ARTICLE IN PRESS

Human Movement Science xxx (2009) xxx–xxx

Contents lists available at ScienceDirect

Human Movement Science

journal homepage: www.elsevier.com/locate/humov

Bimanual coordination as task-dependent linear control policies

Jörn Diedrichsen *, Noreen Dowling

Wolfson Centre for Cognitive and Clinical Neuroscience, School of Psychology, University of Wales, Adeilad Brigantia, Bangor Gwynedd LL57 2AS, UK

ARTICLE INFO

PsycINFO classification:
2330

Keywords:
Motor
Feedback control
Modelling
Bimanual

ABSTRACT

When we perform actions with two hands in everyday life, coordination has to change very quickly depending on task goals. Here, we study these task-dependent changes using a bimanual reaching task in which participants move two separate cursors to two visual targets, or move a single cursor, displayed at the average position of the two hands, to a single target. During the movement, one of the hands is perturbed in a random direction using a viscous curl field. We have previously shown that feedback control, the structure of noise, and adaptation change between these two tasks as predicted by optimal control theory: feedback control is independent when the hands control two cursors, but becomes dependent when they move one cursor together. The same changes are observed even on trials in which no visual feedback about the cursor position is given. One assumption in this model is that coordinative motor commands can be described as a linear function of the state of the left and right hands. Here we test the assumption by studying the feedback corrections for 25 combinations of force fields applied to the two hands. Our study shows that feedback gains are constant across all levels of force fields strength, providing strong evidence that intermanual coordination for this task can accurately be explained by optimal task-dependent linear feedback gains.

---

* Corresponding author. Tel.: +44 01248 383768; fax: +44 01248 382599.
E-mail address: j.diedrichsen@bangor.ac.uk (J. Diedrichsen).

0167-9457/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.humov.2008.10.003

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003

1. Introduction

One of the remarkable facts about bimanual coordination is that we can use our two hands to solve a wide variety of tasks. We tie our shoes, play musical instruments, applaud, open wine bottles, and wash the dishes. The coordination requirements between the two hands in these situations are quite different. During applauding, simple mirror symmetric movements are appropriate. When playing an instrument we have to coordinate the two hands in time, yet often make spatially highly asymmetric movements. When manipulating a single object with two hands, each hand needs to take into account the forces that are exerted by the other hand to maintain object stability. How can such flexibility in control be achieved? How does feedforward and feedback control change under different task constraints? We have recently suggested the application of optimal control theory to analyze such task-dependent changes (Diedrichsen, 2007).

Optimal control theory (Todorov, 2007; Todorov & Jordan, 2002) provides a general and flexible framework to understand motor control (Fig. 1). A plant (the body, the two hands, etc.) that is represented with a collection of state variables (vector x) can be controlled with a vector of motor commands (u). Some aspects of the state can be experienced through sensory feedback (y), however, only with a certain time delay and noise. Therefore, it is beneficial for the system to predict the state of the plant using a forward model (Wolpert, Ghahramani, & Jordan, 1995), and to integrate this prediction continuously with the sensory inflow to arrive at a good estimate of the system state. A control policy then computes the new motor commands based on such state estimate. In this framework, there is no qualitative difference between feedforward and feedback control. In both cases the control policy specifies the motor commands based on a state estimate, which can be determined to varying degree by sensory signals or internal predictions.

The system is optimal in two respects. First, it optimally combines new sensory information with a forward model prediction, weighting each of these information sources at every time step by the inverse of their variances (Vaziri, Diedrichsen, & Shadmehr, 2006), therefore acting like a Kalman filter (Wolpert et al., 1995). Optimality is here defined in terms of the quality of the estimator. Second, the control policy is defined to be optimal in respect to a certain cost function, which may include kinematic measure (Flash & Sejnowski, 2001), or endpoint variance (Harris & Wolpert, 1998). These cost functions are thought to be minimized in the course of evolution and/or individual learning history. The approach is in some respects non-falsifiable: for every behavior we can find a cost function, for which this behavior is optimal. However, the approach gains predictive power if we can explain complex behaviors as being the optimal solution for simple cost functions that reflect objective constraints

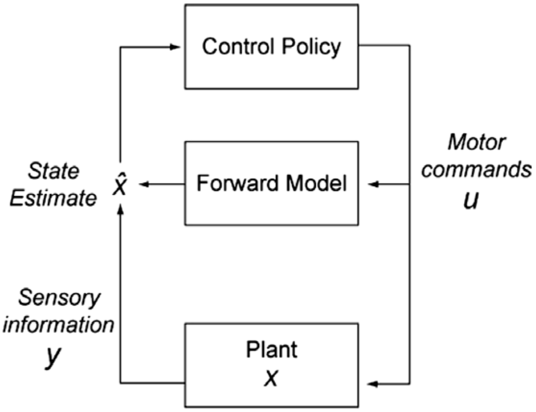

Fig. 1. Control architecture of optimal control theory. The controlled system is described by a vector of state variables $\mathbf{x}$. The brain receives noisy and delayed sensory information $\mathbf{y}$, which is integrated with the prediction from a forward model to arrive at the best concurrent estimate of the state of the plant $\hat{\mathbf{x}}$. The control policy is a rule of how to generate motor commands $\mathbf{u}$ based on this state estimate.

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003

of the body and the environment. In this case we can derive testable predictions from optimal control theory.

We have recently (Diedrichsen, 2007) used a bimanual reaching task to test whether this approach can be fruitfully applied to the problem of bimanual coordination. In this experiment participants hold on to a robotic device with each hand and view visual feedback on a screen. In the two-cursor task (Fig. 2A), two cursors are displayed indicating the position of the two hands. Two targets appear and participants are required to move the cursors as fast and as accurately as possible to these target locations. In a second condition (Fig. 2D) only one cursor is presented at the spatial midpoint between the two hands and has to be moved to a single target. Participants can therefore reach exactly as in the two-cursor condition. However, because only the position of the cursor but not the individual locations of the hands is task-relevant, this tasks has an additional degree of freedom.

