment. L.L. performed research. J.P.G., D.M.W., and J.R.F. analyzed and interpreted data. J.P.G., L.L., D.M.W., and J.R.F. wrote the paper.

Competing Financial Interests

The authors declare no competing financial interests.

Introduction

In natural environments we are often faced with a multitude of action possibilities. In such situations it has been suggested that the brain prepares, in parallel, multiple potential movements prior to deciding between and implementing one of them 1,2. This compelling idea—referred to as the affordance competition hypothesis 3—receives empirical support from neurophysiological investigations in sensorimotor areas of the brain showing the parallel encoding of multiple potential reach, grasp and saccade targets prior to an animal deciding between, and then making a corresponding movement towards, one of these targets 1,4,5. Psychophysical support for this framework has come from spatial averaging behaviour, whereby individuals, when required to launch a reaching or saccadic eye movement prior to knowing which of several potential targets will be selected (i.e., ‘go-before-you-know’), are found to aim their initial movements towards the midpoint of the target distribution, consistent with an averaging of the multiple competing movement directions 6–8. Recent evidence has further indicated that, in addition to movement direction, the reach paths 9 and hand orientations 10 required to contact the multiple, competing targets may also be averaged. Taken together, the available neural and behavioural evidence suggests that the brain, when presented with multiple available actions, directly maps, in parallel, the competing options into associated motor variables.

Although basic movement-related variables like direction, path and orientation can provide a useful characterization of some aspects of motor planning, it constitutes only a fraction of the actual parameterizations required for skilled behaviour. A major component of any fully elaborated movement plan is the specification of a sensorimotor control policy, which refers to the setting of feedback gains that determine how the movement evolves as a function of the state of the motor system and also how the motor system handles errors that can occur due to noise or external perturbations 11–15. Indeed, according to several contemporary models of motor control, the initial movement trajectory is not explicitly planned but, rather, emerges from the specification of a movement goal and feedback gains 16. Feedback gains have been shown to be flexibly adapted in accordance with the features of the task and environment 17–21 and, according to the theory of optimal feedback control (OFC), are governed by a policy of minimum intervention, whereby the sensorimotor system responds strongly to errors that endanger the goal(s) of the task but less vigorously to those that do not 13,14,22. Consistent with the predictions of the OFC framework, feedback gains, in response to either mechanical or visual perturbations, are heightened when reaching towards a narrower compared to wider target 23,24 and can be rapidly updated (~100 ms), within a single movement, based on a change in task parameters (e.g., when the target jumps location, see 25).

Given the clear importance of sensorimotor control policies in movement planning and control, the affordance competition hypothesis might predict that feedback gains should be specified for each of the alternative movement plans formed for each potential reach target. If so, then just as the initial movement in go-before-you-know tasks may reflect an average of the reach directions to the potential targets, the feedback gains should also reflect an average of the gains specified for each potential target. Here, by measuring participants’

feedback gains while performing reaches towards multiple potential targets, we test this novel and critical prediction of the parallel encoding framework.

In our task, participants performed target-directed reaches towards two competing, superimposed targets of different widths (one narrow and one wide) and we varied the time, relative to reach onset, at which the final target was cued. We explored how visuomotor feedback gains—measured via the magnitude of involuntary rapid corrective responses to rapid shifts in the visual position of the hand halfway through the movement—were modulated as a function of uncertainty about target width during the reach. On the majority of trials, the final target was selected before or immediately after the reach was initiated (single target and early selection trials), providing participants with full certainty about the final target width very early into the movement. Critically, on the remaining trials, the final target was selected only near the very end of the movement (late selection trials), after the corrective reflex response under the situation of target uncertainty (i.e., when each target was equally likely) had already been probed. Importantly, we found that the feedback gains for reaches towards multiple competing targets closely resembled an average of the gains on trials in which the final target (narrow or wide) was known very early into the movement (early selection trials). This was despite the fact that participants, on late selection trials, could just have easily performed the task by implementing the corrective responses associated with the narrow target, as this default control policy would have always ensured task success (given the overlapping potential targets). Rather, these results are consistent with the notion that, when presented with multiple action options, individuals prepare competing sensorimotor control policies in advance of selecting between alternatives and suggest that, during planning, the brain represents, in parallel, complex movement parameters above and beyond kinematics alone.

Results

In brief, our task required participants to perform target-directed reaching movements by moving the handle of a robotic manipulandum (Fig. 1a), which controlled the position of a cursor on a screen, to contact one of two potential, superimposed targets that shared the same centroid, one of which was narrow, the other wide (thus requiring for different accuracy constraints). A countdown procedure following target presentation (5 auditory beeps, each spaced 600 ms apart) provided the instruction for participants to prepare (beeps 1–3), move (4th beep) and then arrive at the final target (5th beep). Targets were either cued (filled-in) prior to movement onset (single target trials) or at two different points (early or late) after the movement had been already launched (two-target trials). Critically, on some trials, the visual position of the cursor on the screen was perturbed midway through the movement (under a visual occluder), requiring a rapid corrective response by participants to contact the target (Fig. 1b). To assess participants' feedback gains in a manner uncontaminated by limb dynamics, we incorporated a force channel on a minority of these trials (Fig. 1c), which mechanically constrained the handle (and thus the hand) to a direct path between the start location and target 22,26. By measuring the corrective forces generated on these channel trials, we were able to probe the visuomotor feedback gains that participants adopted to deal with cases in which the final target was certain at the time of the

cursor perturbation (in one-target and early selection two-target trials) versus uncertain (in late selection two-target trials).

Voluntary corrective responses across one- and two-target trials

The average cursor paths of a representative participant on non-channel trials demonstrates the appropriate corrective responses for cases when the cursor was perturbed to the left or right beneath the occluder, as well as the absence of these corrective responses on trials in which the cursor remained unperturbed (Fig. 2a). These reach paths also generally demonstrate that, as expected, the movement corrections implemented on the narrow target trials (black traces) are greater than those implemented on the wide target trials (green traces), particularly on the single and early target selection trials (see differences in reach traces at the final target locations). This latter observation becomes clearer when examining the reach endpoint distributions across these different target selection conditions (see Fig. 2b,c).

In the single target trials for this participant, we observed more lateral (x) variation in the reach end-point distribution for the wide compared to narrow targets (see Fig. 2b), consistent with previous results 23,24. In addition, this participant exhibited a small, but noticeable increase in endpoint variability when reaching for the narrow target as the duration that target information was withheld from them during the reach increased (i.e., comparing single to early to late selection trials, see Fig. 2b, left panel). In contrast, in the wide target trials, the participant's lateral endpoint distribution became increasingly narrower as a function of increasing target uncertainty duration (see Fig. 2b, right panel). To quantify lateral endpoint variability, we computed, for each participant, target selection condition, and target width, the difference between the mean x endpoint positions for the leftward and rightward cursor jumps. We found a significant effect of target selection condition (Single, Early and Late) on the mean endpoint positions ($F_{2, 14} = 88.231$, $P < 0.0001$), as well as a significant effect for target width (Narrow and Wide; $F_{1, 7} = 193.482$, $P < 0.0001$). We also found a significant interaction between target width and selection condition on endpoint variability ($F_{2, 14} = 160.238$, $P < 0.0001$). The difference in variability between the narrow and wide targets was significant for each of the three conditions (Bonferroni corrected T-Tests; Single: $t_7 = -19.34$, $P < 0.0001$, Early: $t_7 = -11.31$, $P < 0.0001$, and Late: $t_7 = -4.31$, $P = 0.009$; see Fig. 2c) even though the difference was substantially smaller in the late selection trials compared to the single and early selection trials. Likewise, there was a significant effect of selection condition for each target width (Bonferroni corrected ANOVAs; Narrow: $F_{2, 14} = 5.19$, $P = 0.042$, Wide: $F_{2, 14} = 161.9$, $P < 0.0001$). Whereas for the wide target, variability was markedly less in the late selection condition, for the narrow target, variability was slightly greater.

