

Contralateral manual compensation for velocity-dependent force perturbations

Jackson, Carl; Miall, Rowland

DOI:
[10.1007/s00221-007-1179-6](https://doi.org/10.1007/s00221-007-1179-6)

Citation for published version (Harvard):

Jackson, C & Miall, R 2007, 'Contralateral manual compensation for velocity-dependent force perturbations', *Experimental Brain Research*, vol. 184, pp. 261-267. <https://doi.org/10.1007/s00221-007-1179-6>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Contralateral manual compensation for velocity-dependent force perturbations

Carl P. T. Jackson · R. Chris Miall

Received: 10 July 2007 / Accepted: 12 October 2007 / Published online: 1 November 2007
© Springer-Verlag 2007

Abstract It is not yet clear how the temporal structure of a voluntary action is coded allowing coordinated bimanual responses. This study focuses on the adaptation to and compensation for a force profile presented to one stationary arm which is proportional to the velocity of the other moving arm. We hypothesised that subjects would exhibit predictive coordinative responses which would co-vary with the state of the moving arm. Our null hypothesis is that they develop a time-dependent template of forces appropriate to compensate for the imposed perturbation. Subjects were trained to make 500 ms duration reaching movements with their dominant right arm to a visual target. A force generated with a robotic arm that was proportional to the velocity of the moving arm and perpendicular to movement direction acted on their stationary left hand, either at the same time as the movement or delayed by 250 or 500 ms. Subjects rapidly learnt to minimise the final end-point error. In the delay conditions, the left hand moved in advance of the onset of the perturbing force. In test conditions with faster or slower movement of the right hand, the predictive actions of the left hand co-varied with movement speed. Compensation for movement-related forces appeared to be predictive but not based on an accurate force profile that was equal and opposite to the imposed perturbation.

Keywords Human · Motor control · Prediction · Temporal delay · Bimanual coordination

Introduction

Many studies have investigated motor adaptation to velocity-dependent forces during unimanual point-to-point reaching movements (e.g. Scheidt et al. 2005; Shadmehr and Mussa-Ivaldi 1994). Adaptation is thought to involve the training of an internal model which predicts the sensory consequences of actions (Miall and Wolpert 1996), although bimanual coordination during and after adaptation is less well understood. In a recent study it has been shown that subjects can learn to adapt to a force on one arm which is proportional to the velocity of the other arm (Bays and Wolpert 2006). That study examined the generalisation of learning to a new joint configuration of the arm experiencing the force, but did not touch on the temporal relationships between the two arm actions. In contrast, Diedrichsen et al. (2007) recently demonstrated a dissociation between state- and time-dependent control. They used a task where subjects made reaching movements with one hand and were instructed to press a button with the other hand at some point during the movement. When the two movement components overlapped, temporal scaling of the arm movement appeared to result in temporal scaling of the button press (state-dependent control); when the movements were separated by a time delay, this scaling did not occur (time-dependent control). However, this work focuses on a discrete button-press and not a compensatory response to an imposed force.

Experiments on grip force provide evidence of predictive responses. Witney et al. (1999) demonstrated that when subjects hold a virtual object in one hand which is acted on by the other after a temporal delay, the reactive grip force that subjects initially employ to prevent the object from slipping decays over time and is replaced with an anticipatory grip force prior to the onset of the load

C. P. T. Jackson (✉) · R. C. Miall
Behavioural Brain Sciences, School of Psychology,
University of Birmingham, Edgbaston,
Birmingham B15 2TT, UK
e-mail: c.p.jackson.1@bham.ac.uk

force. When the load force was not generated by the subject's own action, no predictive grip force was found even when the motion was cued by a tone. In the same way, subjects have been found to exhibit predictive grip forces when exposed to a self-generated sinusoidal load force but not when a similar sinusoidal load force is imposed externally, even though this pattern is predictable and easy to learn (Blakemore et al. 1998). These results demonstrate the importance of efference copy in the production of anticipatory responses.