We can capture the behavior of the participants in the two-cursor task using a simple cost function (J).

$$J = \sum_{t=1}^{T} w_{p,t}(\|p_{L,t} - g_L\|^2 + \|p_{R,t} - g_R\|^2) + w_{v,t}(v_{L,t}^2 + v_{R,t}^2) + w_u(u_{L,t}^2 + u_{R,t}^2) \qquad (1)$$

The first term of the cost function reflects the requirement to minimize the distance between the left and right hand p_L and p_R and their respective targets g_L and g_R. The second term forces the system to stop at the target, i.e., reach zero velocity (v). The weights for these terms increase exponentially

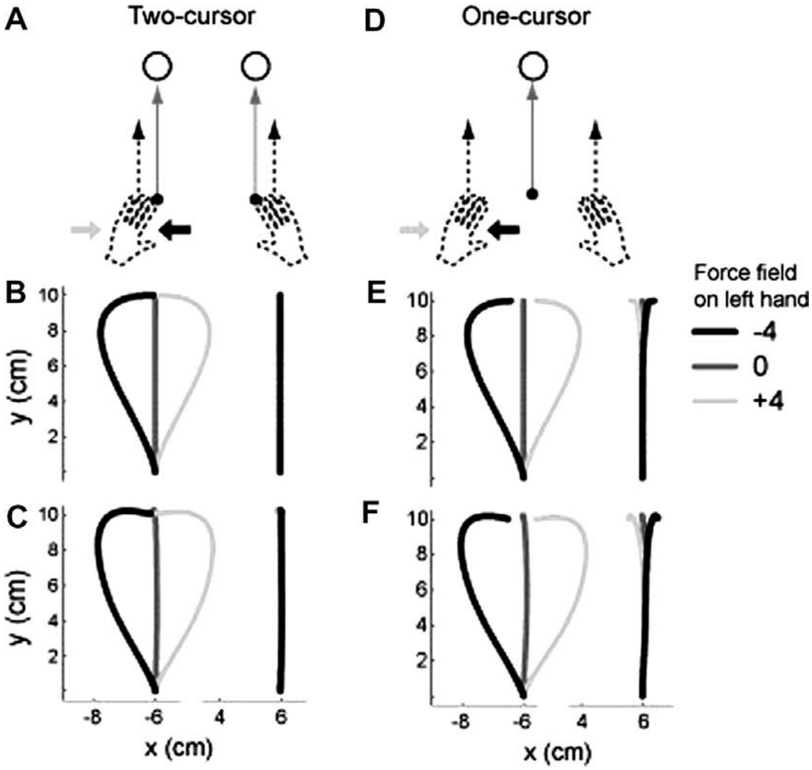

Fig. 2. Simulations and results from Diedrichsen (2007). In the two-cursor task (A) participants reached for two targets with two cursors. In the one-cursor task (D) they reached with one cursor towards one target. One of the hands was perturbed using a velocity-dependent force field to the left (−4, dark gray) or to the right (+4, light gray). Optimal control predicts independent feedback control for the two-cursor task (B) but dependent feedback correction when using one cursor (E). Average movement trajectories (C,F) confirm these predictions.

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003

over the movement time, rather than increasing with a step-function. In such a way, movement time can be treated as a soft constraint and speed-accuracy tradeoffs can be modeled. The last term of the cost function is the sum of the squared motor commands (u), weighted evenly across the time of movement. This term minimizes the expended energy and signal-dependent noise (Harris & Wolpert, 1998; Todorov, 2005) during the movement.

To derive testable predictions from optimal control theory, we needed to generate a reasonably realistic model of the arm movements. We started from a locally linear approximation to a suitable model of the human arm (Shadmehr & Mussa-Ivaldi, 1994; Todorov, 2005), and computed the control policy that minimizes Eq. (1), which can be expressed as a linear feedback rule with time-varying gains, i.e., $u_t = -L_t\hat{x}_t$. By applying control gains L to this system, we arrived at predictions for the optimal kinematic trajectories of both hands (for details, see Methods). The weights of the cost functions were then estimated to match the movement speed and the slightly asymmetric velocity profile of participants during the two-cursor task. In the experiment, we also applied a viscous curl field to one of the arms, which led to a lateral deviation of the arm trajectory and a subsequent correction. To capture the fast initial reactions of the arm to the perturbation, arising from passive muscle stiffness and spinal reflexes, we added a spring-like stiffness of the arm to the model, and estimated the stiffness parameter to match the observed lateral velocity profile. Using this model, the predicted trajectory of the arm perturbed by the force field (Fig. 2B) matches closely the averaged observed trajectory (Fig. 2C).

Up to this point we have used the cost function approach only to fit the behavior of the participants. Now, however, we can start to test predictions arising from this model. First, the control policy for the two-cursor task is block-diagonal, i.e, the commands to the left hand u_L only depend on the estimated state of the left hand $\hat{x}_L$, but are independent of the state of the right hand.

$$\begin{bmatrix} u_L \\ u_R \end{bmatrix}_t = -\begin{bmatrix} L_{LL} & 0 \\ 0 & L_{RR} \end{bmatrix}_t \begin{bmatrix} \hat{x}_L \\ \hat{x}_R \end{bmatrix}_t \tag{2}$$

This predicts that feedback control of the two hands is independent, i.e., only the hand that was perturbed by a force field should show a feedback correction, but not the other hand (Fig. 2B). The data show that this is indeed the case; for the two-cursor condition feedback control is independent (Fig. 2E).

Second, we can predict the behavior in the one-cursor task. To do this, we use the same system and cost function parameters as for the two cursor task. The only change is that the cost function now reflects the new task requirement: to bring one single cursor to a single target, rather than two cursors into separate targets.

$$J = \sum_{t=1}^{T} 2w_{p,t}\left\|\frac{p_{L,t}+p_{R,t}}{2} - g\right\|^2 + w_{v,t}(v_{L,t}^2 + v_{R,t}^2) + w_u(u_{L,t}^2 + u_{R,t}^2) \tag{3}$$

The cost function predicts that when one hand is perturbed both hands are responding with corrections in the same spatial direction. This behavior arises from the shape of the optimal control policy, which now contains off-diagonal nonzero terms:

$$\begin{bmatrix} u_L \\ u_R \end{bmatrix}_t = -\begin{bmatrix} L_{LL} & L_{LR} \\ L_{RL} & L_{RR} \end{bmatrix}_t \begin{bmatrix} \hat{x}_L \\ \hat{x}_R \end{bmatrix}_t \tag{4}$$