Rapid involuntary corrective responses across one- and two-target trials

In agreement with previous work 22,24, we found that when we examined the forces exerted by the handle in channel trials, over a 180–230 ms time window following the perturbation (gray vertical bars in Fig. 3a), that the corrective force in unperturbed trials was near zero whereas the corrective forces produced in response to the cursor perturbations were in the appropriate direction (i.e., counteracting cursor displacement; Fig. 3a and b shows the

single-trial and median force data from a representative participant, respectively, whereas Fig. 3c shows the mean of the median values across all participants). Moreover, we found that the corrective responses in the single-target and early selection trials appropriately scaled with target size (i.e., larger for narrow targets and smaller for wide targets, see also 23,24). Consistent with past work 22,25, we further found that the corrective forces began showing reliable modulation approximately 150–180 ms after the cursor perturbation onset, validating our use of the 180–230 ms time window for analysis.

In order to test for effects at the group level, we computed, for each target width and selection condition, the average of each individual's median forces for the +30 and –30 perturbations, after negating the latter (negative) forces (see Fig. 4a for the means of these 'rectified' forces across participants). To investigate the effects of different aspects of target uncertainty, we carried out two targeted two-way repeated measures ANOVAs. We first focused on the single-target and early selection trials to assess the influence of initial target uncertainty (present on single-target but not early selection trials) under conditions in which participants had ample time to adjust feedback gains prior to the perturbation. We then focused on the early selection and late selection trials to directly assess the influence of multiple, competing targets at the time of the perturbation (present in late but not early selection trials), while equating for any residual effects of initial target uncertainty (present in both the late and early selection trials).

Single vs. Early Selection Conditions—A 2 (Single-target and Early selection) x 2 (Narrow and Wide target) repeated measures ANOVA revealed main effects of selection condition ($F_{1,7} = 53.413$, $P = 0.0002$) and target type ($F_{1,7} = 70.693$, $P < 0.0001$) on the rectified forces. Notably, however, the interaction between these factors was not significant ($F_{1,7} = 2.191$, $P = 0.182$). We followed up the results of this ANOVA using paired sample two-tailed t-tests. Here, as expected, we found a significant difference in the corrective forces associated with the narrow and wide targets in both the single-target ($t_7 = 7.31$, $P = 0.0002$) and early selection ($t_7 = 6.68$, $P = 0.0003$) conditions (see Fig. 4a).

These results show that whereas participants maintained the same sensitivity in their handling of the visual perturbation with regards to narrow versus wide targets, their overall feedback gains were up-regulated in early selection compared to single-target trials (see Fig. 4a, left and middle plots). This suggests that the initial target uncertainty present on early selection trials, despite the capacity of individuals to rapidly adapt their reflex gains within ~100 ms 25—a latency before participants actually experience the visual perturbation (see Experimental Timing, Online Methods)—had a lasting, residual effect on shaping their rapid corrective responses.

Early vs. Late Selection conditions—A 2 (Early and Late selection) x 2 (Narrow and Wide target) repeated measures ANOVA did not reveal a main effect of selection condition ($F_{1,7} = 3.467$, $P = 0.105$) on rectified force but did reveal a main effect of target size (Narrow or Wide; $F_{1,7} = 33.768$; $P = 0.001$) as well as a significant interaction ($F_{1,7} = 33.285$, $P = 0.001$). This suggests that whereas participants displayed sensitivity to target-size in their feedback gains on early selection trials (as noted in the paired sample t-tests performed above), this was not the case on the late selection trials. A paired sample t-test

between the corrective forces implemented on narrow versus wide targets on late selection trials confirmed that the forces generated did not depend on target width ($t_7 = -0.61$, $P = 0.561$). This finding is to be expected given that the feedback gains on these late selection trials were measured in a time window (180–230 ms) before the target had actually been cued (and made available to the participant).

Averaging of feedback gains in late selection trials

If participants generated an averaged feedback gain under conditions of target uncertainty, then we could expect the gain in late selection trials to be intermediate between the gains for the wide and narrow targets in early selection trials. To test this, we compared, using paired t-tests, the average rectified force generated on late selection trials (i.e., the mean of responses to the narrow and wide target, where a difference was neither expected nor observed) to those generated on each of the early selection trials (i.e., for each the narrow and wide target). We chose the early selection trials as the basis for our comparison given the prominent effect of initial target uncertainty (i.e., the up-regulation of feedback gains) observed for the early- compared to single-target selection conditions (noted above). Notably, this analysis showed that the corrective forces produced on the late selection trials were significantly different from those associated with both the wide ($t_7 = 2.62$, $P = 0.034$) and narrow ($t_7 = 5.10$, $P = 0.001$) targets on early selection trials. Thus, the visuomotor feedback gains observed on the late selection trials, when the final target remains uncertain at the time of perturbation, lies somewhere in between those observed on the early selection trials, when there is full target certainty at the time of perturbation and ample time to fully adjust the gains (see Fig. 4a, compare middle and right plots).

In light of previous evidence suggesting that individuals, when presented with multiple competing reach targets and required to act before knowing the final target location, launch reaches that correspond to the average of movement directions towards each target individually 6,7,9,10,27–29, here we further explored whether individuals on the late selection trials also implement feedback gains that resemble the average of those generated on the early selection trials. To explicitly test this idea, for each participant we computed, from the early selection trials, a ‘synthetic’ average force—the force that would be expected had participants explicitly averaged across the reflex gains associated with the narrow and wide potential targets—and then examined how this synthetic average force relates to the mean force (averaged across narrow and wide target trials) implemented on late selection trials. We derived this latter mean force by computing the average of participants’ rectified forces on narrow and wide target early selection trials. Importantly, for the majority of participants, these data-derived synthetic average force values (black symbols) fall along the identity line, and the average absolute deviation from the identity line was smaller for this average force ($M = 0.163$ N) than for the force generated on either narrow ($M = 0.313$ N, white symbols) or wide ($M = 0.536$ N, green symbols) targets from the early selection trials (see Fig. 4b). When considering the signed deviations from the identity line, t-tests revealed that only the deviations of the synthetic average force values did not significantly differ from zero (Average: $t_7 = 1.862$, $P = 0.105$; Narrow: $t_7 = -2.621$, $P = 0.034$; Wide: $t_7 = 5.096$, $P = 0.001$). These findings suggest that participants, on late selection trials, average across the reflex gains separately computed for the narrow and wide targets on early selection trials.

Feedback gain averaging appears at the onset of testing

To what extent was this average feedback gain on late selection trials something that was implemented immediately at the outset of the experiment and consistently on a trial by trial fashion, rather than something that was gradually learned over the course of testing? When we examined the rectified force as a function of trial block (i.e., a binned average of 5 trials) we found that the gain of the corrective response in late selection trials was—right from the very beginning of experiment—consistently intermediate between the gains for the narrow and wide targets in the early selection trials (see Fig. 4c). This indicates that the averaging of feedback gains was not some strategy that participants learnt to implement over the course of the experiment.

Discussion

Here we explored how individuals regulate their visuomotor feedback gains under conditions of target uncertainty. Notably, rather than using changes in task parameters (e.g., shifts in target location) to introduce target uncertainty on a trial-by-trial fashion 25,30, here we instituted this uncertainty by presenting participants with multiple potential targets and manipulating the time at which one was cued. We report two main findings. First, we show that on trials in which the final target information is not given until very early in the reach (i.e., early selection trials), individuals, despite showing similar target sensitivity to single-target trials, nevertheless exhibit an up-regulation of their feedback gains. This residual and relatively long-lasting (i.e., > 100 ms) effect of target uncertainty on reflex gains has not, to our knowledge, been described elsewhere. Second, and more importantly, we show that on trials in which the final target is selected only after the visual cursor has been perturbed, individuals exhibit feedback gains that are well approximated by the average of the feedback gains implemented towards each of those targets on the early selection trials. This finding, in light of previous neural and behavioural evidence showing that individuals prepare multiple competing reach movements 1,31 and average across their spatial directions when executing movements towards potential targets 6,9, suggests that motor averaging may constitute part of a general mechanism that the brain utilizes when dealing with environmental uncertainty.

