A very fast time scale of human motor adaptation: within movement adjustments of internal representations during reaching

Abbreviated Title: Adaptive Control of Reaching

F. Crevecoeur[1,2], J.-L. Thonnard[1,2,3], P. Lefèvre[1,2]

[1]*Institute of Communication and Information Technologies, Electronics and Applied Mathematics (ICTEAM), University of Louvain (UCLouvain), Belgium*

[2]*Institute of Neuroscience, UCLouvain, Belgium*

[3]*Cilniques Universitaires Saint-Luc, UCLouvain, Belgium*

Corresponding Author:

F. Crevecoeur

4 avenue Georges Lemaitre

1348 Louvain-la-Neuve, Belgium

email: frederic.crevecoeur@uclouvain.be

Summary

Humans and animals adapt motor commands to predictable disturbances within tens of trials in laboratory conditions. However, adaptation occurs continuously in natural conditions, and how quickly the brain learns to handle novel disturbances has remained unexplored. Because motor commands and sensory feedback together carry continuous information about limb dynamics, we hypothesized that the nervous system could adapt to unexpected disturbances online. We tested this hypothesis in two reaching experiments during which opposing velocity-dependent force fields were randomly applied. We found that within-movement feedback corrections gradually improved, despite the fact that the perturbations were unexpected. Moreover, when participants were instructed to stop at a via-point, the application of a force field prior to the via-point induced mirror-image after-effects after the via-point, indicative of within-trial adaptation to the unexpected dynamics. These findings highlight a fast time-scale of motor learning, which supplements feedback control and supports adaptation of an ongoing movement.

Keywords: reaching control; motor adaptation; motor learning, internal models; sensorimotor control.

Introduction

Neural plasticity in the sensorimotor system enables adaptive internal representations of movement dynamics and acquisition of motor skills with practice (Dayan and Cohen, 2011). In the context of reaching movements, studies have documented that healthy humans and animals can learn to anticipate the impact of a force field applied to the limb, and recover straight reach paths within tens to hundreds of trials (Krakauer et al., 1999; Lackner and DiZio, 1994; Shadmehr and Mussa-Ivaldi, 1994; Singh and Scott, 2003; Thoroughman and Shadmehr, 2000; Wagner and Smith, 2008). Importantly, studies on motor learning have consistently highlighted the presence of an after-effect, which mirrors the initial movement deviation and indicates that adaptation was supported by a novel internal model of movement dynamics (Lackner and DiZio, 2005; Shadmehr et al., 2010; Wolpert et al., 2011).

To date, motor learning and adaptation have been exclusively studied on a trial-by-trial basis, highlighting context-dependent learning rates and memory dynamics across movements (Diedrichsen et al., 2010; Gonzalez Castro et al., 2014; Kording et al., 2007; Smith et al., 2006; Wei and Koerding, 2010). Although this approach has revealed several fundamental properties of sensorimotor systems, it has left unexplored the problem of online control during early exposure to the altered environment. More precisely, it remains unknown how the nervous system controls reaching movements in the presence of unexpected dynamics, or errors in the internal models, as occurs when a force field is introduced experimentally.

It is often assumed that unexpected disturbances during movements are automatically countered by the limb's mechanical impedance and by reflexes (Burdet et al., 2000; Franklin and Wolpert, 2011; Shadmehr and Mussa-Ivaldi, 1994). However, recent work suggests that muscles viscoelastic properties, as well as the gain of the spinal reflex (latency ~20ms for upper limb muscles) are low at spontaneous levels of muscle activation (Crevecoeur and Scott, 2014). Furthermore, long-latency reflexes (latency ≥50ms) are also based on internal models of the limb and environmental dynamics (Cluff and Scott, 2013; Crevecoeur and Scott, 2013; Kurtzer et al., 2008), even when

disturbances are very small (Crevecoeur et al., 2012). Thus, the presence of model errors is equally challenging for rapid feedback control as it is for movement planning, and yet healthy humans can handle unexpected disturbances relatively well.

Thus, the outstanding question is whether and how the nervous system controls movements online when exposed to unexpected dynamics. In theory, an approximate internal model can be deduced during movement because motor commands and sensory feedback together carry information about the underlying dynamics. This problem was studied in the framework of *adaptive control* (Bitmead et al., 1990; Fortney and Tweed, 2012; Ioannou and J, 1996). The idea is to adapt the model continuously based on sensory feedback about the unexpected movement error. It is important to distinguish adaptive control from standard (extended) stochastic optimal control, which is based on the assumption that there is no fixed model error, and that the disturbances in the process follow a known distribution (Izawa et al., 2008; Todorov, 2005). The goal of adaptive control is to correct for the possible presence of model errors and unexpected movement deviations continuously.

In light of the theory of adaptive control and considering the very rich repertoire of rapid sensorimotor loops (Scott, 2016), we explored the possibility that mechanical disturbances to the limb update the internal representations of movements online. An online estimation of task-related parameters was reported in the context of virtual visuomotor rotations (Braun et al., 2009), where participants had to estimate the mapping between hand and cursor motion. Here we address this problem in the context of reaching in a force field to probe specifically the internal representations of limb dynamics. We show that healthy humans adapt internal representations within a movement, and illustrate how their behaviour can be captured in the framework of adaptive control. We also highlight the associated computational challenges for the nervous system that have yet to be fully characterized.

Results

Experiment 1

Participants grabbed the handle of a robotic device and performed forward reaching movements towards a visual target (Fig. 1a). In the first experiment, clockwise and counterclockwise orthogonal force fields (FF) were randomly applied as catch trials, such that neither the occurrence nor the direction of the force field could be anticipated. Reach paths were of course strongly impacted by the presence of the force field, however the online correction became smoother with practice (compare trials #1 and #20 in Fig. 1b). To quantify this, we extracted the maximum hand deviation along the direction of the force field, and the target overshoot near the end of the movement (Fig. 1b: Max. and Osh., respectively). We found that the maximum deviations as well as their timing did not evolve significantly across force field trials (Fig. 1c, $P>0.05$), whereas the overshoot displayed a significant exponential decay (Fig. 1d, $P<0.001$, $R^2 = 0.17$ and 0.56 for clockwise and counterclockwise force fields, respectively). The absence of a clear change in the maximum deviation confirms that there was no anticipation, or default strategy that could counter the unexpected disturbance.

We also observed that exposure to force field trials induced standard after-effects on the next trial, as previously reported in both standard and a random adaptation experiments (Lackner and DiZio, 1994; Scheidt et al., 2001; Shadmehr and Mussa-Ivaldi, 1994). To observe this, we separated trials performed in the null field dependent on whether they were preceded by force field trials or by at least 3 null field trials (Fig 1e). Because we were interested in online control, the use of force channels in our experiment was not suitable (that is, a virtual wall that constraints the reach path, see e.g. (Smith et al., 2006)). Instead, we extracted the x-coordinate of the hand path at the moment of lateral peak hand force (Fig. 1f) (this moment was well defined as it was induced by the inertial anisotropy of the KINARM handle). The lateral coordinate was clearly impacted by the occurrence of a force field trial, and the effect was opposite to the perturbation encountered in the preceding trial as expected (Fig. 1g, rmANOVA: $F_{(2,26)}=18.3$, $P<10^{-4}$).