The matrices L_{LR} and L_{RL} determine the amount and direction of the coordination between the two hands. Why would this coordinative control policy (Eq. (4)) be better than the independent control policy (Eq. (2))? The latter clearly also constitutes a satisfactory solution for the one-cursor task. However, a coordinative strategy is optimal under the cost function of the one-cursor task (Eq. (3)), because it distributes the motor corrections across the two hands, minimizing the term $u_{L,t}^2 + u_{R,t}^2$. The exact distribution of the online correction to the two hands is mostly determined by the stiffness of the arm, which we have estimated from perturbed trials in the two-cursor condition. Based on this estimate, the predicted and observed amount of bilateral correction matches well (Figs. 2E, F). Importantly, the corrective response was unchanged when visual feedback of the cursor was withdrawn during the course of a whole trial. Thus, the task-dependent coordination did not depend on the visual feedback, but rather on an internal estimate of the position of the cursor based on proprioceptive feedback from both hands.

The proposed model predicts accurately another important feature of coordinative movements. It is often observed that the variability of multiple effectors is structured in a task-dependent fashion; variability is minimized along the task-relevant dimensions, but allowed to accumulate along task-irrelevant dimensions. This characteristic has been termed the minimum-intervention principle (Braun & Wolpert, 2007; Todorov & Jordan, 2002) or the uncontrolled manifold hypothesis (Domkin, Laczko, Jaric, Johansson, & Latash, 2002; Latash, Li, & Zatsiorsky, 1998; Latash, Scholz, & Schöner, 2002). In the one-cursor task, the system has an additional degree of freedom, given by the fact that the end position of the hands is not task-relevant, as long as the cursor is close to the target. As predicted, in unperturbed trials in the one-cursor task, participants show a negative correlation between the x position of the movement endpoints. In the two-cursor task this correlation is absent (Fig. 2, Diedrichsen, 2007).[1] This task-dependent shift in variability is naturally predicted by the optimal control model. When we simulate the model using optimal control policies (Eq. (2) or Eq. (4)) with the parameters used to explain movement corrections in the force field and add signal-dependent noise to the motor commands (the noise sources being independent for each hand), the model accurately predicts the task-dependent co-variance structure of the movement endpoints. The behavior arises naturally from the coordinative feedback policy, this time not reacting to external force fields but to small internal errors.

In this article we test a further prediction of the suggested optimal control model of coordination (Diedrichsen, 2007), namely that the feedback gains are linear. In Eq. (4), the motor commands to the left hand are calculated as the sum of two functions, one depending on the state of the left hand, one depending on the state of the right hand. Thus, when we perturb the right hand in the same direction as the left hand, we should observe stronger feedback corrections than when both hands are perturbed in opposite directions. More precisely, the influence of a force field applied to the right hand onto the correction of the left hand should be the same, no matter what force field is presented to the left hand.

The prediction of linear feedback gains arises from a number of modeling decisions taken to make the optimal control problem mathematically tractable. First, we used a linear approximation to a non-linear arm system, motivated by the fact that the arm is indeed fairly linear over the course of a 10 cm movement. Second, we posited that participants were trying to minimize the squared difference between target and cursor, which in turn is a linear function of the positions of the hands. Are these assumptions reasonable? Alternatively, the arms may be coordinated in a non-linear fashion, with different feedback gains for when the hands are perturbed in the same or opposite directions.

To test the prediction of linear feedback gains, we applied 25 combinations of force fields with five levels to each hand under the two- and one-cursor condition. We estimated cost function and model parameters from the data in the two-cursor task and then predicted quantitatively the distribution of corrections across the two hands in the one-cursor task.

2. Method

2.1. Participants

Eleven participants (5 female, 6 male) volunteered for the experiment. All were right-handed with a mean age of 23 years (SD = 3.4). All experimental procedures were approved by the ethics committee of the School of Psychology at Bangor University.

2.2. Apparatus and stimuli

Participants were seated in front of the apparatus while holding on to a robotic device with each hand (Phantom 3.0, SensAble Technologies). Movements were performed in the natural reaching space in an upward-forward direction, involving shoulder and elbow movements, with the elbow

---

[1] This task-dependent difference in correlation also appears in the y direction (r = .35 in the two- and r = -.23 in the one-cursor task). Here all correlations are offset by a constant positive amount, resulting from common variability of the initial motor commands.

pointing downwards. A simulated spring (150 N/m) restricted the movements to a fronto-parallel plane, 20° from vertical. A horizontal crossbar stabilized the upper body and minimized interaction torques between left and right arm movements. Participants used a chin rest for head stabilization. Using two mirrors mounted at 90° to each other, participants viewed one monitor with the left and one monitor with the right eye. This stereoscopic display was calibrated with the robotic devices, such that stimuli could be displayed at their veridical 3D locations. Cursors and targets were displayed as 8 mm spheres.

2.3. Procedures

To start a trial, participants moved the two cursors into the starting spheres, displayed 6 cm to the left and right of the body midline at breast height. In the two-cursor condition, two targets were presented 10 cm in forward-upward direction above the starting spheres. Participants were instructed to reach to the targets by moving both hands rapidly upwards. In the one-cursor condition, a single target was presented at body midline, and a single cursor was presented at the spatial average position of the two hands. A movement started when the hand reached a speed of 3 cm/s and ended when the speed fell below 1.5 cm/s for 30 ms. Movement times of less than 700 ms with a movement endpoint less than 5 mm from the target position were awarded one point.

Participants performed one session in the one-cursor condition and one session in the two-cursor condition. Sessions were separated by at least one day, and the order was counterbalanced between participants. Each session consisted of 10 blocks of 50 trials. The first block was a training block with only unperturbed trials, followed by 9 blocks with random force fields applied to each hand. During these force field trials the robotic arm produced a sideways force that was proportional to the upward velocity on the movement plane with a proportionality constant of −4, −2, 0, 2, or 4 Ns/m. Each of the 25 combinations of right and left force fields was presented two times during each block. For the last 8 blocks of each session, visual feedback of the cursor was withdrawn between the start and end of the movement on half the trials. The trials were randomly chosen, such that participants did not know until after the movement started whether they would receive visual feedback or not.