Residual Effects of Target Uncertainty on Feedback Gains

Previous studies have shown that the visual presentation of stimuli 32, as well as visual shifts in background 33,34, target location 35,36, and representation of hand position 37,38 all elicit rapid motor responses. Notably, corrective movements in response to these visual displacements (e.g., changes in the position representation of the hand or target) do not require that participants be consciously aware of such changes 35,39. This is consistent with the suggestion that the early components of these visually induced reflexive motor responses rely on involuntary rather than voluntary mechanisms 22,33,34,36. Despite their involuntary nature however, previous work shows that these rapid responses can be flexibly specified prior to movement onset, in accordance with task goals 19–22 and gradually updated based on the learnt dynamics of an environment 18. Although few studies to date have examined modulations of the rapid motor response within the context of a single goal-directed movement, Dimitriou et al 25 found that the central nervous system (CNS) can intelligently modify its feedback responses, based on jumps in the location of a reach target, within 100

ms. In light of this previous evidence, it is notable, though perhaps not surprising, that on our early selection trials (in which the target was cued, on average, 190 ms before the perturbation) participants were able to appropriately up- or down-regulate their feedback gains in accordance with the width of the cued target (to a level of sensitivity observed on single-target trials). What is less clear, however, is why, despite showing this goal-related sensitivity, participants also exhibited an overall up-regulation in their reflex gains on these trials.

Previous work has suggested that the optimal response to increased uncertainty is to decrease one's feedback gains 30,40,41. In these previous studies, however, task uncertainty was introduced through sensory noise (e.g., visual uncertainty of hand position or target location), and thus there is no advantage to setting up robust reflexive responses when the location of the target is unknown. Here, task uncertainty was introduced by presenting multiple potential targets of different widths—both of which were fully known in advance of movement (c.f. 1). Thus, uncertainty was a result of participants' lack of knowledge about which target would be selected, rather than in the sensory processing of those targets. This is an important distinction, as the optimal setting of feedback gains depends on where the uncertainty exists in a given system 42. If uncertainty is coupled to the sensory system, then the optimal control policy will be to reduce feedback responses 18. However, if the uncertainty is coupled to the external world (as in our task), or in one's own model of that world (i.e., the internal model), then the optimal response will be to increase the feedback gains of the system²² (and increase co-contraction, see 43) while also reducing the contributions of predictive control 44. In our experiment, the up-regulation of feedback gains on early selection trials would be consistent with this optimal response. Nevertheless, given previous results 25, we find it surprising that participants still exhibit heightened reflex gains so late into the movement after target cuing (i.e., ~190 ms). This indicates that the uncertainty introduced by multiple competing targets may have a much longer residual effect on the adjustment of reflex gains than previously demonstrated.

Evidence for the parallel encoding of motor plans

Mounting evidence suggests that, in situations affording several possible actions, multiple potential movement plans are represented in parallel. Neural recordings from brain areas involved in eye movements 45, as well as areas involved in arm movements 1,2, both show the simultaneous encoding of multiple competing targets prior to the decision to make an eye or reach movement, respectively, towards one of those locations. Consistent with these observations, psychophysical studies have shown that when individuals are required to initiate an eye (for review, see 8) or reach movement 46,47 prior to knowing which of several potential targets will be cued, their initial movement vector corresponds to a spatial average of the movements performed toward each target separately. Together, this suggests that a basic mechanism by which the brain deals with a dynamic world is to prepare multiple potential actions to available targets, presumably allowing each to be implemented in a moment's notice 3. If this is the case, however, then beyond merely encoding multiple movement directions, one would predict, given the importance of sensorimotor control policies to skilled motor behaviour, that the brain may actually specify, in parallel, the feedback gains associated with each potential movement option. Indeed, according to OFC

models 13,14,16, explicit motor planning involves setting higher-level goals (e.g., get the hand to the target) and specifying the parameters of the feedback controller before each movement. And, at the neural level, it is reasonable to assume that preparatory activity in motor areas, given direct spinal projections, encodes all aspects of the planned movement, including feedback gains associated with the control policy 48.

Here we found that the visuomotor feedback gains implemented on late selection trials (when each target still represents a potential reach option) are well approximated by the average of the gains on the narrow and wide targets on early selection trials. This finding not only provides a novel dimension of the averaging phenomenon described above but, more generally, provides a new line of evidence in support for the parallel encoding framework suggested by the affordance competition hypothesis 3. Specifically, our results are consistent with the idea that the CNS prepares fully elaborated movements, complete with control policies governing feedback gains, for alternate reach options.

Importantly, the fact that an intermediate feedback gain was specified in late selection trials suggests that participants were not overly concerned with maximizing the probability of hitting the target within the specified movement time (i.e., based on automatic corrections), in which case they should have always set their gains for the narrow target. Note that implementing an intermediate gain may lower the cost of control associated with maintaining a high gain 49.

Although spatial averaging behaviour in go-before-you-know tasks is often taken as evidence that the brain encodes multiple potential movements in parallel 9, it has been recently suggested that it may instead constitute a deliberative strategy for minimizing movement-related costs 27. That is, launching a single movement in the spatially averaged direction of potential targets tends to minimize, on average, the cost of in-flight corrective actions that must be taken when one of the potential targets is cued 29. It is difficult to imagine how people would similarly ‘aim towards’ an average feedback gain, given that gains are concerned with the evolution of the movement and are not spatial in nature. It is conceivable that participants indirectly specify an intermediate gain by first constructing an average visual target (with an average width). However, we think this is unlikely as there is strong evidence that, when required to launch a movement towards two potential targets in different spatial locations, participants do not construct a visually averaged target to aim towards 9. Furthermore, we see no evidence that intermediate gains result from learning (i.e., average gains are seen from the earliest trials), which might be expected if participants developed a deliberate strategy of constructing an average visual target.

Our results are in accord with a recent model in which an optimal feedback control policy is calculated independently for each potential target, and a weighted average of these policies (i.e., feedback gains) is computed at each point in time based on the relative desirability of each target 50. Importantly, this model, which predicts averaging of feedback gains, can also account for spatial (i.e., trajectory) averaging in go-before-you-know tasks. We submit that our result showing feedback gain averaging, coupled with previous work demonstrating trajectory averaging, provides strong support for the compelling idea that the CNS, under

cases of target uncertainty, encodes in parallel multiple motor plans, along with their associated control policies, for competing action options.

Online Methods

Participants

Eleven participants (5 men and 6 women aged 18-23) participated in the experiment, with 8 being included in data analysis (see exclusion criteria below). A target sample size of 10-12 participants was specified in advance based on previous studies in this area and our expectation that, if the main experimental effect was present, it should be observed in almost all participants. All participants were right-handed, had normal or corrected-to-normal vision, were neurologically healthy, and naïve to the purpose of the study. They were compensated for their time with a cash payment of \$20-\$32 (see Experimental Timing). The study was conducted with an experimental protocol approved by Queen's University Research Ethics Board, which adhered to the principles of the Canadian Tri-council Policy Statement on Ethical Conduct for Research Involving Humans and the principles of the Declaration of Helsinki (1964).

Experimental Apparatus

Participants were seated in a chair and used their right hand to grasp the handle of a robotic manipulandum (vBOT; Howard et al. 2009) that could move freely in a horizontal plane (Figure 1A). The robot measured the position of the handle and could apply forces to the hand via the handle. Participants were instructed to place their left hand comfortably in their lap. Targets and a cursor representing the position of the handle were displayed on a horizontal computer monitor located above the robot. The participant viewed these stimuli through a mirror located halfway between the monitor and handle, such that the stimuli appeared in the horizontal plane of the handle. The mirror prevented direct visual feedback of the participant's arm and the handle of the manipulandum. The position of the handle was calculated using joint position sensors on the vBOT, which were sampled at 1 KHz.