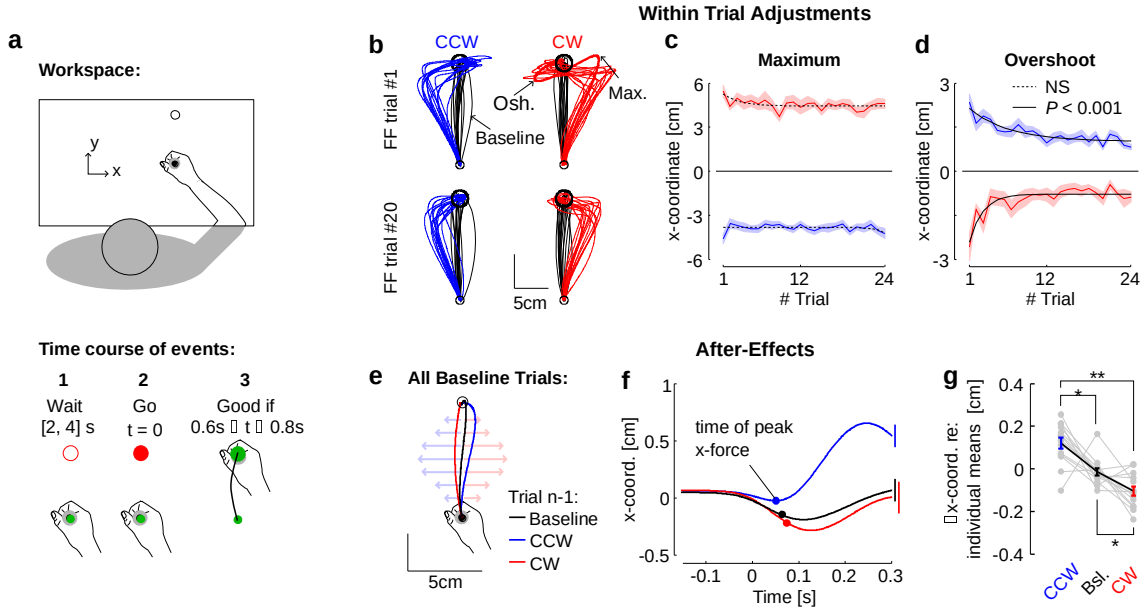

Figure 1: Experiment 1, Behavior. a: Illustration of a successful trial in Experiment 1 (n=14). Participants were instructed to wait for the go signal, and then to reach for the goal target within 0.6 s to 0.8 s of the go signal. Feedback about movement timing was provided to encourage participants to adjust their movement speed, but all trials were included in the dataset. b: Hand paths from all participants for the first force field trial (top) and the force field trial number 20 (bottom), were selected to illustrate the change in online control. Counter-clockwise (CCW) and clock-wise (CW) FF trials are depicted in blue and red, respectively. The black traces (baseline) are randomly selected trials from each participant. Maximum deviation in the direction of the force field (Max.) and the maximum target overshoot (Osh.) are illustrated. c: Maximum hand path deviation for CCW and CW force fields across trials. The dashed traces illustrate the non-significant exponential fits. d: Maximum target overshoot across FF trials. Solid traces indicate significant exponential decay. Shaded blue and red areas around the means indicate the SEM across participants. e: Illustration of the after effect. Baseline trials were categorized dependent on whether they were preceded by CCW or CW force fields, or by at least 3 baseline trials. f: Continuous traces of x-coordinate as a function of time and extraction at the moment of peak lateral hand force (dots). Error bars at 300ms illustrate the SEM across participants. g: individual averages of lateral hand coordinate (gray), and group averages (mean±SEM), after subtracting each individual's grand mean for illustration). One (two) star(s) indicate significant difference at the level $P < 0.05$ (0.01) based on paired t-test, adjusted for multiple comparison with Bonferroni correction.

The reduction in target overshoot resulted from a smoother correction for the lateral displacement, in the sense that the target overshoot were reduced. Indeed, there was no significant change in the peak forward hand velocity across trials for both clockwise (rmANOVA, $F_{(23,299)} = 0.56$, $P = 0.9$) and counterclockwise force fields ($F_{(22,286)} = 0.87$, $P = 0.62$). Thus the perturbation applied to the limb remained statistically similar. However, the absolute peak hand force near the end of the

movement decreased significantly. This is illustrated in Fig. 2, where the average lateral displacement as a function of time is shown (Fig. 2a) along with the lateral forces (Fig. 2b). We observed a significant exponential decay of the terminal peak hand force (Fig. 2c, R^2 =0.17 and 0.11 for clockwise and counterclockwise perturbations). As a consequence, the correlation between the measured lateral force and the force field perturbation, corresponding to the forward velocity multiplied by the scaling coefficient of the force field, significantly increased across trials (Fig. 2d). In the case of perfect compensation, the measured force should equal the action of the force field, and a correlation of 1 is expected. Thus, the increase in correlation between the measured force and the force field within trials indicated that the online correction became progressively tuned to the properties of the force field.

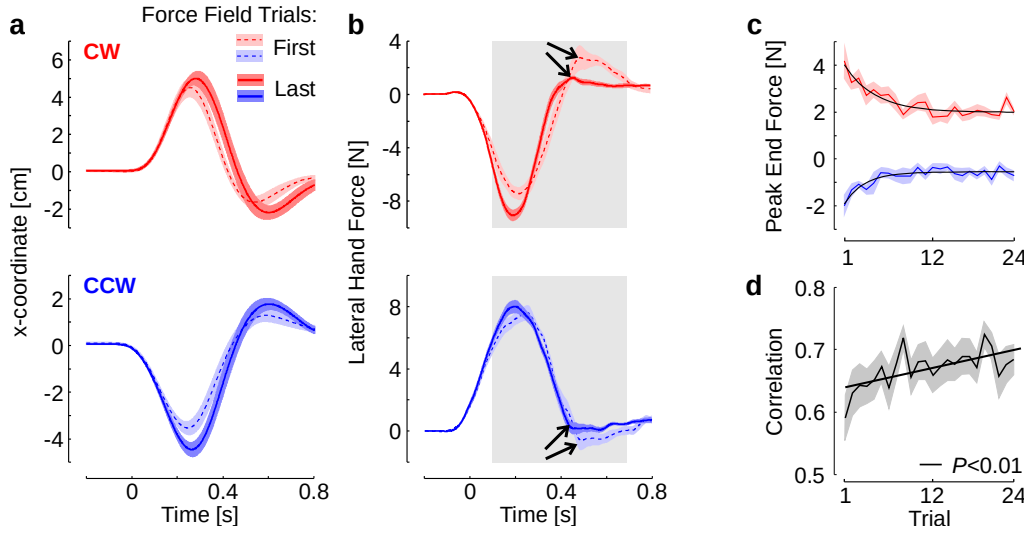

Figure 2. Improvement of Online Corrections. a: lateral coordinate as a function of time for the first (dashed/light) and last (solid/dark) force field trials. Observe that the difference of ~1cm in peak hand displacement for CCW trials was visible in Fig. 1c, but there was no significant change for the following trials. **b**: Same as panel a for the lateral force. Shaded areas represent one SEM across participants (n=14) with the same color code as in Fig.1. The reduction in peak end force is highlighted with the black arrows. **c**: Peak end force across force field trials (mean±SEM). The solid traces show the significant exponential decay (P<0.005). **d**: Correlation between the lateral hand velocity and the force generated by the force field during the time interval corresponding to the gray rectangle in panel a (from 100ms to 700ms). The correlations were averaged across participants; the shaded area represents one SEM. The significant linear regression of the correlations across blocks is illustrated.

The results of Experiment 1 suggest that the controller changed during movement. Indeed, the absence of an effect on the maximum lateral deviation indicated that participants started the movement without anticipating the force field ("baseline" controller, $C(B)$), and then changed to a controller that at least partially accounted for the force field ($C(F_i)$ for the ith force field trial). It is important to realize that if the nervous system always switched to the same controller during force field trials, it would have been impossible to identify a change in representation from a standard feedback response. The evidence for a change in representation emerges from the fact that $C(F_i)$ itself changed across force field trials ($C(F_i) \neq C(F_j), i \neq j$), without anticipation and hence based on sensory information collected during the ongoing trial. Interestingly the presence of an after-effect indicated that these within-trial changes in control are likely linked to the fast time scales of motor learning (Smith et al., 2006).

Experiment 2

The results of Experiment 1 are consistent with adaptive control. However it can be objected that participants learned to produce smooth corrective movements by altering movement kinematics without learning about the force field specifically, or by altering the mechanical impedance of the limb without changing their internal model (Burdet et al., 2001). To address this, we sought a more direct link between the online correction and the force field by using a via-point located at two thirds of the reach path (see Methods). We reasoned that if the feedback correction during force field trials reflected an online update of the internal models, then these corrections should evoke almost instantaneous after-effects for the remainder of the movement (from the via point to the goal target).