Optimal control model. The model used was identical to that used in previous work (Diedrichsen, 2007). The optimal control problem was made tractable by using a local linear approximation to a two-joint arm (Shadmehr & Mussa-Ivaldi, 1994). In the model a vector of 10 variables captures the state of each hand. The state vector is composed of the x and y Euclidian coordinates of position (p), velocity (v), and force (f), muscle activation (h), and target position (g). The kinematics are modeled with

$$\begin{aligned} p_{t+1} &= p_t + \Delta t v_t \\ v_{t+1} &= v_t + \frac{\Delta t}{I}(f_t - kp_t) \end{aligned} \tag{5}$$

where Δt is the size of the discrete time step, t the current time step, and I is the inertia of the arm at the endpoint (0.5 kg). To account for the natural stiffness of the arm against the force field perturbations resulting from muscle properties and low-level reflexes, we added a spring-like stiffness with a spring constant of $k = 104$ N/m for the x direction. The spring constant was estimated to match the size of the feedback corrections in the two-cursor task.

We simulated the delay in the muscles by using two-coupled first-order low-pass filters with time constants $\tau_1 = \tau_2 = 40$ ms (Todorov, 2005) that related the force (f) to the motor command u. Thus, the produced force is a low-pass filtered version of the motor command u.

$$\begin{aligned} f_{t+1} &= f_t + \Delta t/\tau_2(h_t - f_t) \\ h_{t+1} &= h_t + \Delta t/\tau_1(u_t - h_t) \end{aligned} \tag{6}$$

To simulate sensory delay, we expanded the state-space with a series of four coupled first-order filters for the sensed position, velocity and force.

$$x_{t+1}^{(j)} = x_t^{(j)} + \Delta t/\tau_s(x_t^{(j-1)} - x_t^{(j)}), \text{ with } j = 1, \ldots, 4. \tag{7}$$

The current state of the system is $x^{(0)}$, and the state that can be sensed is $x^{(4)}$, where x is used as a placeholder for p, v, or f. The time constant for each filter was $\tau_s = 15$ ms.

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003

In summary, the state vector for each hand consisted of:

$$\mathbf{x}_t = \begin{bmatrix} p_t^{(0)} & v_t^{(0)} & f_t^{(0)} & h_t & g & p_t^{(1)} & v_t^{(1)} & f_t^{(1)} & \ldots & p_t^{(4)} & v_t^{(4)} & f_t^{(4)} \end{bmatrix}^T \tag{8}$$

And the sensed state is

$$\mathbf{y}_t = \begin{bmatrix} p_t^{(4)} & v_t^{(4)} & f_t^{(4)} \end{bmatrix} \tag{9}$$

As Todorov (2005), we calculated the Kalman gain of the system assuming sensory noise with zero-mean Gaussian distribution and standard deviations of 0.02 m, 0.2 m/s, and 1 N for p, v, and f, respectively.

By defining the appropriate matrices, the whole time-discrete system can be written as

$$\begin{bmatrix} \mathbf{x}_L \\ \mathbf{x}_R \end{bmatrix}_{t+1} = \mathbf{A}\begin{bmatrix} \mathbf{x}_L \\ \mathbf{x}_R \end{bmatrix}_t + \mathbf{B}\begin{bmatrix} \mathbf{u}_L \\ \mathbf{u}_R \end{bmatrix}_t \begin{bmatrix} \mathbf{y}_L \\ \mathbf{y}_R \end{bmatrix}_{t+1} = \mathbf{H}\begin{bmatrix} \mathbf{x}_L \\ \mathbf{x}_R \end{bmatrix}_t \tag{10}$$

A, B and H are "block-diagonal", i.e., there is no physical coupling between the left hand and the right hand.

From this linear system description, the quadratic cost functions (Eq. (1) or Eq. (3)), and the assumption of Gaussian noise, we can derive the optimal control gains $\mathbf{L}_t$ using the Riccati equations (Todorov, 2007). This computation is performed by starting at the end of the movement and iteratively moving backwards in time. We then simulated the reaction of the system under the optimal control policy, while the force field combinations from the experiment were applied to the left and the right hand. The amount of feedback correction was computed on the simulated data as on the experimental data.

3. Results

Before comparing the feedback corrections in the one-cursor task with the model predictions, we need to ascertain another prediction of the optimal control model: the average movement parameters of unperturbed movements under the one- and two-cursor condition should be approximately the same.

3.1. Movement times, accuracy, and temporal coupling

Congruent with this prediction, movement times were on average 740 ms, with no significant differences between the one- and two-cursor conditions, $t(10) = -0.584$, $p > .5$. The mean distance of the cursor (s) from the spatial target (s) in the end of the movement was 7 mm for the one-cursor task and 7.2 mm for the two-cursor task, again not significantly different, $t(10) = 0.412$, $p = .688$. In the two-cursor task the right hand was significantly more accurate (6.6 mm) than the left hand (7.04 mm), $t(10) = 5.23$, $p < .001$.

Movement times were approximately 40 ms longer (721 ms vs. 760 ms) when visual feedback was provided during that trial, and visual feedback improved the endpoint accuracy from 8.5 mm to 5.7 mm, $t(10) = 5.48$, $p < .001$, and these effects did not interact with cursor condition.

The movements of the two hands showed tight temporal coupling. The average interval between the movement start of the left and right hand was close to zero, with an average SD of 19 ms. Again, there was no difference in this measure between the one- and two-cursor condition, $t(10) = .873$, $p = .40$. Thus, the basic kinematic parameters of the average movements were, as predicted, very similar across the two task conditions.