In a similar vein, Blakemore and colleagues conducted a series of bimanual experiments on sensory cancellation of self-produced tickle sensations (e.g. Blakemore et al. 2000). They found that when subjects used a robot arm controlled by one hand to tickle themselves on the other hand, their subjective sensations of tickliness decreased during synchronous movements and increased as the delay between the action of the hand controlling the arm and the onset of the tickle increased (Blakemore et al. 1999). These findings can also be interpreted in the context of a sensorimotor system which attenuates synchronous sensation based on action prediction.

Thus, there is evidence of sensory prediction based on continuous voluntary action, and evidence of anticipatory responses in coordination of load and grip forces, but whether these have time- or state-dependency is not clear. We wanted to investigate whether varying properties of the force imposed on the resisting arm in both a state-dependent and time-dependent fashion caused comparable changes in the behaviour of the resisting arm. We hypothesised that subjects would exhibit predictive coordinative responses to delayed force feedback in a similar manner to those displayed during the aforementioned grip force studies: after learning to compensate for the force, they should start to make responses before the force is incident on the arm. Furthermore, we expected these responses to co-vary with the state of the arm but not with the timing of the force, because of their origin in the forward model predictions of the arm.

Materials and methods

Participants

Six normal healthy right-handed adults (three male, three female) participated in this study after giving their informed consent. Their ages ranged from 20 to 35 years with a mean age of 24.5 years. None of the subjects had any known motor or sensory abnormalities, and all had normal or corrected-to-normal vision. The experiment was approved by the local ethics committee.

Apparatus

Subjects used both hands on a symmetric pair of robotic manipulanda (vBOTs, cf, Bays and Wolpert 2006) to control the movement of cursors on a flat screen monitor positioned 1 m in front of them. The apparatus was connected to a computer that displayed the stimuli on the screen and that also controlled and received data from the robots. Sensors in the motors of the robots allowed the position and velocity of the manipulanda to be sampled at 1,000 Hz, and forces imposed by the robots were also updated at this frequency.

Procedure

At the start of each experimental trial, subjects were presented with a 5 mm red circular cursor indicating the position of their right hand [moving hand (MH)], and a green circular cursor indicating the position of their left hand [opposing hand (OH)]. They were asked to move the red cursor on to a 5 mm white circle in the bottom centre of the screen and the green cursor into a 20 mm white circle in the bottom left of the screen to start the trial. Once these conditions were met, both the white circles and the green cursor disappeared and a 20 mm blue target circle appeared in the top centre of the screen. Subjects were instructed to move the red cursor to the blue target within a time window of 400–600 ms from movement onset. If these timing conditions were met, then the target turned yellow to indicate appropriate movement duration. If the subjects moved too slowly, the target turned green; if the subjects moved too quickly, the target turned red. While subjects moved the MH to the target, the OH experienced a leftward force proportional to the forward velocity of the MH movement that had its onset delayed either by 0, 250 or 500 ms. These delays were chosen to ensure that, with movements of ~500 ms duration, there was complete, partial or no overlap between the MH and the forces on the OH. Subjects were instructed to keep their left hand as still as possible.

The delays were experienced in three randomised blocks of 250 trials each. For the first 150 trials, subjects were trained with the selected delay to adapt them to the perturbation. The next 20 trials were 'test' trials at the same speed but without colour feedback of movement duration for comparison to future speed-modified test trials. These were followed by 20 'top-up' trials in which the target again changed colour to make sure that subjects were moving at the correct speed. 20 more test trials followed, and this time subjects were either asked to move more quickly or more slowly. Subjects were instructed to move "as fast as possible, about twice the speed" or "as slow as possible, about half the speed" for all 20 movements. They then received

20 more top-up trials to return them to the original speed, and then 20 final test trials with the speed modified in the opposite way to the previous test trials. During top-up and test trials in each delay condition, forces were applied with the same velocity dependence as in the training condition and at the same time. The only difference between top-up and test trials was the lack of a colour change in the target in the latter.