Our results confirmed this prediction. First we reproduced the observations made in the first experiment for the trials without via-point: we found no significant exponential decay in the maximum lateral deviation, whereas the target overshoot displayed significant exponential decay across trials (Fig. 6 a-d). Second, we found a systematic after-effect between the via-point and the

193 goal target that mirrored the impact of the force field prior to the via-point (Fig. 3e for average hand
194 paths and 3f). Stopping at the via-point was necessary for two reasons: first it allowed us to turn off
195 the force field unnoticeably because the velocity was close to zero. Second, it also ensured that the
196 after-effects were not due to momentum following the online correction prior to the via-point,
197 because the velocity was again very close to zero at the via-point. Thus, the after-effects between
198 the via-point and the goal target were due to a genuine re-acceleration in the direction opposite to
199 the force field just experienced. We extracted the minimum lateral hand velocity in the via-point and
200 the lateral hand velocity at the second peak (Fig. 3g: squares and open discs on exemplar traces). As
201 expected, there was no difference across clockwise, baseline, and counterclockwise trials for the
202 minimum velocity (Fig. 3h), which allowed us to verify that there was no difference in momentum
203 across trials at the via-point. The same analysis with the norm of hand velocity in place of lateral
204 velocity gave identical results. In contrast, the lateral hand velocity measured at the second
205 maximum displayed a very strong modulation consistent with an after-effect (Fig. 3h, rmANOVA on
206 average lateral hand speed across force field: $F_{(2,14)} = 17.23$, $P < 0.001$). In addition, we observed that
207 eight out of ten participants displayed the same effect when comparing trial-by-trial hand velocity of
208 CW and CCW trials (Wilcoxon ranksum test, $P < 0.05$).
209

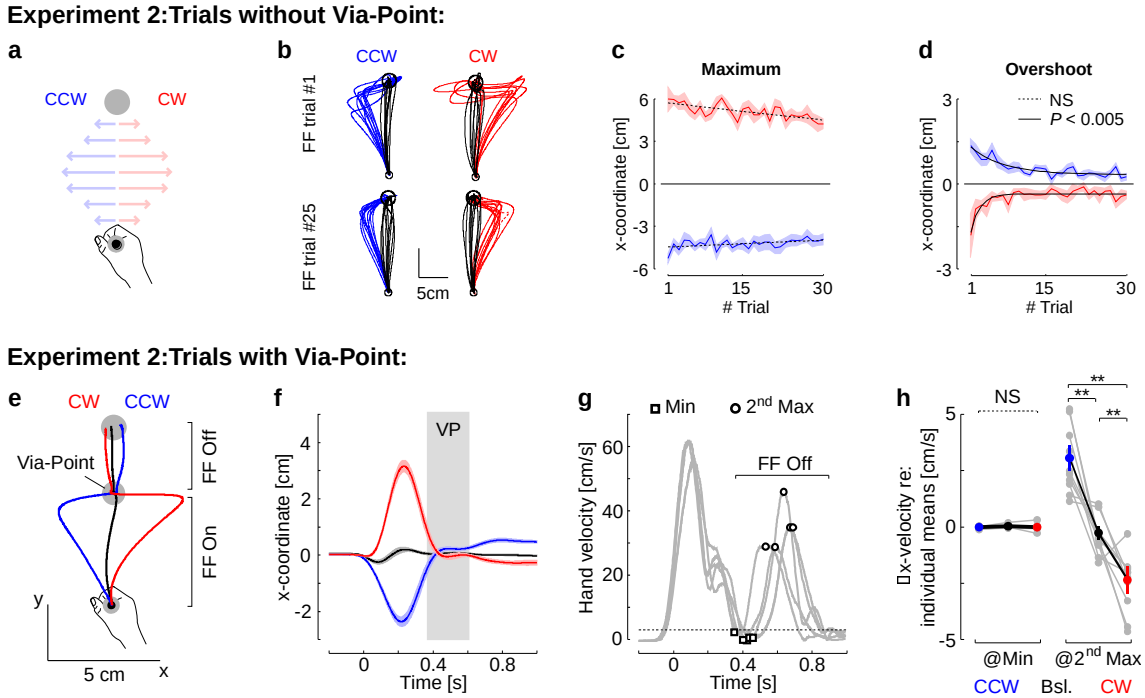

Figure 3: Experiment 2 Behavior. a-d: illustration of the force field trials and behaviour as reported in Fig. 1b-d. e: Average hand traces from Experiment 2 for trials with a via-point with the same color code as in Experiment 1 (black: null-field, red: CW, blue: CCW). For all trials, the second portion of the reach, from the via-point to the goal target, was performed without force field (FF Off). f: Average x-coordinate as a function of time across participants (n=10). Shaded areas represent one SEM across participants. The gray rectangle illustrates the approximate amount of time spent within the via-point. Observe the displacement following the via-point in the opposite direction as the perturbation-related movement. g: Illustration of the trial-by-trial extraction of parameters on five randomly selected trials from one participant. The force field was turned off when the hand speed dropped below 3 cm/s (dashed trace). The parameters extracted were the lateral hand velocity at the moment of the hand speed minimum (squares), and the lateral hand velocity at the moment of hand speed maximum (open discs). h: Lateral hand velocity at the moment of hand speed minimum or at the moment of the second hand speed maximum averaged across trials for each participant (gray dots), and averaged across participants (blue, black and red, mean±SEM). Individual means across the three trial types were subtracted from the data for illustration purpose. The two stars represent significant pair-wise differences at the level $P<0.01$ with Bonferroni correction for multiple comparisons.

By forcing participants to stop at the via-point, we were able to control for the effect of momentum experimentally. We conducted additional analyses to address whether movement kinematics prior to the via-point had an impact on the trajectories after the via-point. To address this issue, we regressed the lateral velocity at the second peak as a function of lateral velocity extracted at the first velocity peak, prior to the via-point (Fig. 4a). We fitted a linear mixed model on individual trials from all participants while including a random intercept to account for idiosyncrasy. The

regression was computed first on the baseline trials, and we found no significant correlation between the lateral velocity before and after the via-point ($F_{(244)}$ = -0.81, P = 0.41, Fig. 4b, dashed line). We then included force field trials and found a significant relationship between velocity prior and after the via-point ($F_{(633)}$ = 157, P < 0.0001). Importantly, the specific impact of the force field was established by the presence of a significant interaction between the first velocity peak, and the trial type (mixed model ANOVA, : $F_{(2,633)}$ = 4.24, P=0.014, Fig. 4b: observe the ellipses above and below the regression for baseline trials). Thus we found a clear modulation of hand velocity after the via-point compatible with an after effect, while controlling for momentum experimentally, and for non-specific effects linked to movement kinematics statistically.

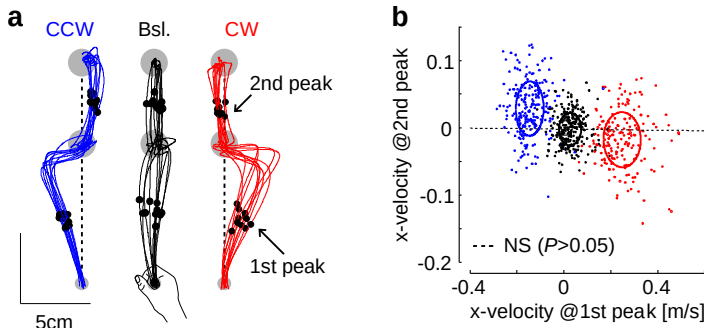

Figure 4. Kinematics prior to the via-point. a. Extracted parameters. The lateral velocity (x-direction) was extracted at the moment of the first and second peaks of hand speed prior and after the via-point, repectively (dots and arrows). Vertical dashed lines were added to highlight the systematic deviation opposite to the perturbation during force field trials. We selected 10 traces randomly per condition from one representative participant. b. Lateral velocity at the second peak as a function of the lateral velocity at the first peak. Dots are all individual trials from all participants and dispersion ellipses illustrate one standard deviation along the main variance axes. Observe that the statistical analyses were performed by taking each individual's offset into account.