3.2. Task-dependency of feedback corrections

We now turn to the analysis of how the two hands corrected for initial errors induced by velocity-dependent force fields applied to either hand. Fig. 3 presents average trajectories for 9 of the 25 force field combinations. In the two-cursor task the left correction was basically independent of the force

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003

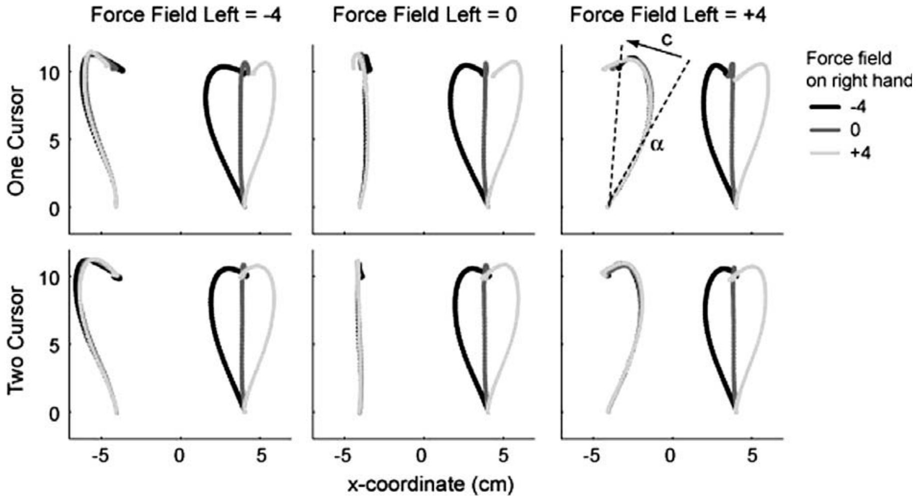

Fig. 3. Movement traces of the left and right hand, averaged across participants. Displayed are trials that had a particular level of force field applied to the left hand (columns), while a force field of −4 Ns/m (dark gray), 0 Ns/m (medium gray), or 4 Ns/m (light gray) was applied to the right hand. For each movement we calculated the initial movement error as the angular deviation from a straight-ahead movement (α), as well as the amount of angular correction (c).

field on the right hand. In the one-cursor task, however, the corrective movement of the left hand depended on the force field presented to the right hand.

For each hand and trial, we computed the initial movement direction (α) and the difference between initial and final movement direction (c). This measure of the online correction is displayed in Fig. 4 for the one- and two-cursor conditions for feedback and non-feedback trials, averaged across the hands. For the two-cursor condition this confirmed the prediction that the correction performed by each hand is largely independent of the force field applied to the other hand. From the data from the two-cursor condition we estimated an arm stiffness of 104 N/m to optimally match the size of the lateral perturbation caused by each force field. The resulting fit for the corrections is displayed in dashed lines.

We then predicted the amount of bilateral corrections for the one-cursor condition by exchanging the cost function for the task (Eq. (3) instead of Eq. (1)). All other model parameters were kept exactly identical. The model predicts that in the one-cursor condition, the correction of each hand is influenced by the force field applied to the other hand (dashed lines). The data clearly conforms to this prediction. To test this further we expressed the correction of each hand as a function of the initial direction error on the own and the other hand on that trial.

$$\begin{bmatrix} c_L \\ c_R \end{bmatrix}_n = \begin{bmatrix} C_L & C_{LR} \\ C_{RL} & C_R \end{bmatrix} \begin{bmatrix} \alpha_L \\ \alpha_R \end{bmatrix}_n \tag{11}$$

The correction rate C_L and C_R are within-hand correction rates and C_{LR} and C_{RL} are between-hand correction rates. A correction rate of 1 means that a hand corrects fully for the initial error, a correction rate of 0 implies that the correction was independent of the initial error.

Fig. 5 shows the correction rate for trials with and without visual feedback, and the left and the right hand. As can be seen, within-hand correction rates dropped from 1 in the two-cursor condition to .89 in the one-cursor condition, $F(1,10) = 104.71$, $p < 0.001$. However, it was still significantly different from the predicted optimal rate of .81, $t(10) = 6.69$, $p < .001$.

The between-hand correction rate showed the opposite pattern. It increased significantly in the one-cursor condition, $F(1,10) = 144.62$, $p < .001$, however not quite as high as the predicted value of .174, $t(10) = 3.99$, $p = .003$. Thus, we clearly replicated our previous result (Diedrichsen, 2007) and showed that bimanual feedback correction rates can be strongly modulated by task requirements.

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003

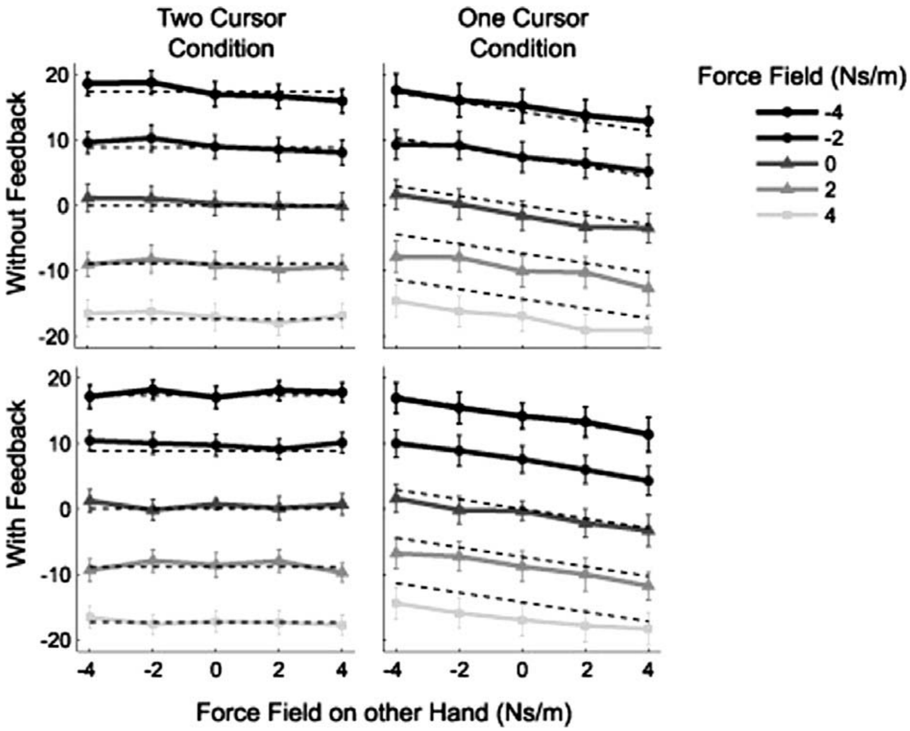

Fig. 4. Size of correction (in degrees) as a function of the force field on that hand (separate lines) and of the force field presented to the other hand (x axis). Results are averaged across both hands. In the two-cursor condition the correction only depends on the perturbation on this hand, while in the one-cursor condition it also depends on the force field presented to the other hand. The dashed lines indicate the predictions of the optimal control model.