Experimental Design

Using the robotic handle to control the cursor, participants placed the cursor over a start location (both of which were represented as 1 cm diameter circles) positioned ~20 cm in front of their chest. Two potential, superimposed targets were located 25 cm directly in front of the start location: a narrow 2 cm x 2 cm outlined square or a wide 8 cm x 2 cm outlined rectangle (see black square and green rectangle, Fig. 1B). A 30 cm x 5 cm visual occluder (colored gray) was located midway between the start position and targets. On all trials, the cursor passed behind the occluder, emerging at the exact midpoint of the movement (12.5 cm from the start position). On cursor jump trials, the cursor exited the occluder displaced by 3cm to the left or right ($\pm x$ direction) of the hands position. We chose to perturb the cursor at the mid-point of target distance because it has been previously shown that participants' reflex gains elicit the highest sensitivity at this point in the movement 25. Also, given that visuomotor gains are modulated as a function of distance to the target 25,51, the common perturbation point allowed for a direct comparison of the corrective responses across the different experimental conditions.

Target Selection

Our experiment consisted of both one- and two-target trials. In the one-target trials, only one of the two potential targets appeared at the beginning of each trial and appeared filled-in. In the two-target trials, both potential targets appeared at the beginning of the trial and were initially displayed as an outlined (unfilled) square and rectangle (superimposed) with a border thickness of 3 mm. On early selection trials, once the hand-cursor reached 2.5 cm of target distance, one of the two targets filled-in as the other simultaneously disappeared. On late selection trials, the sequence of events was identical with the exception that the target was only filled-in once the cursor had reached 22.5 cm of target distance. Thus, when the cursor was perturbed (at 12.5 cm distance to target) on these late selection trials, participants could correct for the cursor perturbation prior to actually knowing which target would be eventually cued. Critically, on both early and late selection trials, each target had an equal likelihood of being cued.

Force Channel and Non-Channel Trials

On the majority of trials, the motion of the vBOT handle, and thus the cursor on the screen, was entirely controlled by the participant. Thus, on trials in which the cursor was perturbed, in order for the participant reach the target, they needed to implement corrective responses by moving the handle (Figure 1B). To clearly assess participants' reflex gains, we also, on a minority of trials, incorporated a force channel, which has been used in previous research to examine corrective actions uncontaminated by limb dynamics (e.g., 22,25,26). The force channel was used to mechanically constrain the handle (and thus the participant's hand) to the direct path between the start location and the target. The two 'walls' of the channel were simulated as stiff damped springs (with a stiffness of 4000 N/m and a viscosity of 80 N/m/s) that prevented the handle from moving laterally (see dashed lines in Fig. 1C). The force applied by the robot in order to keep the hand mechanically constrained in the channel, which is equal and opposite the lateral force that the participant applies to the wall of the channel in response to a cursor perturbation, provides a direct read-out of the visuomotor feedback gain that is uncontaminated by lateral motion of the hand. In force channel trials with a cursor perturbation (see example in Fig. 1C), the cursor remained offset (after emerging from the occluder) for 250 ms and then automatically returned to its midline position, allowing participants to successfully complete the task. The time of 250 ms was chosen on the basis of previous work 25, and allowed us an adequate time window in which to measure the automatic reflex response to the perturbation. Whereas the vast majority of participants, at post-experiment debriefing, seemed to have no explicit knowledge that the cursor correction (at 250 ms post-perturbation) was not driven by their own movement, all participants appeared to be unaware that their hand movement path had in fact been constrained (to a straight-ahead movement) on that subset of trials. This is consistent with previous work that has used channel trials to explore feedback gain modulation 22,25 and force-field adaptation 26, and may also reflect the fact that our force channel cursor perturbation trials occurred relatively infrequently during testing (for details, see below).

Experimental Timing

Once the cursor was positioned over the start location for 200 ms, a progression of 5 beeps, each spaced 600 ms apart, were played. Participants were instructed to leave the start position on the 4th beep and arrive at the target on the 5th beep. If the reach was initiated too quickly (> 200 ms before the fourth beep) or too slowly (> 200ms after the fourth beep), participants, upon completion of the trial, were given the text feedback “Too Early” or “Too Late”, respectively, on the screen. If the total movement duration (time from start position to target) was greater than 800 ms, participants were given the text feedback “Too Slow” following the trial. If the movement duration was less than 400 ms, participants were given the text feedback “Too Fast” following the trial. On trials in which participants did not commit these timing errors, and depending on whether participants hit or missed the target (the criterion for a hit being whether the pixels of the cursor overlapped with the pixels of the target), the text feedback “Good” or “Miss”, respectively, was displayed following the trial. This feedback, in addition to encouraging similar accuracy demands across participants, encouraged consistent timing across trials and participants. For instance, the amount of time it took (in channel trials) for participants to move, in the direction of the target, from 0 to 25 mm (i.e., from the start position to the point where the early target selection occurred), from 25 mm to 125 mm (the point where the perturbation occurred), from 125 mm to 240 mm (the front edge of the target) and from 225 mm to 240 mm (from the point where the late target selection occurred to the front edge of the target) were as follows (range of the lowest-to-highest and average movement duration, based on participant medians): 0 to 25 mm: 136–143 ms, M = 139 ms; 25 mm to 125 mm: 186–192 ms, M = 190 ms; 125mm to 240mm: 260–275 ms, M = 269 ms; 225mm to 240mm: 55–62 ms, M = 60 ms). When adding each of the first three mean times together, we obtained a 597.3 ms movement time, only 2.7 ms less than 600 ms, the time interval between the 4th and 5th beeps (the first providing the cue for participants to move and the latter providing the desired target contact time).

Given the observation by Dimitriou et al.²⁵ that individuals can reliably update their reflex gains within 100 ms, this consistency in the experimental timing of our task is important in several ways. First, it allowed participants, on early selection trials, adequate time (mean of 190 ms) between target cuing (at 25 mm of reach distance) and cursor perturbation (at 125 mm of reach distance) to update their visuomotor feedback gains in accordance with the cued target size. Second, given that the average duration between the time of cursor perturbation and late target selection is approximately 210 ms (i.e., 269 ms minus 60 ms), it is highly unlikely that the visual perturbation-induced corrective responses measured on these two-target trials (at 180–230 ms post perturbation onset) is contaminated by any of the subsequent corrective responses associated with target cuing that occurs later on during the movement (which would require, at minimum, an additional ~100 ms to be implemented, i.e., 310 ms). Third, the time between target cuing and target contact on the late selection trials (mean of 60 ms) also makes it highly unlikely that participants could have deliberately adjusted their reflex gains in accordance with the cued target size, even if they wished to do so.

All participants were paid \$20 for participation and, in addition, earned \$0.01 for every “Good” trial, allowing them to earn \$20–\$32 based on their performance. The person with the highest number of “Good” trials also won a gift card to a popular local restaurant.

Experimental Conditions

We had 4 general experimental conditions, based on whether the force channel was present or absent and whether the cursor was perturbed or unperturbed. The entire experiment contained 1200 trials in total and was presented in 5 blocks of 240 trials each, with each experimental block having the same proportions of experimental conditions. The trials in each block were fully randomized. For non-channel and non-perturbation trials, participants performed 60 trials for each of the one-target and two-target early and late selection conditions (30 for each target size; 180 trials total); for non-channel and perturbation trials, participants performed 240 trials for each of the three target selection conditions (120 for each combination of target size and selection condition; 720 trials total); for channel and non-perturbation trials, participants performed 20 trials for each of the three target selection conditions (10 for each target size; 60 trials total); and lastly, for the key channel and perturbation trials, participants performed 80 trials for each of the three target selection conditions (40 for each target size; 240 trials total). Experimental testing was completed over two days. On day one, participants performed 75 practice trials, in order to familiarize themselves with the task, and Blocks 1 and 2. On day 2, they performed experimental Blocks 3–5. Testing on each day lasted approximately 1.5 hours.