We were then interested to see how early the modulation of the lateral hand velocity following the via-point occurred. In addition to the main effect of the force field, we found a significant interaction between the force field and the block number (Fig. 5a, $F_{(10,70)}$ = 2.34, P=0.018). The repeated measures ANOVAs were performed on the data from the 8 participants who completed at least one successful via-point trial in each block to balance the test. The paired comparisons were performed with the data from 8 participants in the first block, and 9 in block number 6. Strikingly, the

effect reported in Fig. 3 was already present in the first block (Fig. 5b, paired t-test on data from block 1: $t_8 = 3.6$, $P = 0.007$, effect size: 1.2, data from block 6: $t_9 = 5.8$, $P < 0.001$, effect size: 1.95). Again we verified that there was no difference across force-field conditions and blocks for the minimum hand velocity ($F_{(10,70)} = 0.59$, $P = 0.8$), and thus no residual momentum, which could have confounded the modulation of hand velocity after the via-point.

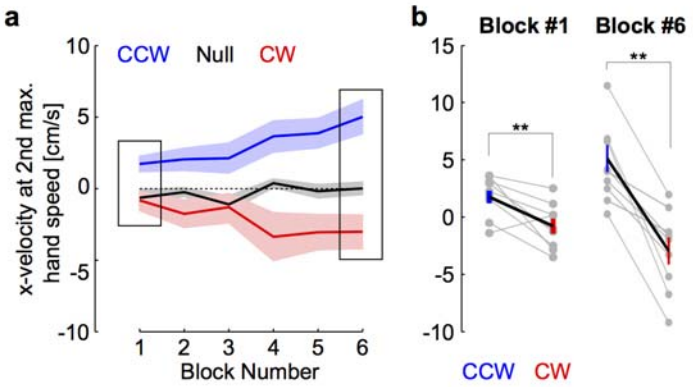

Figure 5: Learning across Blocks. a: Lateral hand velocity measured at the second peak hand speed following CCW and CW force-fields (red and blue, respectively), and baseline trials (black) across blocks. There were blocks in which participants did not succeed on any force-filed trial, thus displays are the mean±SEM across participants in each block with $8 \leq n \leq 10$. The insets highlight the first and last block. b: Individual average across trials for each participant (gray dots), and group mean±SEM in the first and 6th block (n = 9 in both cases). The two stars highlight significant difference based on paired t-test ($P < 0.01$).

The presence of a significant modulation in the first block indicated that online adaptation occurred during the first few trials (there were 5 trials with via-point and switching-off of the force field per direction and per block). This striking result warranted further investigation: can these adjustments occur within the very first trial? Our dataset cannot provide a definitive answer because, due to randomization, some via-point trials were preceded by force field trials without via-point for some participants. Nevertheless, we used a statistical model of the modulation of lateral velocity as a function of trial number to obtain the beginnings of an answer.

We first changed the sign of velocity during CW trials such that an after effect would correspond to a positive velocity for all trials. Then, we accounted for idiosyncrasy by including

286 participants as a random factor in a mixed linear model. We found a significant positive correlation
287 between the lateral hand velocity and the trial index. Most importantly, the regression included an
288 offset that was significantly distinct from 0. The intercept value was $0.014ms^{-1}$, and the *P*-values for
289 the slope and intercept were 0.014 and 0.006, respectively. The same regression performed on the
290 baseline trials, where no force field was applied prior to the via point, revealed an offset of -0.007,
291 which was not significant ($P = 0.06$). Thus, we can reject the hypothesis that the lateral hand velocity
292 was zero in the very first trial. In other words, our dataset supports the hypothesis that online
293 adjustments occurred even during the very first trial. Lastly, we compared the lateral hand velocity
294 from the very first successful trial in each direction and from each participant, and also found a
295 significant modulation of hand velocity following the via-point (CCW>CW, one-tail paired t-test, $t_{(8)}$ =
296 1.91, $P = 0.044$).

297 Experiment 2 suggested that the internal representation dynamics was updated during the
298 movement prior to the via-point, and could modify the previously planned movement from the via-
299 point to the goal target in a way that was consistent with an after-effect. Statistical analyses showed
300 that these adjustments may have occurred during the very first trial. It is possible that participants
301 used three internal models, one for null-field trials and one per force-field direction, and switched
302 between them dependent on the ongoing movement. But in order to explain the data, it is necessary
303 to account for the fact that participants did not previously adapt to the force field, and that the after
304 effect increased across the blocks (Fig. 5). A more compelling view in our opinion, which explains the
305 results of the two experiments as well as the increase in modulation reported in Experiment 2, is the
306 hypothesis that the nervous system uses adaptive control to track model parameters online. This is
307 illustrated in the next section.

308

309 **Adaptive Control Model**

310 A first important distinction to make is that adaptive changes in the internal models differ
311 from the sensorimotor loops that correct for perturbations to the state of the limb (Fig. 6a: inner

loop, solid). In the inner loop, sensory feedback is used to estimate the state of the system through an observer (e.g. Kalman filter), and the controller generates motor commands that correct for motor errors. This loop is responsible for adjusting motor commands following, for instance, mechanical disturbances to the limb or target jumps.

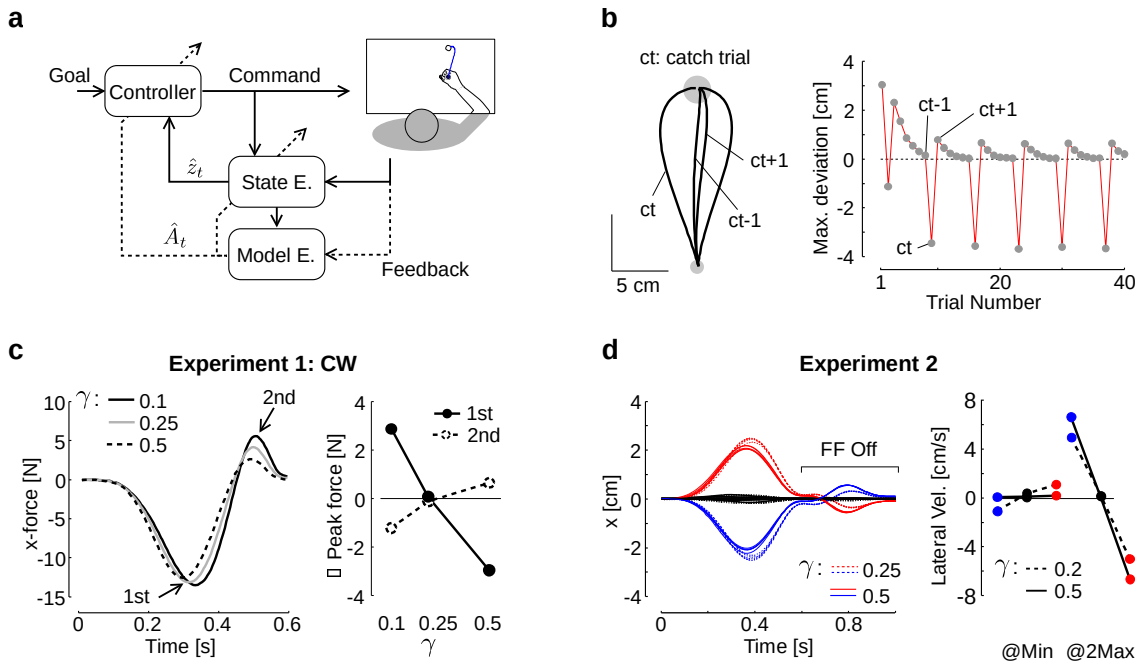

Figure 6. Model Simulations. a: Schematic illustration of an adaptive controller: sensory feedback is used to estimate both the state of the system and the model, which in turn updates the controller and estimator online. The solid arrows represent the inner loop correcting for state deviations and the dashed arrows represent the outer loop correcting for model errors. b: Reproduction of a standard learning experiment (hand paths and maximum lateral deviation across trials), with mirror-image catch trials (ct), and unlearning following the catch trials (difference between ct-1 and ct+1). The trial-by-trial changes were reproduced with online computations exclusively (inspired by Thoroughman and Shadmehr (Thoroughman and Shadmehr, 2000)). c: Simulation of the results from Experiment 1 with three distinct values of online learning rate (0.1 in black, 0.25 in gray, and 0.5 in dashed), which reduces the second peak hand force while the first displays smaller changes across the tested range of γ. The inset represent the change in peak force relative to the mean across simulations, observe that the first peak hand varies by less than 2N, while the second peak displays changes >4N. See the main text for correspondences with experimental data. d: Simulations of behavior from Experiment 2, the force field was turned off at the via-point (t = 0.6s), and the second portion of the reach displays an after effect (see Methods). The inset highlights the lateral velocity measured at the minimum and at the second peak hand speed as for experimental data. The traces were simulated for distinct values of online learning rate (0.25 in dashed, and 0.5 in solid).