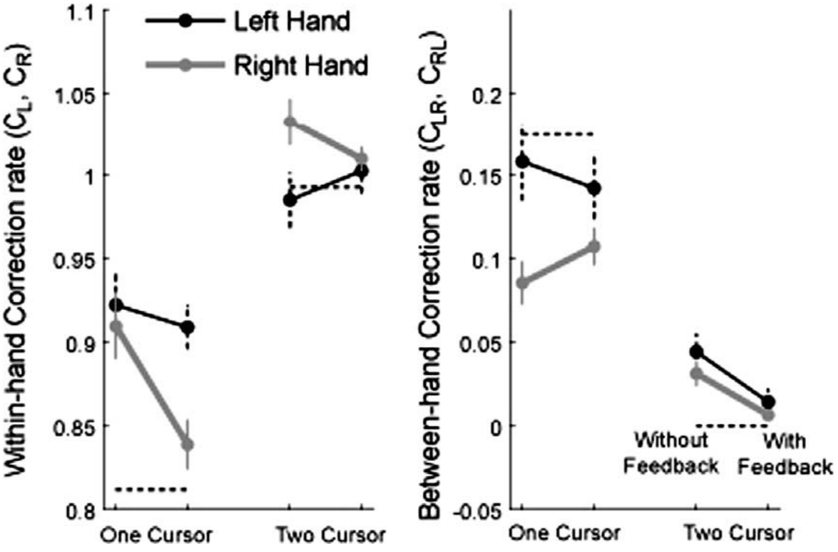

Fig. 5. Correction rates, across different task conditions. Within-hand correction rates indicate the proportion of the initial error (α) that is corrected by the hand itself. The between-hand correction rates indicate the proportion of error corrected by the other hand. Dashed lines are predictions of the optimal control model.

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003

3.3. Linearity of coordinative feedback control

The central interest for this article is to test the assumption of our model that interactions between hands arise from a time-varying linear feedback policy, in which the movement commands to one hand are a linear function of the state of the own hand, and to a lesser degree, of the state of the other hand. The linearity of the feedback control policy predicts that the influence of the left force field on the right hand correction is independent of the force field applied to the right hand, and vice versa. This means the slope of the lines in Fig. 4 should be the same for each force field level, i.e., the between-hand correction rate should be constant across all levels of force fields applied to the correcting hand. The prediction appears to be well matched by the empirical data.

To test this assumption quantitatively, we estimated the between-hand correction rate for each of the tasks (one- vs. two-cursor), feedback conditions, hands, and levels of force field on the own hand. A repeated measures ANOVA revealed no main effect of perturbation, nor an interaction with any of the other factors (all $F[4,40] < 1.667$, $p > .21$). Conversely, we estimated the within-hand correction rates for every level of the perturbation on the other hand, task condition, hand and feedback condition. Again, the factor of the perturbation was not significant and did not interact with any of the other factors (all $F[4,40] < 2.077$, $p > .10$). Thus, the degree to which each hand corrected for initial movement error of its own was independent of the force field applied to the other hand, and the degree to which it corrected for initial movement errors of the other hand was independent of the force field applied to the own hand. Thus, a linear coordination rule appears to be appropriate to model human behavior in this task.

We found some interesting differences between the hands in the one-cursor condition. The left hand showed a significantly higher within-hand correction rate than the right hand, $t(10) = 2.399$, $p = .037$, as well as a higher between-hand correction rate, $t(10) = 2.675$, $p = .023$. Thus, it appears that in the one-cursor condition participants relied slightly more on their left, non-dominant hand to correct for force fields on either hand.

3.4. Learning of coordinative feedback control

The one- and two-cursor tasks were performed in separate sessions with a total of 500 trials each. Thus, it is possible that participants slowly learned the optimal feedback control policy in the course of the experiment. To test for this possibility, we estimated the within-hand and between-hand correction rates for each block separately. During the first block of each session, participants performed 50 trials without any force field present. Thus, it was not possible to estimate correction rates on this block. However, from the beginning of block 2 (see Fig. 6), there was a significant difference between correction rates in the one- and two-cursor conditions, which do not appear to change systematically across the course of the experiment. Thus it appears that 50 movements without perturbation were enough to establish a new feedback control policy.

4. Discussion

In the present study we replicate and extend previous findings (Diedrichsen, 2007) and show that feedback control changes dependent on the task goals. In the two-cursor condition, in which separate effectors have to be moved to separate targets, feedback control is nearly independent across hands. In the one-cursor condition the movement corrections become dependent on the initial error of the other hand.

Here we test whether coordination can be described with a simple set of locally linear feedback gains (Eq. (4)), such as they are predicted when using a Linear Quadratic Gaussian (LQG) description of the optimal control problem (Todorov, 2007). Note that in this paper we are less concerned of whether the underlying arm system can be described by a single linear approximation. Possibly non-linearities – such as they become increasingly important when longer movements are studied – can be modeled by using time-varying system matrices for each hand, resulting in time-varying, but locally linear feedback matrices for each hand L_L and L_R (Todorov & Li, 2005). We are here more

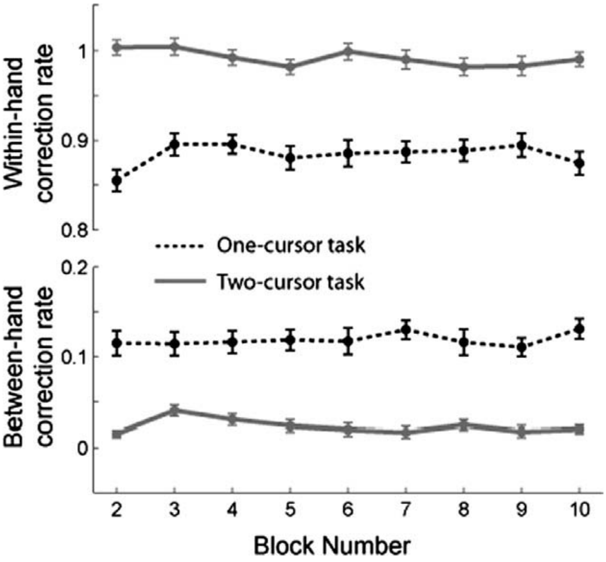

Fig. 6. Within-hand and between-hand correction rates as a function of block (50 trials each). Correction rates were averaged across hands and feedback conditions. Correction rates could not be estimated on the first block, as no perturbations were present in this block.

concerned with the linearity of coordination, the hypothesis that the motor commands to each hand can be written as the sum of two functions (possibly non-linear) – one depending on the left hand, one depending on the right hand, without any multiplicative interactions between the states of the two hands. The data conforms to this prediction; the influence of one hand onto the other was the same across all force field levels tested here.