Data analysis

To measure participant feedback gains, we used the forces generated on the channel trials. Participants were included for analysis only if the data on their channel trials met the following basic criteria: 1) They exhibited less median force on the channel in non-perturbation than perturbation trials, and, 2) Their median direction of corrective force counteracted the cursor perturbation, and, 3) They exhibited stronger force responses for narrow than wide targets. Three of 11 participants (1 women and 2 men) did not meet these criteria and were excluded from further analysis. Two of these participants were excluded for not meeting the 3rd criteria (i.e., exhibiting stronger force responses for narrow targets than for wide targets on one-target trials) and the remaining participant was excluded for not meeting the 2nd criteria (i.e., their corrective force did not counteract the cursor perturbation direction on one-target trials).

To calculate each participant’s visuomotor feedback gains, we computed, for each channel trial, the mean force exhibited over the time window of 180–230 ms following perturbation onset in line with 22,25—a time window over which corrective responses are thought to be uncontaminated by voluntary responses c.f. 22. Then, for each participant and for each of the 18 experimental conditions [Target type (2; Narrow or Wide) x Selection condition (3; Single, Early, or Late) x Perturbation direction (3; –30mm, 0, +30mm)], the median of these mean forces was computed. We used participant medians rather than means to guard against outliers. Next, to derive participant force measures that are independent of the direction of cursor perturbation, for each participant we computed rectified median forces by multiplying

participants' median force responses to the -30 mm perturbation by -1 , and then averaging them across the corresponding median force responses for the $+30$ mm perturbation.

To obtain a measure of participant's movement endpoints, we computed the x and y position of the participant's hand on non-channel trials once their velocity slowed to 20 mm/s. Mean x and y values (in mm) were then calculated for endpoints corresponding to the leftward ($-x$) and rightward ($+x$) perturbations for each participant over the 3 selection conditions and 2 target types. The mean difference in the lateral plane was then calculated ($+x - (-x)$) so as to provide an approximation of the overall width of the endpoint distribution, and thus a measure of the amount of correction implemented on non-channel trials.

Statistical Analysis

No statistical methods were used to predetermine sample sizes, but our sample sizes are similar to those reported in previous publications 52–56. Data distribution was assumed to be normal but this was not formally tested. Significance level was set at $P < 0.05$ and all data are reported as mean \pm SEM across subjects.

Acknowledgements

This research is supported by the Natural Sciences and Engineering Research Council of Canada; The Wellcome Trust; the Human Frontiers Science Program; and by the Royal Society Noreen Murray Professorship in Neurobiology (to D. M. W.). J.P.G. was supported by CIHR postdoctoral fellowship award. We thank D. Franklin for providing helpful advice and suggestions.

References

1. Cisek P, Kalaska JF. Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron*. 2005; 3:801–814.
2. Thura D, Cisek P. Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron*. 2014; 81:1401–1416. DOI: 10.1016/j.neuron.2014.01.031 [PubMed: 24656257]
3. Cisek P. Cortical mechanisms of action selection: the affordance competition hypothesis. *Phil Transact R Soc Lond B*. 2007; 362:1585–1599.
4. Baumann MA, Fluet MC, Scherberger H. Context-specific Grasp Movement Representation in the Macaque Anterior Intraparietal Area. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2009; 29:6436–6448. [PubMed: 19458215]
5. McPeck RM, Han JH, Keller EL. Competition between saccade goals in the superior colliculus produces saccade curvature. *Journal of neurophysiology*. 2003; 89:2577–2590. [PubMed: 12611995]
6. Chapman CS, et al. Reaching for the unknown: multiple target encoding and real-time decision-making in a rapid reach task. *Cognition*. 2010; 116:168–176. DOI: 10.1016/j.cognition.2010.04.008 [PubMed: 20471007]
7. Ghez C, et al. Discrete and continuous planning of hand movements and isometric force trajectories. *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale*. 1997; 115:217–233. [PubMed: 9224851]
8. Van der Stigchel S, Meeter M, Theeuwes J. Eye movement trajectories and what they tell us. *Neurosci Biobehav Rev*. 2006; 30:666–679. [PubMed: 16497377]
9. Stewart BM, Gallivan JP, Baugh LA, Flanagan JR. Motor, not visual, encoding of potential reach targets. *Current biology : CB*. 2014; 24:R953–954. DOI: 10.1016/j.cub.2014.08.046 [PubMed: 25291634]

10. Stewart BM, Baugh LA, Gallivan JP, Flanagan JR. Simultaneous encoding of the direction and orientation of potential targets during reach planning: evidence of multiple competing reach plans. *Journal of neurophysiology*. 2013; doi: 10.1152/jn.00131.2013
11. Shadmehr R, Smith MA, Krakauer JW. Error correction, sensory prediction, and adaptation in motor control. *Annual review of neuroscience*. 2010; 33:89–108. DOI: 10.1146/annurev-neuro-060909-153135
12. Wolpert DM, Diedrichsen J, Flanagan JR. Principles of sensorimotor learning. *Nature reviews. Neuroscience*. 2011; 12:739–751. DOI: 10.1038/nrn3112 [PubMed: 22033537]
13. Scott SH. Optimal feedback control and the neural basis of volitional motor control. *Nature reviews. Neuroscience*. 2004; 5:532–546. DOI: 10.1038/nrn1427 [PubMed: 15208695]
14. Todorov E, Jordan MI. Optimal feedback control as a theory of motor coordination. *Nature neuroscience*. 2002; 5:1226–1235. DOI: 10.1038/nrn963 [PubMed: 12404008]
15. Wolpert DM, Flanagan JR. Motor learning. *Current biology : CB*. 2010; 20:R467–472. DOI: 10.1016/j.cub.2010.04.035 [PubMed: 20541489]
16. Todorov E. Optimality principles in sensorimotor control. *Nature neuroscience*. 2004; 7:907–915. DOI: 10.1038/nrn1309 [PubMed: 15332089]
17. Cluff T, Scott SH. Rapid feedback responses correlate with reach adaptation and properties of novel upper limb loads. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2013; 33:15903–15914. DOI: 10.1523/JNEUROSCI.0263-13.2013 [PubMed: 24089496]
18. Franklin S, Wolpert DM, Franklin DW. Visuomotor feedback gains upregulate during the learning of novel dynamics. *Journal of neurophysiology*. 2012; 108:467–478. DOI: 10.1152/jn.01123.2011 [PubMed: 22539828]
19. Diedrichsen J. Optimal task-dependent changes of bimanual feedback control and adaptation. *Current biology : CB*. 2007; 17:1675–1679. DOI: 10.1016/j.cub.2007.08.051 [PubMed: 17900901]
20. Dimitriou M, Franklin DW, Wolpert DM. Task-dependent coordination of rapid bimanual motor responses. *Journal of neurophysiology*. 2012; 107:890–901. DOI: 10.1152/jn.00787.2011 [PubMed: 22072514]
21. Pruszynski JA, et al. Primary motor cortex underlies multi-joint integration for fast feedback control. *Nature*. 2011; 478:387–390. DOI: 10.1038/nature10436 [PubMed: 21964335]
22. Franklin DW, Wolpert DM. Specificity of reflex adaptation for task-relevant variability. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2008; 28:14165–14175. DOI: 10.1523/JNEUROSCI.4406-08.2008 [PubMed: 19109499]
23. Nashed JY, Crevecoeur F, Scott SH. Influence of the behavioral goal and environmental obstacles on rapid feedback responses. *Journal of neurophysiology*. 2012; 108:999–1009. DOI: 10.1152/jn.01089.2011 [PubMed: 22623483]
24. Knill DC, Bondada A, Chhabra M. Flexible, task-dependent use of sensory feedback to control hand movements. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2011; 31:1219–1237. DOI: 10.1523/JNEUROSCI.3522-09.2011 [PubMed: 21273407]
25. Dimitriou M, Wolpert DM, Franklin DW. The temporal evolution of feedback gains rapidly update to task demands. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2013; 33:10898–10909. DOI: 10.1523/JNEUROSCI.5669-12.2013 [PubMed: 23804109]
26. Smith MA, Ghazizadeh A, Shadmehr R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS biology*. 2006; 4:e179.doi: 10.1371/journal.pbio.0040179 [PubMed: 16700627]
27. Haith AM, Huberdeau DM, Krakauer JW. Hedging your bets: intermediate movements as optimal behavior in the context of an incomplete decision. *PLoS computational biology*. 2015; 11:e1004171.doi: 10.1371/journal.pcbi.1004171 [PubMed: 25821964]
28. Gallivan JP, et al. One to four, and nothing more: nonconscious parallel individuation of objects during action planning. *Psychological science*. 2011; 22:803–811. [PubMed: 21562312]
29. Hudson TE, Maloney LT, Landy MS. Movement planning with probabilistic target information. *Journal of neurophysiology*. 2007; 98:3034–3046. DOI: 10.1152/jn.00858.2007 [PubMed: 17898140]