336 In the framework of adaptive control, there is a second loop (Fig. 6a: outer loop, dashed),
337 which estimates the model based on current sensory data and motor commands (Bitmead et al.,
338 1990; Fortney and Tweed, 2012; Ioannou and J, 1996). This estimation is indirect because there is no
339 measurement of the unknown parameter, thus motor commands and sensory feedback must be
340 used together to deduce the underlying dynamics. The simulations performed in this study were
341 based on iterative least square identification (RLS) (Bitmead et al., 1990), but other techniques such
342 as expectation-maximization may be considered (Ghahramani and Hinton, 1996). However it is
343 important to note that the increase in after-effect observed across the blocks of Experiment 2
344 required a degree of freedom in the model to account for trial-by-trial modulation, which is why we
345 privileged the RLS formalism as it includes such parameter by design (online learning rate, γ, see
346 Methods). The outer loop updates the model used in the estimator and in the controller of the inner
347 loop (Fig. 6a, oblique arrows). The second experiment was designed to highlight the presence of this
348 loop: at the via-point with near-zero velocity, there is no deviation or error in the state of the limb
349 that requires correction by the inner loop. Yet, participants did not simply reach straight ahead
350 because of a change in the internal models performed by the outer loop during the first portion of
351 the reaching movement.

352 Using this framework, we reproduced standard learning curves in computer simulations, with
353 exponential decay of the lateral hand displacement, mirror-image catch trials or after effects as
354 documented in Experiment 1 (Fig. 1 e-g), and unlearning following the catch trials as previously
355 reported in human learning experiments (Fig. 6b) (Thoroughman and Shadmehr, 2000). This model
356 also reproduced the reduction in peak end force occurring near the end of movement as observed in
357 Experiment 1 by assuming that the online learning rate increased across force field trials (Fig. 6c, see
358 also Methods), that is the online adjustments were smaller in the first trials than in the last trials.
359 This idea is consistent with the presence of savings, characterizing a faster re-learning upon exposure
360 to a previously experienced environment (Gonzalez Castro et al., 2014; Smith et al., 2006). Here we
361 simply hypothesized a similar faster relearning when exposed to a previously experienced

perturbation based on the continuous inflow of sensory information and an increased online learning rate (γ parameter, note that γ was fixed within each trial). This adjustment of online learning rate is required to account for the changes in control observed in Experiment 1 (Fig. 6c), and for the increase in lateral hand velocity observed for the via-point trials in Experiment 2 (Fig. 6d). Finally, we also reproduced the presence of after-effects after switching off the force field at the via-point and extracting the lateral velocity exactly as in Experiment 2 (Fig. 3c and 6d).

We did not fit the model parameters to account for all observed results quantitatively (see Methods), instead we were interested to capture these effects qualitatively and highlight that the outer loop was necessary to explain the data. We should stress that the model predicts small variations in the first peak hand force and larger variations for the second peak hand force (Fig. 6c). In our dataset we found no systematic change in the first peak, which can be explained by the fact that there was no delay in the model, whereas the adjustments in practice occurred necessarily later due to the presence of sensorimotor delays. The second peak hand force displayed changes up to 4N with the tested values of γ, thus the variations observed experimentally (~2N) are clearly in a range that can be explained in the model. Likewise, the amplitude of the aftereffect from the simulations of Experiment 2 could explain the range observed experimentally. Observe that increasing the online learning rate in the model not only increased the aftereffect, but also improved the stopping at the via-point as the minimum velocity was closed to zero (Fig. 6d). Finally, we observed that an exponential decay similar to previous adaptation studies was obtained with comparatively lower values of γ ($\gamma = 0.05$ in Fig. 6b).

To summarize, participants improved their online corrections when exposed to unexpected force field perturbations, without previous acquisition of internal representation of these perturbations. Furthermore, similar force field perturbations evoked near instantaneous after effect in the via-point experiment. These effects suggest that the internal representations changed within a trial, which is faster than previously recognized in typical human reaching experiments. Furthermore, the change in representation increased across trials, which is inconsistent with the idea of switching

between a few fixed controllers. Instead, we identified a theoretical framework that captures all of these findings based on online tracking of model parameters.

Discussion

Controlling movements in the presence of model errors is a challenging task for robotics as well as for the nervous system. However healthy humans can handle model errors like those arising when a force field is introduced experimentally during reaching. To understand this, we explored the possibility that the nervous system reduces the impact of model errors online following the principles of adaptive control. We found that participants' feedback corrections to unexpected force fields gradually improved (Experiments 1 and 2), and evoked after-effects during an ongoing movement (Experiment 2). We showed that these effects, as well as standard learning curves, were captured in a simple adaptive control model in which motor commands, state estimates, and sensory feedback are used to update the internal models of dynamics online. The hypothesis that the nervous system learns online could capture our findings qualitatively, as well as the typical exponential decay observed in human motor adaptation. Thus it constitutes a powerful candidate model for linking online control and motor execution with the fast time scales of trial-by-trial motor learning.

We should emphasize that there is no disagreement between adaptive control and stochastic optimal control (LQG, (Todorov, 2005)), but instead there is complementarity between these models. Stochastic optimal control is based on the assumption that the system dynamics are known exactly, and that the process disturbance is known in distribution (Astrom, 1970). The unexpected introduction of a force field introduces a fixed error in the model parameters, which is not a stochastic disturbance. Such disturbances warrant the use of a control design that explicitly considers the presence of errors in the model parameters. Adaptive control aims at solving this problem.

In order to delineate the limitation of our study it is important to identify candidate alternative mechanisms. Our hypothesis is that the internal model of reaching dynamics changed within a trial. Alternative models that assume no change within a trial cannot explain all our experimental results. First, a default modulation of limb mechanical impedance would have impacted the maximum deviation as well as the overshoot, which is inconsistent with the results of Experiment

1. A stiffening or damping mechanism after the peak displacement could have produced smoother online corrections, but such stiffening could not account for the after effect highlighted in Experiment 2, since more stiffening would predict less lateral deviation between the via-point and the goal target (Burdet et al., 2001).

Second, these results seem also compatible with a form of switching from one controller to another, showing that separate motor memories of opposite perturbations can be learned and retrieved within movement. Consistent with this idea, previous work showed that distinct planning conditions were critical to learn opposite perturbations, such as explicit cues (Schaal et al., 2004), distinct representations of movements (Hirashima and Nozaki, 2012), or even distinct prior or follow-through movements associated with the same movement (Howard et al., 2012; Sheahan et al., 2016). In our study, there was no cue given about the upcoming perturbation, and movements' internal representations were identical across baseline and force field trials at movement onset. Thus, it was the physical state of the limb during movement that evoked the changes in online control.

It is important to stress that such a switching between controllers, is in fact compatible with adaptive control with a limited, or finite set of representations. But our results highlighted a more complex mechanism. Indeed, participants' behaviour during force field trials varied from trial to trial in both experiments, showing that the putative switching controllers were not fixed. This result required a change in online learning rate in the model, which is similar to the concept of savings applied to online learning. Furthermore, it is important to recall that the force fields were not learned a priori, thus the switching does not explain how these representations were first acquired.