Is this universally true? Our task was designed to be linear between the hands, i.e., the single cursor was presented at the average location of the two hands. It remains an open question whether the brain could learn tasks in which the coordination is a non-linear function of the states of the two hands, for example in a task where the participants would only need to correct when both hands are perturbed to the right, but not if only one hand is perturbed to the right. In this case motor commands would optimally have to be a non-linear function of the states or both hands. While it is possible that participants would adapt a non-linear control policy under such circumstances, they may also try to approximate control using a linear policy. Latter solution may be preferred because most objects manipulated with two hands have linear dynamics.

How well can this model predict the exact amount or strength of the bilateral corrections? We have derived specific predictions about the size of the bimanual feedback adjustments given a task-dependent cost function. The observed correction rates certainly followed the predicted trends and were in the correct range, but the between-hands correction rates in the one-cursor condition were significantly smaller than predicted. In general, however, the quantitative predictions about distribution of corrections across the two hands in the one-cursor should be taken with some caution, as they depend quite strongly on the estimated arm stiffness (see supplementary materials, Diedrichsen, 2007). Here we have summarized passive muscle resistance against perturbations and spinal feedback control, both of which are assumed to be task-independent, as a simple spring-like stiffness around a desired trajectory (Shadmehr & Wise, 2005). We estimated the stiffness from data from the two-cursor condition and then predicted the optimal correction gains for the one-cursor condition. These simplifications limit our confidence in the exact quantitative predictions arising under this model. In the current experiments the movements were quite unconstrained and the robotic devices were not stiff enough to gain reliable empirical estimates of arm stiffness. While more controlled experimentation is clearly needed to address this issue, the current results clearly show that there are predictable

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003

task-related changes in bimanual feedback control, and that these changes are well-described with a linear feedback control policy.

Independent of the exact formulation of the plant model, it is always the case that if a cost function has separate quadratic terms for the goal of the left and the goal of the right hand, the optimal feedback policy will be block-diagonal (Eq. (2)). This implies that control of the two hands should be independent when they try to achieve independent goals. While this was clearly supported for the current task, a large number of studies have highlighted phenomena in which there is dependence between the hands, even when the task requires them to perform independent movements. For example, when attempting to draw to separate figures at the same time with the two hands, there are substantial interference phenomena between the hands (Franz, Zelaznik, & McCabe, 1991). The tendency of the two hands to produce mirror symmetric movements is easily observable during activities like hand circling (Semjen, Summers, & Cattaert, 1995) or oscillatory finger movements (Kelso, 1984). Furthermore, there is tight temporal coupling between the onsets of movements and, to a slightly lesser degree, between the movement ends (Kelso, Putnam, & Goodman, 1983; Marteniuk, MacKenzie, & Baba, 1984). In all these cases optimal control theory would have predicted independent control of the hands. So how can this contradiction be reconciled?

First of all it is important to consider the ecological validity of the tasks studied. Control is optimized through evolutionary pressures and in the learning history of a particular individual. Thus, tasks that are not part of our normal behavioral repertoire will be solved by the brain in an ad-hoc fashion and will likely not be controlled optimally. We think that it is insightful that we can observe the tendency to mirror symmetric movements during a task like bimanual circle drawing, something that we do not have to do very often, but that these tendencies are absent during the planning (Diedrichsen, Hazeltine, Kennerley, & Ivry, 2001) and online control (Diedrichsen, Nambisan, Kennerley, & Ivry, 2004) of bimanual goal directed reaching.

On the one hand, optimal control approaches have to take into account lower-level constraints of the system, which might have developed over a longer evolutionary time scale and as an optimization process to other cost functions. Clearly we cannot overcome these constraints on a moments notice. In determining the origin of these constraints, however, we need to be careful whether these can be really ascribed to the existing low-level coupling between the hands on the level of brain stem or spinal cord (Carson, 2005). Recent finding have shown that the tendency towards mirror symmetric movements is modifiable by visual feedback (Mechsner, Kerzel, Knoblich, & Prinz, 2001; Swinnen, Walter, Lee, & Serrien, 1993), or the allocation of attention (de Poel, Peper, & Beek, 2006; Swinnen, Jardin, & Meulenbroek, 1996). Furthermore, the tendency towards mirror symmetric bimanual movements disappears largely after callosotomy (Diedrichsen, Hazeltine, Nurss, & Ivry, 2003; Franz, Eliassen, Ivry, & Gazzaniga, 1996; Kennerley, Diedrichsen, Hazeltine, Semjen, & Ivry, 2002), suggesting a cortical rather than subcortical origin.

Finally, the optimal control model presented here has theoretical similarities with models that conceptualize motor control and planning as occurring on the level of anticipated movement goals (Prinz, 1990). We have chosen here to capture the differences between the task in terms of the cost function, while leaving the controlled plant (the two hands) and the forward model the same across different task. A slightly different way of accounting for the data is to include the controlled object – either the one or two cursors – explicitly into the forward model and the estimated state. The cost function and consequently the control of the hands would then be based strongly on the state of the cursor – the anticipated movement outcome. While such a model makes mathematically equivalent predictions to the formulation chosen here, it formalizes the idea that movement control and coordination is determined by the anticipated consequences of those movements. This is congruent with findings that highlight the importance of goal-level representations in determining the coordination between the hands (Kunde & Weigelt, 2005; Mechsner et al., 2001; Swinnen et al., 1993).