30. Izawa J, Shadmehr R. On-line processing of uncertain information in visuomotor control. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2008; 28:11360–11368. DOI: 10.1523/JNEUROSCI.3063-08.2008 [PubMed: 18971478]
31. Gallivan JP, Barton KS, Chapman CS, Wolpert DM, Flanagan JR. Action plan co-optimization reveals the parallel encoding of competing reach movements. *Nature Communications*. 2015; 2015doi: 10.1038/ncomms8428
32. Corneil BD, Olivier E, Munoz DP. Visual responses on neck muscles reveal selective gating that prevents express saccades. *Neuron*. 2004; 42:831–841. [PubMed: 15182721]
33. Saijo N, Murakami I, Nishida S, Gomi H. Large-field visual motion directly induces an involuntary rapid manual following response. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2005; 25:4941–4951. DOI: 10.1523/JNEUROSCI.4143-04.2005 [PubMed: 15901775]
34. Gomi H, Abekawa N, Nishida S. Spatiotemporal tuning of rapid interactions between visual-motion analysis and reaching movement. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2006; 26:5301–5308. DOI: 10.1523/JNEUROSCI.0340-06.2006 [PubMed: 16707782]
35. Prablanc C, Martin O. Automatic control during hand reaching at undetected two-dimensional target displacements. *Journal of neurophysiology*. 1992; 67:455–469. [PubMed: 1569469]
36. Day BL, Lyon IN. Voluntary modification of automatic arm movements evoked by motion of a visual target. *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale*. 2000; 130:159–168. [PubMed: 10672469]
37. Saunders JA, Knill DC. Visual feedback control of hand movements. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2004; 24:3223–3234. DOI: 10.1523/JNEUROSCI.4319-03.2004 [PubMed: 15056701]
38. Brenner E, Smeets JB. Perceptual requirements for fast manual responses. *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale*. 2003; 153:246–252. [PubMed: 14610635]
39. Goodale MA, Pelisson D, Prablanc C. Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature*. 1986; 320:748–750. [PubMed: 3703000]
40. Kording KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature*. 2004; 427:244–247. DOI: 10.1038/nature02169 [PubMed: 14724638]
41. Braun DA, Aertsen A, Wolpert DM, Mehring C. Learning optimal adaptation strategies in unpredictable motor tasks. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2009; 29:6472–6478. DOI: 10.1523/JNEUROSCI.3075-08.2009 [PubMed: 19458218]
42. Orban G, Wolpert DM. Representations of uncertainty in sensorimotor control. *Current opinion in neurobiology*. 2011; 21:629–635. DOI: 10.1016/j.conb.2011.05.026 [PubMed: 21689923]
43. Franklin DW, Osu R, Burdet E, Kawato M, Milner TE. Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *Journal of neurophysiology*. 2003; 90:3270–3282. DOI: 10.1152/jn.01112.2002 [PubMed: 14615432]
44. Crevecoeur F, McIntyre J, Thonnard JL, Lefevre P. Movement stability under uncertain internal models of dynamics. *Journal of neurophysiology*. 2010; 104:1301–1313. DOI: 10.1152/jn.00315.2010 [PubMed: 20554851]
45. Munoz DP, Wurtz RH. Saccade-related activity in monkey superior colliculus. II. Spread of activity during saccades. *Journal of neurophysiology*. 1995; 73:2334–2348. [PubMed: 7666142]
46. Song JH, Nakayama K. Hidden cognitive states revealed in choice reaching tasks. *Trends in cognitive sciences*. 2009; 13:360–366. [PubMed: 19647475]
47. Gallivan JP, Chapman CS. Three-dimensional reach trajectories as a probe of real-time decision-making between multiple competing targets. *Frontiers in neuroscience*. 2014; 8:215.doi: 10.3389/fnins.2014.00215 [PubMed: 25100941]
48. Churchland MM, Santhanam G, Shenoy KV. Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. *Journal of neurophysiology*. 2006; 96:3130–3146. DOI: 10.1152/jn.00307.2006 [PubMed: 16855111]

49. Manohar SG, et al. Reward Pays the Cost of Noise Reduction in Motor and Cognitive Control. *Current biology* : CB. 2015; 25:1707–1716. DOI: 10.1016/j.cub.2015.05.038 [PubMed: 26096975]
50. Christopoulos V, Schrater PR. Dynamic Integration of Value Information into a Common Probability Currency as a Theory for Flexible Decision Making. *PLoS computational biology*. 2015; 11:e1004402.doi: 10.1371/journal.pcbi.1004402 [PubMed: 26394299]
51. Liu D, Todorov E. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *The Journal of neuroscience : the official journal of the Society for Neuroscience*. 2007; 27:9354–9368. DOI: 10.1523/JNEUROSCI.1110-06.2007 [PubMed: 17728449]
52. Ingram JN, Flanagan JR, Wolpert DM. Context-dependent decay of motor memories during skill acquisition. *Current biology* : CB. 2013; 23:1107–1112. DOI: 10.1016/j.cub.2013.04.079 [PubMed: 23727092]
53. Huang VS, Haith A, Mazzoni P, Krakauer JW. Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron*. 2011; 70:787–801. DOI: 10.1016/j.neuron.2011.04.012 [PubMed: 21609832]
54. Galea JM, Mallia E, Rothwell J, Diedrichsen J. The dissociable effects of punishment and reward on motor learning. *Nature neuroscience*. 2015; 18:597–602. DOI: 10.1038/nn.3956 [PubMed: 25706473]
55. Baugh LA, Kao M, Johansson RS, Flanagan JR. Material evidence: interaction of well-learned priors and sensorimotor memory when lifting objects. *Journal of neurophysiology*. 2012; 108:1262–1269. DOI: 10.1152/jn.00263.2012 [PubMed: 22696542]
56. Ingram JN, Howard IS, Flanagan JR, Wolpert DM. Multiple grasp-specific representations of tool dynamics mediate skillful manipulation. *Current biology* : CB. 2010; 20:618–623. [PubMed: 20346672]