Data analysis

All data analysis took place in Matlab. Since subjects' instructions were to keep the OH stationary during the trial, we defined 'error' as any deviation from the initial hand position. The root mean square error and end-point error of the OH were therefore defined as the RMS distance moved by the hand during the trial and the difference between the hand position at the start and the end of the trial, respectively. The *x*-direction was defined as the left-right horizontal plane, and the *y*-direction was defined as the forward-back horizontal plane.

Pilot experiment

Prior to the current experiment, 12 subjects participated in a similar study where they moved a cursor to a target with their MH and received a velocity-dependent perturbation on their OH, delayed by either 0 or 250 ms, blocked and counterbalanced across subjects. They first adapted to the perturbation for 50 trials, and then took part in three counterbalanced blocks at slow, medium and fast speeds. The first 20 trials of each block contained no perturbation and

were included to allow subjects to train at a particular speed. During these training trials subjects received colour feedback about movement duration as described above. The remaining 50 trials were test trials at the same speed as the preceding training trials, this time with a perturbation as well as the colour feedback. Five no-force 'catch' trials were presented pseudorandomly during each block of these test trials. In the main experiment we removed the colour feedback during test trials to reduce subjects' cognitive load and also removed the catch trials so as not to interfere with trial-to-trial performance (Donchin et al. 2003). However, the catch trial data from the pilot experiment proved to be a useful comparison to the data from the current experiment.

Results

By the final 35 trials of the training sequence, subjects had adapted to the perturbation such that the slope of the group mean RMS error in the *x*-direction over these trials in each delay condition (Fig. 1a) was not significantly different from zero (Pearson's *r*, 0 ms: $r = -0.101$, $P = 0.562$; 250 ms: $r = -0.237$, $P = 0.171$; 500 ms: $r = 0.066$, $P = 0.706$). Within-subject analysis in each condition found that only two slopes were significantly different from zero. As both were small and positive, we assume this indicates fatigue rather than adaptation to the perturbation. The mean velocity in the *x*-direction of the OH (solid lines) and the *y*-direction of the MH (dashed lines) over the 20 test trials at the medium speed for each delay condition are plotted in Fig. 1b. The MH velocity is shown time-delayed to illustrate




Fig. 1 Movement data. **a** RMS error of the OH over training sequence for 500 ms delay condition, averaged across all subjects. An exponential fit to the data is shown to demonstrate learning. By the final 35 trials adaptation was complete, as the slope of the RMS error across these trials was not significantly different from zero. **b** Mean velocities across subjects for the test trials at the medium speed for each of the delay conditions (black curves: 0 ms; blue curves: 250 ms; red curves: 500 ms). Velocity curves for the MH (dashed lines) are bell-shaped and time-shifted for each delay condition to illustrate the timing of the forces

incident on the OH. Predictive responses can be seen for the 250 and 500 ms delay conditions for the OH (solid lines) as positive-going peaks before the onset of the MH velocity and these are followed by a deflection and then an overshoot. **c** Mean position data for the OH (solid lines) and MH (dashed lines) across subjects for the test trials at the medium speed for each of the delay conditions. Position curves for the MH are time-delayed. The OH is always pushed back (negative) by the force before returning to the original zero position. Anticipatory peaks are difficult to make out in this case but examples can be seen in Fig. 2a

the impact of the velocity-dependent force on the OH at each time delay. The MH velocity curves exhibit the usual bell-shaped profile characteristic of point-to-point reaching movements. The OH curves for the 250 and 500 ms delay conditions show a small, positive velocity peak prior to the peak of the imposed force. This peak represents anticipatory rightward motion of the OH. Anticipatory responses in the 250 and 500 ms delay conditions were also observed for the slow and fast speed-modified test trials (not shown). In all delay conditions the OH is pushed back by the force before returning to the original position, as illustrated in Fig. 1c; similar predictive peaks are in evidence for the delay conditions in the position data. No evidence of anticipatory responses was seen in the synchronous (zero delay) condition.