Finally, alternative hypotheses that assume no update of movement representation in ≤500ms cannot account for the results of both experiments. A fair objection is that the instantaneous after effect after the via-point can be viewed as a second sub-movement, and hence that this effect is not distinct from previous reports on motor adaptation. It is clear that we cannot make any claim about the underlying mechanism and it is possible it is the same process that generates online

adjustments as the process responsible for trial-by-trial learning. Regardless of the number of sub-movements and of the mechanism, it is an empirical fact that adjustments occurred in less than 0.5ms, which is much faster than previously identified time scales of motor adaptation. Our main contribution is thus to show that motor learning is perhaps better considered as a real-time and online process instead of a discrete process that influences motor performances on a trial-by-trial basis. We identified the framework of adaptive control to account for our results. The strength of the model is to show that it is not necessary to assume that previously learned internal models are retrieved, because parameter tracking is possible continuously, and in theory in real time. We suggest that online tracking in real-time combined with a change in online learning rate provides the most comprehensive account of our data.

It remains unclear how quickly these adjustments occurred within movements, and this represents an important but challenging question for future research. Anatomically, long-latency corrections for perturbation to the limb (~60ms) are supported by a distributed network through cerebellum (Hore and Flament, 1988), primary somatosensory area, motor and pre-motor cortices, and parietal cortex (Omrani et al., 2016; Pruszynski et al., 2011; Scott, 2016). This network clearly overlaps with the main regions associated with short-term plasticity (Dayan and Cohen, 2011). Behaviorally, long-latency feedback exhibits exquisite sensitivity to even very small disturbances (Crevecoeur et al., 2012). Thus, the network of brain regions engaged in trial-to-trial learning is also likely recruited very early during unexpected force field trials, and it is conceivable that rapid adjustments occurred within this rapid pathway.

Computationally, however, the correction of internal models involves a less direct mapping between sensory feedback and the internal representations of movement dynamics, in addition to a change in the neural controller. This computational load may require longer processing times associated with model updates in comparison with feedback responses following learned perturbations to the limb or to the movement goal. As an absolute upper bound, we may consider that adjustments in Experiment 2 were performed between reach onset and the moment when

468 participants exited the via-point, which was around 0.5s. Similarly, visual inspection of the forces
469 from Experiment 1 indicates that the reduction in lateral force, consistent with the adaptive control
470 model, occurred in about 0.4s. We focus on this timing because these changes in force were directly
471 linked to the reduction in target overshoot, but it remains possible that changes occurred earlier.
472 However the precise timing of change in feedback correction cannot be unambiguously determined
473 from the resultant force since this variable also depends on the forward kinematics. We expect that
474 future work might aim at measuring the latency of this outer loop more accurately.

475 Another clear challenge for future work is to unravel the underlying neural mechanism. The
476 key question is to understand how the brain generated a near instantaneous after effect like those
477 observed in Experiment 2. Indeed, at the via-point with near-zero velocity across all trials, there was
478 no motor error to correct; yet the subsequent motor output mirrored the previously encountered
479 disturbances. One possibility to account for this result is through synaptic plasticity (Dayan and
480 Cohen, 2011): that is the gain of the same sensorimotor loop changes following re-afference of
481 sensory prediction errors, and associated changes in synaptic strength. Another possibility is that the
482 unexpected force field generated changes in the state of the sensorimotor system along a dimension
483 which does not influence the motor output (Kaufman et al., 2014; Stavisky et al., 2017), allowing that
484 a similar biomechanical state (i.e. stopping at the via-point) be followed by a distinct motor output
485 after the via-point. Such mechanism does not necessarily engage changes in synaptic weights, which
486 is computationally advantageous (Fortney and Tweed, 2012). Unraveling the latency of the model
487 updates will likely set physiological constraints on the candidate neural underpinnings.

488 Finally, our study also opens questions from a theoretical standpoint. Indeed, we must recall
489 that our model was excessively simple: we considered the translation of a point-mass in the plane
490 with full state information (see Methods and Appendix). We thus expect that future modeling work
491 unveil the computational challenges that arise when considering more realistic models of the neuro-
492 musculoskeletal system, including nonlinear dynamics, sensorimotor noise, and temporal delays.

493

494 # Methods

495 **Experimental Procedures**

496 Healthy volunteers (n=24) provided informed written consent and the experimental
497 procedures were approved by the ethics committee at the host institution (UCLouvain). Participants
498 grabbed the handle of a robotic device (KINARM, BKIN Technologies, Kingston, Canada) and were
499 instructed to perform reaching movements towards a visual target. The movement onset was cued
500 by filling in the goal target (Fig. 1a), and feedback was provided about movement timing (good trials
501 were between 600ms and 800ms between the go signal and the moment when they entered the
502 goal target).

503

504 *Experiment 1.* The goal of this experiment was to characterize online control in the presence of
505 unexpected changes in reach dynamics. Participants (n= 14) performed baseline trials randomly
506 interleaved with orthogonal force field trials mapping the forward hand velocity onto a lateral force
507 ($F_x = \pm L\dot{y}$, $L = 13\ Nsm^{-1}$). Movements consisted of 15cm reaches in the forward direction. To
508 make the task less predictable, other trial types were also randomly interleaved including trials with
509 a via-point situated on the side of the straight reach path, and trials with a constant background load
510 applied to the hand ($F_x = \pm 4N$), with the same frequency as the force field trials. Participants
511 performed six blocks of 80 trials including 56 baseline trials and 4 force field trials per direction of the
512 force field and per block, summing to a total of 24 force field trials per direction and participant.

513

514 *Experiment 2.* This experiment was designed to address whether the improvement in online
515 correction observed in the first experiment could evoke a near instantaneous after effect.
516 Movements consisted of 16cm reaches in the forward direction, and a via-point (radius: 1cm) was
517 located at 10cm on the straight line joining the start and goal target. Participants (n=10) performed
518 six blocks of 80 trials composed of 45 baseline trials, 5 force field trials, 5 baseline trials with a via-
519 point, 5 force field trials with a via-point, and 5 force field trials with a via-point and with a switching

520 off of the force field. All trials and force field directions were randomly interleaved. Participants were
521 given scores for good trials (1 point), and bonuses when they stopped at the via-point (3 points). The
522 bonus was awarded if the hand velocity inside the via-point dropped below 3cm/s. The experimental
523 setup monitored the hand velocity in real time allowing to turn off the force field if (1) the hand
524 cursor was in the via-point and (2) the hand velocity dropped below 3cm/s while in the via-point
525 (determined based on preliminary testing). Feedback about a successful stopping at the via-point
526 was given online. This task was difficult and we recorded an average of 21, 25 and 17 trials for
527 counterclockwise, baseline and clockwise via-point trials meeting these two conditions, respectively
528 (range across subjects: [5, 30], [18, 29], [2, 30]).
529
530 **Analysis and Statistical Design**
531 Position and force signals were sampled at 1kHz. Velocity signals were derived from
532 numerical differentiation (2nd order centered finite difference algorithm). The presence of significant
533 main effects was assessed with repeated-measures ANOVAs performed on individuals' mean value of
534 the variable of interest (rmANOVA). Post-hoc tests were performed based on paired t-test with
535 Bonferroni correction for multiple comparisons. In experiment 1, one participant missed a force-field
536 trial such that the analyses were performed on 23 trials for the corresponding force field direction.
537 The exponentials were fitted in the least-square sense and the significance of the fit was determined
538 based on whether the 99.9% confidence interval of the exponent responsible for the curvature of the
539 fit included or not the value of zero for Experiment 1 (P < 0.001), and 99.5% confidence for
540 Experiment 2 (P<0.005). The difference was likely due the number of participants in each experiment
541 but the level of 0.005 remains sufficiently low to limit the risk of false positive. The non-significant
542 fits were associated with P > 0.05. An R^2 statistics was derived as follows:
543 $$R^2 = 1 - \frac{V_{RES}}{V_{TOT}}, \quad (1)$$
544 where V_{RES} denotes the variance of the residuals and V_{TOT} is the variance of the data. We adopted
545 Cohen's d to quantify the effect size for paired comparison as the mean difference between two

populations divided by the standard deviation of the paired difference between the populations (Lakens, 2013).