In sum, the results presented here raise three important issues. First, the findings constitute a clear example of task-dependent coordination. This stresses the importance of studying a range of ecologically valid tasks when trying to extract general principles of bimanual coordination. Second, the results show the applicability of optimal control theory to bimanual control. Coordination can be understood in terms of a set of locally linear feedback gains, which express the motor commands to one effector as being dependent on the state of another effector. The results also highlight the importance of incorporating interaction between lower-level mechanisms (the stiffness of the hand) and higher-level control (task-dependent control policies) into optimal control models. Finally, the presented model formalizes the idea that coordination arises from an optimal solution of achieving a specific distal goal.

Acknowledgements

This work was supported through grants from the Biotechnology and Biological Sciences Research Council (BBSRC, BB/E009174/1) and the National Science Foundation (NSF, BCS-0726685).

References

Braun, D. A., & Wolpert, D. M. (2007). Optimal control: When redundancy matters. *Current Biology, 17*, R973–975.

Carson, R. G. (2005). Neural pathways mediating bilateral interactions between the upper limbs. *Brain Research Brain Research Reviews, 49*, 641–662.

de Poel, H. J., Peper, C. L., & Beek, P. J. (2006). Laterally focused attention modulates asymmetric coupling in rhythmic interlimb coordination. *Psychological Research, 72*, 123–137.

Diedrichsen, J. (2007). Optimal task-dependent changes of bimanual feedback control and adaptation. *Current Biology, 17*, 1675–1679.

Diedrichsen, J., Hazeltine, E., Kennerley, S., & Ivry, R. B. (2001). Moving to directly cued locations abolishes spatial interference during bimanual actions. *Psychological Science, 12*, 493–498.

Diedrichsen, J., Hazeltine, E., Nurss, W. K., & Ivry, R. B. (2003). The role of the corpus callosum in the coupling of bimanual isometric force pulses. *Journal of Neurophysiology, 90*, 2409–2418.

Diedrichsen, J., Nambisan, R., Kennerley, S. W., & Ivry, R. B. (2004). Independent online control of the two hands during bimanual reaching. *European Journal of Neuroscience, 19*, 1643–1652.

Domkin, D., Laczko, J., Jaric, S., Johansson, H., & Latash, M. L. (2002). Structure of joint variability in bimanual pointing tasks. *Experimental Brain Research, 143*, 11–23.

Flash, T., & Sejnowski, T. J. (2001). Computational approaches to motor control. *Current Opinion in Neurobiology, 11*, 655–662.

Franz, E. A., Eliassen, J. C., Ivry, R. B., & Gazzaniga, M. S. (1996). Dissociation of spatial and temporal coupling in the bimanual movements of callosotomy patients. *Psychological Science, 7*, 306–310.

Franz, E. A., Zelaznik, H. N., & McCabe, G. (1991). Spatial topological constraints in a bimanual task. *Acta Psychologica, 77*, 137–151.

Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature, 394*(6695), 780–784.

Kelso, J. A. S. (1984). Phase transitions and critical behavior in human bimanual coordination. *American Journal of Physiology, 246*(6 Pt 2), R1000–R1004.

Kelso, J. A. S., Putnam, C. A., & Goodman, D. (1983). On the space-time structure of human interlimb coordination. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 35*, 347–375.

Kennerley, S., Diedrichsen, J., Hazeltine, E., Semjen, A., & Ivry, R. B. (2002). Callosotomy patients exhibit temporal and spatial uncoupling during continuous bimanual movements. *Nature Neuroscience, 5*, 376–381.

Kunde, W., & Weigelt, M. (2005). Goal congruency in bimanual object manipulation. *Journal of Experimental Psychology: Human Perception and Performance, 31*, 145–156.

Latash, M. L., Li, Z.-M., & Zatsiorsky, V. M. (1998). A principle of error compensation studied within a task of force production by a redundant set of fingers. *Experimental Brain Research, 122*, 131–138.

Latash, M. L., Scholz, J. P., & Schöner, G. (2002). Motor control strategies revealed in the structure of motor variability. *Exercise and Sport Science Reviews, 30*, 26–31.

Marteniuk, R. G., MacKenzie, C. L., & Baba, D. M. (1984). Bimanual movement control: Information processing and interaction effects. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 36*, 335–365.

Mechsner, F., Kerzel, D., Knoblich, G., & Prinz, W. (2001). Perceptual basis of bimanual coordination. *Nature, 414*(6859), 69–73.

Prinz, W. (1990). A common coding approach to perception and action. In O. Neumann & W. Prinz (Eds.), *Relationships between perception and action* (pp. 167–201). Berlin: Springer.

Semjen, A., Summers, J. J., & Cattaert, D. (1995). Hand coordination in bimanual circle drawing. *Journal of Experimental Psychology: Human Perception & Performance, 21*, 1139–1157.

Shadmehr, R., & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience, 14*(5 Pt 2), 3208–3224.

Shadmehr, R., & Wise, S. P. (2005). *The computational neurobiology of reaching and pointing*. Cambridge: MIT Press.

Swinnen, S. P., Jardin, K., & Meulenbroek, R. (1996). Between-limb asynchronies during bimanual coordination: Effects of manual dominance and attentional cueing. *Neuropsychologia, 34*, 1203–1213.

Swinnen, S. P., Walter, C. B., Lee, T. D., & Serrien, D. J. (1993). Acquiring bimanual skills: Contrasting forms of information feedback for interlimb decoupling. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*, 1328–1344.

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

Todorov, E. (2007). Optimal control theory. In K. Doya, S. Ishii, A. Pouget, & R. P. N. Rao (Eds.), *Bayesian brain. Probabilistic approaches to neural coding*. Boston: MIT Press.

Todorov, E., & Li, W. (2005). A generalized iterative LQG method for locally-optimal feedback control on constrained non-linear stochastic systems. In *Proceedings of the American control conference*.

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

Vaziri, S., Diedrichsen, J., & Shadmehr, R. (2006). Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and the actual sensory feedback. *Journal of Neuroscience, 26*, 4188–4197.

Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science, 269*, 1880–1882.

Please cite this article in press as: Diedrichsen, J. & Dowling, N. Bimanual coordination as task-dependent linear ... *Human Movement Science* (2009), doi:10.1016/j.humov.2008.10.003