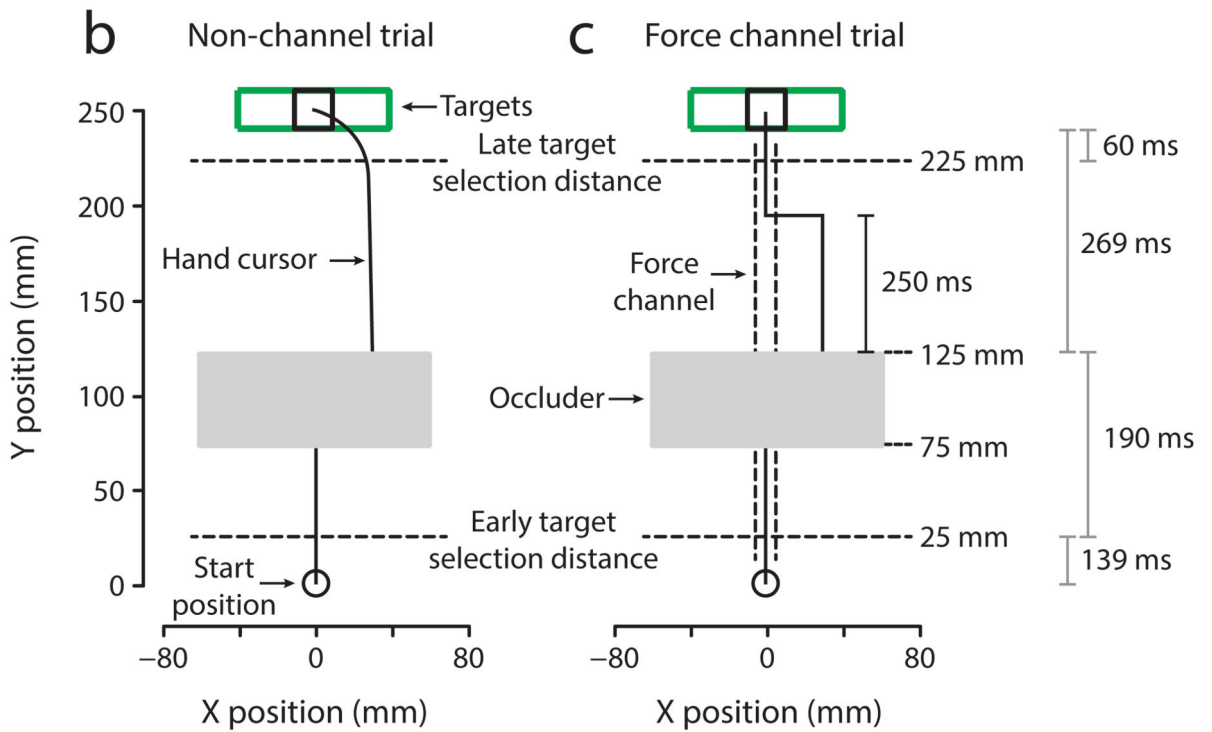



Fig. 1. Experimental Methods.

A) Experimental setup. Participants ($N = 8$) were seated in a chair and grasped, with their right hand, the handle of a robotic manipulandum (vBOT) that controlled the position of a cursor on a monitor. Visual feedback from the monitor was viewed through a mirror, located in the plane of movement, and which prevented view of the participant’s arm. **B-C)** Examples of experimental conditions. On a subset of trials, the cursor was visually perturbed (to the left or right; right perturbation shown) exactly half-way through the movement (125 mm), after it passed under an occluder (gray horizontal bar). In the Non-Channel condition

(B), participants had to use the handle to correct the position of the cursor in order to reach the target. In the Force Channel condition (C), which constrained the participants movement to a straight ahead path (denoted by dashed black vertical lines) the cursor position automatically jumped back after 250 ms. In B and C, the black and green boxes show the positions of the narrow and wide targets, respectively. Target selection (filling-in) could either occur before movement onset (single target trials), early in the movement (at 25 mm Y position; early selection trials) or late in the movement (at 225 mm Y position; late selection trials). The average movement durations (across participants and conditions, in ms) associated with different epochs of the trial are displayed at right. Note: the x-length of the occluder has been shortened for display purposes. All other dimensions are drawn exactly to scale.

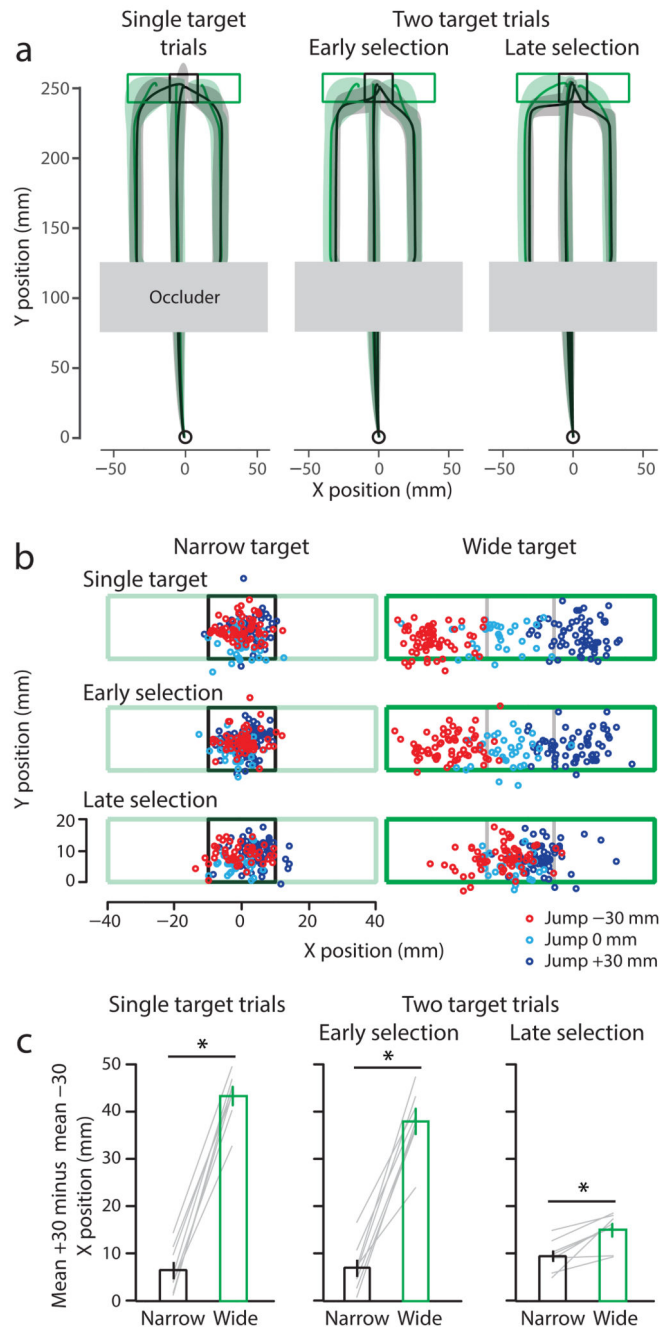


Fig. 2. Hand paths on non-channel trials and movement endpoint variance for each target selection condition

A) Averaged trajectory traces for a representative participant. Left panel: Single target condition. Black rectangle and trajectory traces represent the narrow target (20 mm wide) and the associated movement paths to that target, respectively. Green rectangle and trajectory traces represent the wide target (80 mm wide) and the associated movement paths to that target, respectively. Gray horizontal bar represents the occluder. Middle panel: Early target selection condition. In these trials, target selection (filling-in) occurs when the

participant's cursor passes 25 mm of reach distance. Right panel: Late target selection condition. In these trials, target selection occurs when the participant's cursor passes 225 mm of reach distance. **B)** Reach endpoints, for the participant shown in A, plotted as a function of robot x-y position for each of the three experimental conditions. Endpoints are color-coded according to cursor jump direction (left and right directions, in red and dark blue, respectively, and no jump, in light blue). **C)** Mean group (N = 8) differences in reach endpoint, plotted as a function of differences in mean x position (mean x position of the rightward perturbations minus the mean x position of the leftward perturbations), for each of the three experimental conditions. The lines represent means from individual participants and error bars denoted ± 1 SEM. *denotes significance at $p < 0.05$.