One aim of the experiment was to investigate whether the predictive responses co-varied with the state of the moving arm, so to investigate how the movement of the OH changed with MH speed and with delay a 3×2 repeated-measures ANOVA (speed: slow, medium and fast; delay:

250 and 500 ms) was performed on the amplitude of the predictive position peak of the OH. The zero delay condition was not included because no predictive peaks were evident in this condition. We found that the amplitude of the predictive position peaks increased significantly with both speed [$F(2,10) = 14.214$, $P = 0.001$] and delay [$F(1,5) = 39.729$, $P = 0.014$]. There was also a trend towards an interaction between speed and delay [$F(2,10) = 3.876$, $P = 0.057$]. These results indicate that the magnitude of the anticipatory movement by the OH increased as speed and delay increased (Fig. 2a, b). Data from the catch trials in the pilot experiment show similar behaviour in the 0 ms condition when no perturbing force was incident on the OH (Fig. 2c, d).

The movement strategy subjects used during the task was investigated by examining both the RMS error and the end-point error (displacement between the start and end of each trial) in the OH for the last 35 trials of the training period for each delay condition (Table 1). Repeated-measures ANOVAs demonstrated that there was no significant




Fig. 2 Evidence against time-dependent control. **a** Position data for the OH averaged over all subjects for the first trial at the slow (*black*), medium (*blue*) and fast (*red*) speeds in the 250 ms delay condition. As the speed increases, the magnitude of the predictive peak also increases, but the onset time is not affected. **b** Average predictive position peak magnitude (and standard error) in the 250 and 500 ms delay conditions for all three movement speeds. **c** Velocity data for the OH averaged over all catch trials and all subjects in the pilot experiment at the

slow (*black*), medium (*blue*) and fast (*red*) speeds in the 250 ms delay condition. Again, the onset time is not affected by the increase in speed. The **bold lines** up to maximum velocity show that the traces are scaled versions of each other before corrective movements are made (from maximum velocity onwards; fine lines). **d** Average predictive position peak magnitude (and standard error) in the 0 and 250 ms delay conditions for all three movement speeds in the catch trials of the pilot experiment

Table 1 Means (and standard deviations) of end-point and RMS error for the last 35 trials of the training period for each delay condition

Delay	End-point error	RMS error
0 ms	0.29 (0.41) cm	1.04 (0.40) cm
250 ms	0.12 (0.46) cm	1.15 (0.27) cm
500 ms	−0.00 (0.62) cm	1.57 (0.45) cm

difference in the end-point error between delay conditions [$F(2,10) = 1.762$, $P = 0.221$] but that RMS error increased significantly with delay [$F(2,10) = 5.518$, $P = 0.024$]. Post hoc tests revealed that the 500 ms delay condition had a significantly larger RMS error than the 0 ms condition, though the 0 and 250 ms conditions and the 250 and 500 ms conditions did not differ significantly. Furthermore, two-tailed one-sample t -tests showed that the end-point error was not significantly different from zero for any of the delay conditions [0 ms: $t(5) = 1.721$, $P = 0.146$; 250 ms: $t(5) = -0.005$, $P = 0.997$; 500 ms: $t(5) = 0.472$, $P = 0.657$]. These results demonstrate that subjects were not able to keep their OH still as instructed during the experiment, but instead suggest that they were acting to return to the original position. The increased RMS error for the 500 ms delay condition indicates that subjects found it more difficult to stabilise their hand with a higher delay.

Discussion

During the experiment we found that subjects were largely successful at opposing the force on their left hand generated by the movement of their right hand whether it was delayed or not. There are several ways they could have compensated for the imposed force, and by testing their responses at two different speeds we can distinguish between these methods.

First, they could simply have increased the mechanical impedance of the opposing arm by stiffening the muscles. This has been shown to be at least a partially effective strategy for perturbation tasks (Franklin et al. 2003) but it was not the case here: the existence of the predictive peaks for the delay (250, 500 ms) conditions showed that subjects were actively moving to oppose the force before it was incident rather than just stiffening up and waiting until it impacted on the arm. Pilot data using catch trials on which there was no perturbation showed that subjects were indeed moving their OH at expected incidence of the force in the synchronous (0 ms) condition. This behaviour demonstrates that the subjects were not exclusively increasing stiffness.