The data of experiment 2 was also analyzed in more detail to extract information at single trial level. We used a mixed linear model of lateral hand velocity as a function of the lateral hand velocity prior to the via-point, and also as a function of the trial number based on standard techniques (Laird and Ware, 1982). We changed the sign of the lateral velocity for clockwise force field trials such that for all trials, a positive modulation reflected the presence of an after effect. The dependent variable was the lateral hand velocity measured at the moment of the second velocity peak, the fixed factor was the trial index, and the random effect was the participant.

Model

The model is based on the translation of a point-mass close to the mass of an arm ($m = 2.5$ kg) in the horizontal plane. The state-space representation is as follows:

$$m\ddot{x} = -k\dot{x} + F_x + F_{ENV}, \quad (2)$$

$$m\ddot{y} = -k\dot{y} + F_y, \quad (3)$$

$$\tau\dot{F}_x = u_x - F_x, \quad (4)$$

$$\tau\dot{F}_y = u_y - F_y, \quad (5)$$

where k is a dissipative constant (0.1 Nsm^{-1}), Equations 4 and 5 capture the first order muscle dynamics with time constant set to 0.1s, and F_{ENV} represents the unknown environmental perturbation: $F_{ENV} = 0$ for the null field, or $F_{ENV} = \pm L\dot{y}$ for force field trials.

This system of differential equations can be transformed in a discrete-time state space representation allowing us to add stochastic disturbances. Let x^* and y^* denote the coordinates of the goal target relative to the start, $\boldsymbol{z}_t := [x_t,\ y_t,\ \dot{x}_t,\ \dot{y}_t,\ F_{x,t},\ F_{y,t},\ x^*,\ y^*]^T$ represents the state vector at the time step t, and $\boldsymbol{u}_t := [u_{x,t}, u_{y,t}]^T$ represent the two-dimensional control vector, then the system can be written as follows:

$$\boldsymbol{z}_{t+1} = \left(\hat{A}_t + \Delta A\right)\boldsymbol{z}_t + B\boldsymbol{u}_t + \boldsymbol{\varepsilon}_t, \quad (6)$$

where $\hat{A}_t$ is the current (time varying) expected dynamics, ΔA is the unknown model error containing the unmodeled environmental disturbance (F_{ENV}), and $\boldsymbol{\varepsilon}_t$ is a Gaussian disturbance with zero-mean and covariance $\Sigma_\varepsilon \coloneqq BB^T$. We denote the true unforced dynamics by the matrix $A \coloneqq \hat{A}_t + \Delta A$. Details about these matrices can be found in the Appendix.

We assumed for simplicity that the controller has perfect state measurement, meaning that $\boldsymbol{z}_t$ is known (fully observable case). The problem is that the controller does not know ΔA. This problem was approached in the context of sensorimotor control assuming that ΔA followed a known (Gaussian) distribution (Izawa et al., 2008), which induces state-dependent noise. Observe that this assumption implies that the force field intensity varies randomly over time, which is not consistent with our experiment for which the force field varied randomly across trials, but remained fixed within trials. In other words, stochastic optimal control does not apply in the context of this paper because ΔA depends on F_{ENV}, which is not a stochastic variable but instead fixed bias or error in the model. Thus the initial controller is derived assuming that the system dynamics corresponds to A. To further simplify the problem, we assume that the controller knows which parameter is unknown (L). Let $\hat{A}_t$ denote the expected dynamics at time t dependent on the current estimate of the force field intensity noted $\hat{L}_t$, a Linear-Quadratic-Gaussian controller can be derived giving a time series of control signals based on linear feedback gains that depend on $\hat{A}$:

$$\boldsymbol{u}_t = -C_t\left(\hat{A}_t\right)\boldsymbol{z}_t, \qquad (7)$$

where $-C_t(\hat{A}_t)$ is the time series of feedback gains that, when applied to the state vector, defines the optimal control policy (Todorov, 2005).

Now, under this assumption, we may predict the next state vector, which we designate by $\hat{\boldsymbol{z}}_t$. When the true dynamics is distinct from the expected dynamics ($\Delta A = A - \hat{A}_t \neq 0$), there is a prediction error, $\boldsymbol{e}_t \coloneqq \boldsymbol{z}_t - \hat{\boldsymbol{z}}_t$, which can be used to correct the estimate of the model. We follow standard Recursive Least Square identification techniques (RLS), and use the following adaptive rule to update the estimate of L (Bitmead et al., 1990):

$$\hat{L}_{t+1} = \hat{L}_t + \gamma \frac{\partial \hat{\boldsymbol{z}}_{\boldsymbol{t}}}{\partial L} \boldsymbol{e}_t, \tag{8}$$

where γ is the online learning rate. This new value of estimated force field intensity is then used to update $\hat{A}_{t+1}$, and a new series of feedback gains is applied to the next state as control law: $\boldsymbol{u}_{t+1} = -C_{t+1}(\hat{A}_{t+1})\boldsymbol{z}_{t+1}$. In all, the closed loop controller consists in iteratively deriving optimal control gains based on an adaptive representation of the unknown parameters estimated online using Eqn 8. It can be shown in theory that under reasonable assumptions, this adaptive optimal control scheme robustly stabilizes the closed loop system (Bitmead et al., 1990).

The free parameters of the model are the cost-function used to derive the LQG controller, and the online learning rate γ. We fixed the cost-function to generate smooth movements with bell-shape velocity profiles, qualitatively comparable with human movements, and allowed γ to vary in order to capture the greater online adjustments observed across trials in Experiment 1 (see Appendix for a more detailed description). The simulation of a learning paradigm (Fig. 6b) was performed by letting $\hat{A}_t$ evolve during movements and storing the value obtained at the end of the trial for the beginning of the next trial ($\gamma = 0.05$). The reduction in peak end-force observed in Experiment 1 was simulated by using a range of values of γ to show that the range observed experimentally could be explained by this model ($\gamma = 0.1$, 0.25, and 0.5). Finally, Experiment 2 was simulated as a series of null field trials interleaved with force-field trials with switching-off at the via-point. Two values of γ were chosen to illustrate that the modulation of lateral velocity after the via-point increased with this parameter ($\gamma = 0.2$ and 0.5). The simulated baseline trials in the simulation of Experiment 2 were preceded by three baseline trials to induce washout in the model. We extracted the local minimum of lateral velocity and the lateral velocity at the moment of the second peak velocity to illustrate the after effect in the simulations.

Appendix

Definition of System Matrices and Cost-function

This section provides information about the algebraic representation of the system model used for the simulations. The system of differential equations (Eqns. 2-5) was transformed into the algebraic form of Eqn. 6 by using the following definition (the dimension corresponding to the target were omitted to simplify the matrices):

$$A_0 = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -k/m & \theta/m & 1/m & 0 \\ 0 & 0 & 0 & -k/m & 0 & 1/m \\ 0 & 0 & 0 & 0 & -\frac{1}{\tau} & 0 \\ 0 & 0 & 0 & 0 & 0 & -\frac{1}{\tau} \end{bmatrix}, \text{ and } B_0 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ \frac{1}{\tau} & 0 \\ 0 & \frac{1}{\tau} \end{bmatrix}.$$

The parameter θ is either 0 for baseline trials, or $\pm L$ for force field trials, in agreement with the definition of F_{ENV}. Then, the system must be expressed in discrete time, which is achieved by using a first order Taylor expansion over one time step of δt: $A := I + \delta t A_0$, and $B := \delta t B_0$ (I is the identity matrix). We use in the simulation a discretization step of $\delta t = 0.01s$. The model error, ΔA, comes from the unknown value of L, representing the force field direction, at the beginning of each trial.