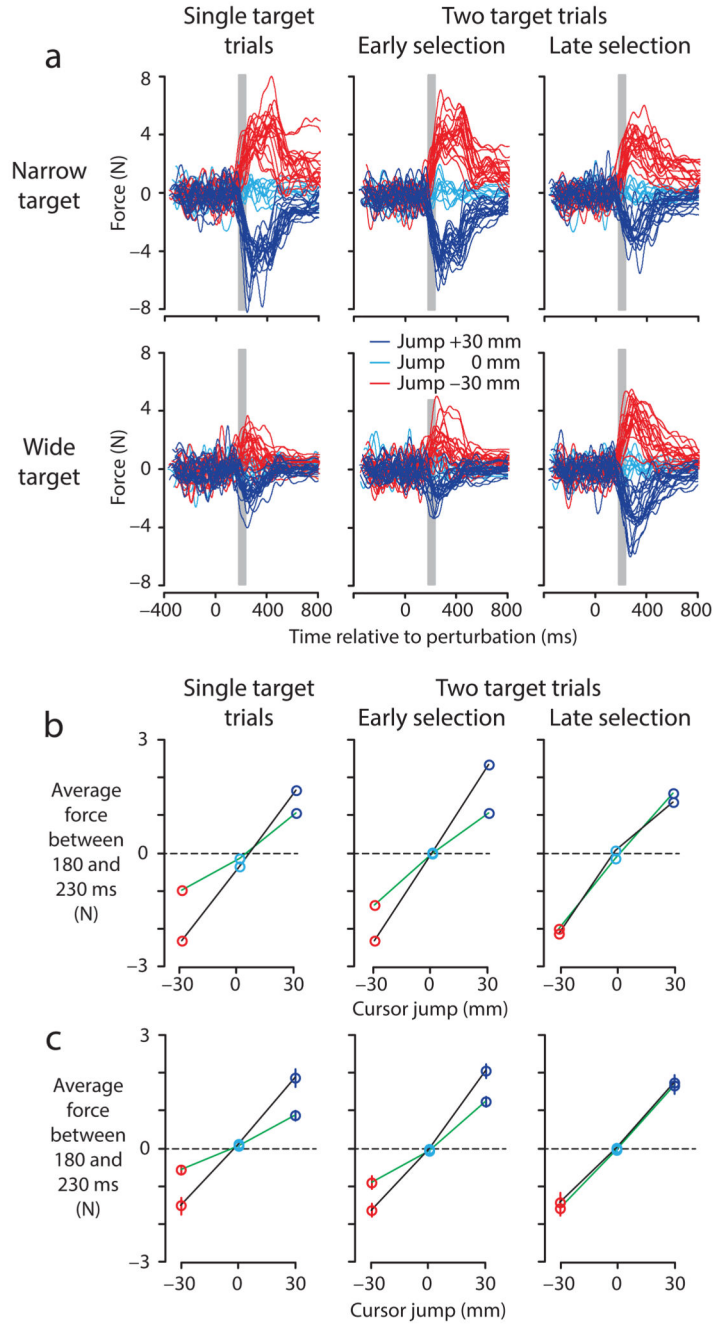


Fig. 3. Scaling of feedback gains (on channel trials) across target selection conditions.
A) Force readings, plotted as a function of time relative to cursor perturbation, for all trials and all conditions for a single, representative participant. 0 ms denotes the onset of cursor perturbation. Shaded gray vertical bars denote 180–230 ms post-perturbation—the time window reflecting the involuntary feedback response. The mean robot force value over this time window was extracted for each trial type and participant, and the median of these mean values were used for group-level analyses. **B)** Median forces for the 180–230 ms time window and same participant shown in A. Force traces in A and open circles in B are color-

coded according to the perturbation condition (light blue: unperturbed, dark blue and red: left and right perturbation, respectively). The color of the lines joining the open circles indicates the associated target size (green: wide target, black: narrow target). **C)** Mean of median forces across participants (N=8), plotted the same as in B. Small vertical lines in the open circles denote ± 1 SEM.

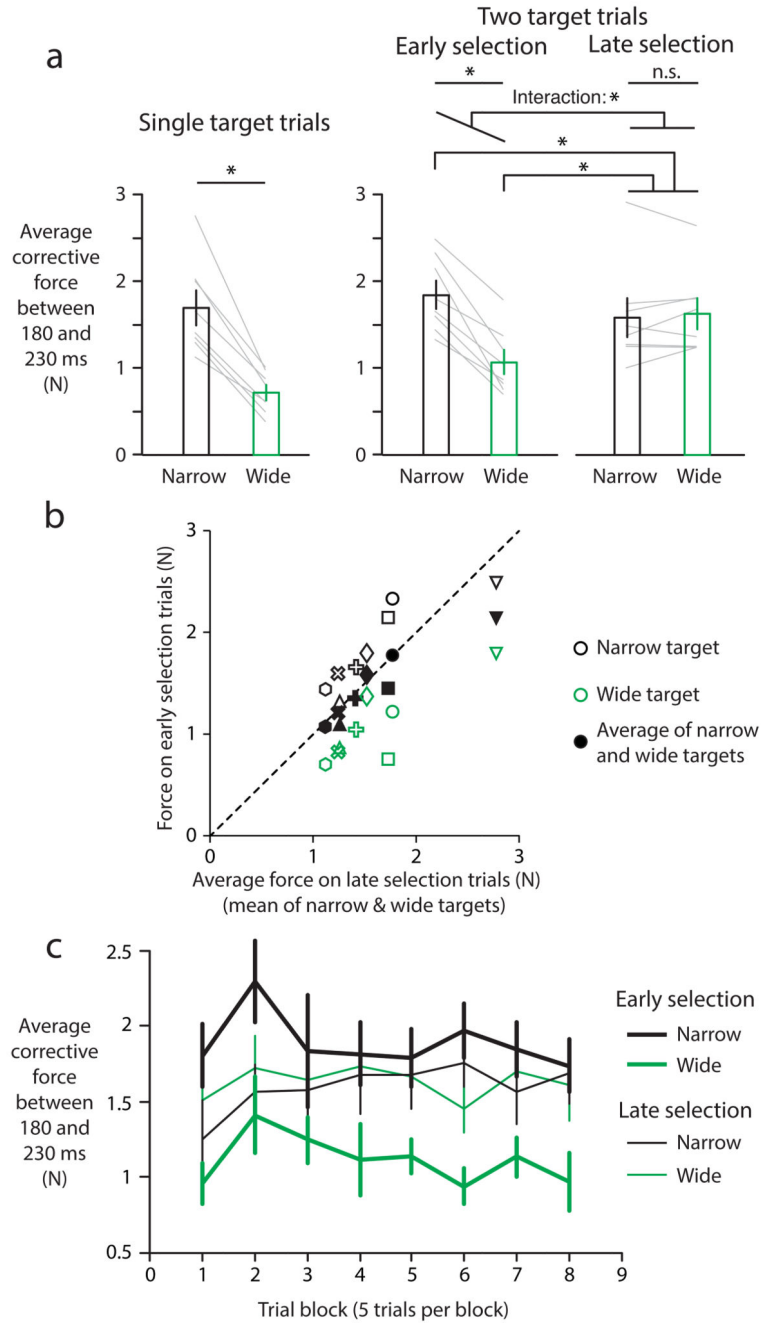


Fig. 4. Group-level analysis of feedback gains.

A) Rectified forces (whereby forces from -30 mm perturbations are multiplied by -1), averaged across participants, for the data shown in Figure 3C. The lines represent medians from individual participants and error bars denoted ± 1 SEM. * denotes significance at $p < 0.05$. **B)** Relationship between mean force on late selection trials (x-axis, from rightmost plot in A) and early selection trials (y-axis, from middle plot in A). Filled black data points denote an average of the force response on narrow (open black points) and wide (open green points) target early selection trials. Different symbols represent different participants.

Dashed line represents the unity line. C) Rectified force as a function of trial block (averaged across bins of 5 trials) for early and late selection trials with each target.