Second, subjects could have learnt the timing and magnitude of the force perturbation as a ‘template’ during the

training phase and produced a force profile that was equal and opposite to the imposed force. They would then be able to reproduce this template with the OH at a fixed time after the onset of the MH. While this would be an effective strategy for those cases where the speed does not change, it does not appear to be the strategy employed in this case. If subjects were simply ‘playing back’ a learnt force profile then we would not expect to see any change in the motion of the OH before the onset of the imposed force. In other words we should see equivalent predictive OH movements for all three test speeds, and it is clear that we do not (Fig. 2b); similar behaviour is observed in the catch trials (Fig. 2c, d). The magnitude of the anticipatory position peak (or, for the pilot catch trials in Fig. 2d, the magnitude of the initial peak) increases for the faster speed and decreases for the slower speed, even in the first trial of each test block (Fig. 2a).

Perhaps then something more sophisticated is going on? If the movement was based on a time-dependent template of the force, subjects should have scaled the OH template in time as they increased or decreased the movement duration of the MH (for the faster or slower speed). However, this behaviour would result in early predictive peaks for the fast delay conditions and later peaks for the slow delay conditions. While very early movements were occasionally seen they appeared across all conditions and were most likely associated with an erroneous simultaneous movement of both arms—usually these brief movements of the OH were quickly aborted and later predictive peaks were observed just prior to the delayed onset of the force. These early peaks appeared only rarely and did not vary systematically with the speed of the movement. Thus, we see no evidence for temporal scaling of a response profile.

A more comprehensive explanation for the anticipatory responses can be found in the context of a state-dependent representation (Diedrichsen et al. 2007). State-dependent control is based on the idea of a predictive forward internal model (Miall and Wolpert 1996) which predicts the sensory consequences of a descending motor command to the MH (Witney et al. 2000). Subjects could then use this information to generate a compensatory response in the OH (Wolpert and Flanagan 2001). The anticipatory responses were dependent on the velocity of the MH when there was a delay between its motion and the force onset. Moreover, the amplitude of these predictive responses scaled with the speed of the MH (Conditt and Mussa-Ivaldi 1999), but their timing did not. This suggests that the responses were related to the state of the MH, rather than to the time of the MH’s action. A state-dependent account of control in this case predicts this scaling as the motion of one hand is dependent on the estimated state of the other hand.

Are aspects of time-dependent control combined with this state-dependent strategy? The fact that the onsets of the

anticipatory responses do not vary with increasing speed suggests that the control exhibited may not be entirely state-dependent. It is possible that subjects learnt to make their responses in a state-dependent fashion to the changing velocity of the MH, but also learnt a time-dependent representation of *when* to make these responses with the OH. While internal models in the motor system do not seem to be time-dependent during the course of an action (Conditt and Mussa-Ivaldi 1999), the initiation of a state-dependent response may well be able to be learnt in a time-dependent fashion. This interpretation ties in well with evidence from functional neuroimaging that timing processes and sensory-motor processes are separate (Sakai et al. 2000). Furthermore, the observation that subjects are less able to stabilise their hand as the delay increases (Table 1) suggests that as delay increases the uncertainty in the initiation of the response also increases, as would be predicted by the scalar timing hypothesis that variability of timing increases with increased interval (Gibbon et al. 1997).

The lack of predictive peaks in the movement of the OH prior to the onset of the force for the 0 ms delay condition is similar to the behaviour observed in grip force studies where grip force increases concurrently with load force (e.g. Witney et al. 1999), but in those studies the same behaviour was seen in the presence of a delay. In the present study, the anticipatory movements that appeared before force onset in the delay conditions differ from those observed in the grip force studies. This apparent contradiction may be resolved firstly by considering the uncertainty in when to initiate the motor command and secondly by looking at the differences between the grip force studies and the present study. When the movement of the OH is initiated at the same time as the movement of the MH, the motor commands for both actions can be sent simultaneously, minimising the uncertainty in when to make the OH movement. When there is a delay between movement of the MH and movement of the OH, the uncertainty increases so it may be better for the motor system to start the movement early.