We use the following cost function to simulate reaching movements without via-point:

$$J(\boldsymbol{z}, \boldsymbol{u}) = (t/N)^{25}(w_1[(x_t - x^*)^2 + (y_t - y^*)^2] + w_2[\dot{x}_t^2 + \dot{y}_t^2]) + \sum_{t=1}^{N-1} 10^{-5}\|\boldsymbol{u}_t\|$$

with $N = 61$, consistent with the average movement duration from Experiment 1. The parameters w_1 and w_2 were set to 1000 and 20 to reproduce smooth movements with bell-shape velocity profiles. The first term produces a polynomial buildup of the terminal cost with value ~0 until about 80% of the terminal step, and a smooth buildup of the terminal cost for the last ~20% of the reach time. For the via-point experiment, the cost function was identical except that we used $N = 100$, and we added the penalty that the system had to be close to the via-point target at $t = 60$ with the same cost as for the last step. There was no quantitative fitting of parameter values to the data.

Funding and Open Data Statement

FC is supported by a grant from the Belgian Fund for Scientific Research (F.R.S.-FNRS, *Chercheur Qualifié*, grant number 1.C.033.18F). This work was supported by grants from PRODEX and by the Interuniversity Attraction Pole of the Belgian Science Policy (BELSPO). The data from each experiment was uploaded to Dryad (doi:10.5061/dryad.53vm1m2).

References

Astrom, K.J. (1970). Introduction to stochastic control theory (New York: Academic Press).

Bitmead, R.R., Gevers, M., and Wertz, V. (1990). Adaptive optimal control : the thinking man's GPC (New York: Prentice Hall).

Braun, D.A., Aertsen, A., Wolpert, D.M., and Mehring, C. (2009). Learning Optimal Adaptation Strategies in Unpredictable Motor Tasks. Journal of Neuroscience 29, 6472-6478.

Burdet, E., Osu, R., Franklin, D.W., Milner, T.E., and Kawato, M. (2001). The central nervous system stabilizes unstable dynamics by learning optimal impedance. Nature 414, 446-449.

Burdet, E., Osu, R., Franklin, D.W., Yoshioka, T., Milner, T.E., and Kawato, M. (2000). A method for measuring endpoint stiffness during multi-joint arm movements. Journal of Biomechanics 33, 1705-1709.

Cluff, Y., and Scott, S.H. (2013). Rapid feedback responses correlate with reach adaptation and properties of novel upper limb loads. Journal of Neuroscience 33, 15903-15914.

Crevecoeur, F., Kurtzer, I., and Scott, S.H. (2012). Fast corrective responses are evoked by perturbations approaching the natural variability of posture and movement tasks. Journal of Neurophysiology 107, 2821-2832.

Crevecoeur, F., and Scott, S.H. (2013). Priors Engaged in Long-Latency Responses to Mechanical Perturbations Suggest a Rapid Update in State Estimation. Plos Computational Biology 9.

Crevecoeur, F., and Scott, S.H. (2014). Beyond Muscles Stiffness: Importance of State Estimation to Account for Very Fast Motor Corrections. PLoS Computational Biology 10, e1003869.

Dayan, E., and Cohen, L.G. (2011). Neuroplasticity subserving motor skill learning. Neuron 72, 443-454.

Diedrichsen, J., White, O., Newman, D., and Lally, N. (2010). Use-dependent and error-based learning of motor behaviors. J Neurosci 30, 5159-5166.

Fortney, K., and Tweed, D.B. (2012). Computational Advantages of Reverberating Loops for Sensorimotor Learning. Neural Computation 24, 611-634.

Franklin, D.W., and Wolpert, D.M. (2011). Computational Mechanisms of Sensorimotor Control. Neuron 72, 425-442.

Ghahramani, Z., and Hinton, G.E. (1996). Parameter estimation for linear dynamical systems. Technical Report CRG-TR-96-2, University of Toronto, Dept of Computer Sceince.

Gonzalez Castro, L.N., Hadjiosif, A.M., Hemphill, M.A., and Smith, M.A. (2014). Environmental consistency determines the rate of motor adaptation. Curr Biol 24, 1050-1061.

Hirashima, M., and Nozaki, D. (2012). Distinct Motor Plans Form and Retrieve Distinct Motor Memories for Physically Identical Movements. Current Biology 22, 432-436.

Hore, J., and Flament, D. (1988). Changes in motor cortex neural discharge associated with the development of cerebellar limb ataxia. J Neurophysiol 60, 1285-1302.

Howard, I.S., Ingram, J.N., Franklin, D.W., and Wolpert, D.M. (2012). Gone in 0.6 seconds: the encoding of motor memories depends on recent sensorimotor states. J Neurosci 32, 12756-12768.

Ioannou, P., and J, S. (1996). Robust Adaptive Control (Prentice Hall Inc.).

Izawa, J., Rane, T., Donchin, O., and Shadmehr, R. (2008). Motor adaptation as a process of reoptimization. Journal of Neuroscience 28, 2883-2891.

Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. Nat Neurosci 17, 440-448.

Kording, K.P., Tenenbaum, J.B., and Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. Nat Neurosci 10, 779-786.

Krakauer, J.W., Ghilardi, M.F., and Ghez, C. (1999). Independent learning of internal models for kinematic and dynamic control of reaching. Nature Neuroscience 2, 1026-1031.

Kurtzer, I.L., Pruszynski, J.A., and Scott, S.H. (2008). Long-latency reflexes of the human arm reflect an internal model of limb dynamics. Current Biology 18, 449-453.

Lackner, J.R., and DiZio, P. (1994). Rapid adaptation to Coriolis-force perturbations of arm trajectory. Journal of Neurophysiology 72, 299-313.

Lackner, J.R., and DiZio, P. (2005). Motor control and learning in altered dynamic environments. Current Opinion in Neurobiology 15, 653-659.

Laird, N.M., and Ware, J.H. (1982). Random-effects models for longitudinal data. Biometrics 38, 963-974.

Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. Front Psychol 4, 863.

Omrani, M., Murnaghan, C.D., Pruszynski, J.A., and Scott, S.H. (2016). Distributed task-specific processing of somatosensory feedback for voluntary motor control. Elife 5.

Pruszynski, J.A., Kurtzer, I., Nashed, J.Y., Omrani, M., Brouwer, B., and Scott, S.H. (2011). Primary motor cortex underlies multi-joint integration for fast feedback control. Nature 478, 387-390.

Schaal, S., Sternad, D., Osu, R., and Kawato, M. (2004). Rhythmic arm movement is not discrete. Nat Neurosci 7, 1136-1143.

Scheidt, R.A., Dingwell, J.B., and Mussa-Ivaldi, F.A. (2001). Learning to move amid uncertainty. Journal of Neurophysiology 86, 971-985.

Scott, S.H. (2016). A Functional Taxonomy of Bottom-Up Sensory Feedback Processing for Motor Actions. Trends Neurosci 39, 512-526.

Shadmehr, R., and Mussa-Ivaldi, F.A. (1994). Adaptive representation of dynamics during learning of a motor task. Journal of Neuroscience 14, 3208-3224.

Shadmehr, R., Smith, M.A., and Krakauer, J.W. (2010). Error Correction, Sensory Prediction, and Adaptation in Motor Control. Annual Review of Neuroscience, Vol 33 33, 89-108

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

Singh, K., and Scott, S.H. (2003). A motor learning strategy reflects neural circuitry for limb control. Nature Neuroscience 6, 399-403.

Smith, M.A., Ghazizadeh, A., and Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. Plos Biology 4, 1035-1043.

Stavisky, S.D., Kao, J.C., Ryu, S.I., and Shenoy, K.V. (2017). Motor Cortical Visuomotor Feedback Activity Is Initially Isolated from Downstream Targets in Output-Null Neural State Space Dimensions. Neuron 95, 195-208 e199.

Thoroughman, K.A., and Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. Nature 407, 742-747.

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

Wagner, M.J., and Smith, M.A. (2008). Shared Internal Models for Feedforward and Feedback Control. Journal of Neuroscience 28, 10663-10673.

Wei, K., and Koerding, K. (2010). Uncertainty of feedback and state estimation determines the speed of motor adaptation. Frontiers in Computational Neuroscience 4.

Wolpert, D.M., Diedrichsen, J., and Flanagan, J.R. (2011). Principles of sensorimotor learning. Nature Reviews Neuroscience 12, 739-751.