It is possible that this pattern of behaviour is not seen in the grip force studies due to differences in the amount of disturbance caused by each task. Specifically, uncertainty in timing in delayed grip force tasks leads to slipping but this can be easily and quickly rectified by increasing grip force. In the present experiment, uncertainty in the timing can cause the arm to be pushed far out of the ‘comfort zone’ of acceptable responses, which may not be so easy to rectify. Motion of the OH before the force is incident on it leads to the idea of anticipatory responses as an efficient way of minimising the total disturbance to the OH. The final distance of the OH from its initial position did not significantly differ from zero over all speeds and delays. Rather than keeping their hands as

still as possible, subjects were apparently attempting to get back to the initial hand position as accurately as possible. It may be a more efficient strategy for the motor system to begin to make a compensatory movement, get pushed back and return to the start position than to keep the arm absolutely still.

In conclusion, we found that when opposing forces incident on one hand were proportional and delayed relative to the velocity of the other hand, subjects produced predictive movements that peaked just before force onset. These predictive peaks were modulated in a state-dependent fashion by the speed of the moving arm. We have therefore shown that state-dependent predictive control is not only used in discrete bimanual tasks but also in opposing continuous forces generated by movement of the opposite arm. Even when the force on one hand is not incident until after the end of the movement of the other hand (the 500 ms delay condition), this state-dependent control is still observed. Questions remain as to whether some aspects of time-dependent control are also combined with this state-dependent strategy, but will need additional experiments to resolve.

Acknowledgements We thank Jörn Diedrichsen and an anonymous reviewer for their helpful comments on the manuscript. This work was supported by the Wellcome Trust [GR069439].

References

- Bays PM, Wolpert DM (2006) Actions and consequences in bimanual interaction are represented in different coordinate systems. *J Neurosci* 26:7121–7126
- Blakemore SJ, Frith CD, Wolpert DM (1999) Spatio-temporal prediction modulates the perception of self-produced stimuli. *J Cogn Neurosci* 11:551–559
- Blakemore SJ, Goodbody SJ, Wolpert DM (1998) Predicting the consequences of our own actions: The role of sensorimotor context estimation. *J Neurosci* 18:7511–7518
- Blakemore SJ, Wolpert D, Frith C (2000) Why can't you tickle yourself? *Neuroreport* 11:R11–R16
- Conditt MA, Mussa-Ivaldi FA (1999) Central representation of time during motor learning. *Proc Natl Acad Sci USA* 96:11625–11630
- Donchin O, Francis JT, Shadmehr R (2003) Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J Neurosci* 23:9032–9045
- Diedrichsen J, Criscimagna-Hemminger SE, Shadmehr R (2007) Dissociating timing and coordination as functions of the cerebellum. *J Neurosci* 27:6291–6301
- Franklin DW, Burdet E, Osu R, Kawato M, Milner TE (2003) Functional significance of stiffness in adaptation of multijoint arm movements to stable and unstable dynamics. *Exp Brain Res* 151:145–157
- Gibbon J, Malapani C, Dale CL, Gallistel CR (1997) Toward a neurobiology of temporal cognition: advances and challenges. *Curr Opin Neurobiol* 7:170–184
- Miall RC, Wolpert DM (1996) Forward models for physiological motor control. *Neural Netw* 9:1265–1279

- Sakai K, Hikosaka O, Takino R, Miyauchi S, Nielsen M, Tamada T (2000) What and when: Parallel and convergent processing in motor control. *J Neurosci* 20:2691–2700
- Scheidt RA, Conditt MA, Secco EL, Mussa-Ivaldi FA (2005) Interaction of visual and proprioceptive feedback during adaptation of human reaching movements. *J Neurophysiol* 93:3200–3213
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224
- Witney AG, Goodbody SJ, Wolpert DM (1999) Predictive motor learning of temporal delays. *J Neurophysiol* 82:2039–2048
- Witney AG, Goodbody SJ, Wolpert DM (2000) Learning and decay of prediction in object manipulation. *J Neurophysiol* 84:334–343
- Wolpert DM, Flanagan JR (2001) Motor prediction. *Current Biology* 11:R729–R732