

Alexander J. McC. Foulkes · R. Chris Miall

Adaptation to visual feedback delays in a human manual tracking task

Received: 11 January 1999 / Accepted: 5 November 1999 / Published online: 16 December 1999
© Springer-Verlag 2000

Abstract The time-course of human adaptation to spatial perturbations of visuomotor function (e.g. with prisms) is very short. However, it is not clear how rapid the adaptation to other aspects of perturbed feedback is. In this paper we report the adaptation to delayed visual feedback. Three groups of six subjects tracked unpredictable, continuously moving targets using a hand-held joystick while visual feedback of the joystick position was delayed (0 ms, 200 ms or 300 ms). Subjects clearly adapted to the delay, with a significant drop in tracking error, but changes in more subtle aspects of their tracking behaviour (such as changes in intermittency and their “impulse response functions”) were not consistently observed. We suggest that the adaptation seen was consistent with the idea of there being a “delay component” in the internal processes used in manual tracking, as proposed in models such as the Smith predictor model.

Key words Feedback delays · Visuomotor adaptation · Manual tracking · Intermittency · Smith predictor model · Cerebellum · Human

Introduction

Over the past two decades, evidence has accumulated that tells us much about the adaptability of our motor system. It has been shown that we can adapt, in motor tasks, to changes in the gain of visual feedback (Keating and Thach 1990; Bock 1992; Kerr et al. 1993; Deutschl et al. 1996), to visual distortion by prisms (Kane and Thach 1989; Thach et al. 1991, 1992; Martin et al. 1996a), to unexpected changes in load (Bock 1993), to Coriolis force perturbations (Lackner and DiZio 1994) and to virtual force fields (Shadmehr and Mussa-Ivaldi 1993;

Shadmehr et al. 1995; Shadmehr and Brashers-Krug 1997).

The motor system requires this capacity to adapt in order to deal with, for example, changes due to muscle growth or wastage, changes in external loads, or changes in the sensory systems guiding movement. As yet, however, there has been little attention paid to the question of adaptation to temporal aspects of sensory feedback. It has been clearly demonstrated that, when subjected to visual feedback delays, both human subjects and monkeys find manual tracking tasks considerably more difficult than when no delay is present (Warwick 1949; Pew et al. 1967; Miall et al. 1985). Furthermore, it is very likely that, just as a person who wears glasses develops two vestibulo-ocular reflexes, one for when wearing glasses and one for when not, so we also learn to control for different situations such as operating machines (for example, driving a car or steering a boat) when, due to mechanical lags in the system, visual feedback of the results of control actions is delayed longer than normal. Thus there is a need to adapt to visual feedback delays, but as yet no investigation has been made into this adaptation process. To demonstrate that the human motor system can adapt to visual feedback delays may also have important implications in the field of telerobotics, for instance in transglobal surgery, where there are inherently lengthy feedback delays due to the long transmission distances involved. This is the aim of the present study.

An indication as to what one might expect during adaptation when visual feedback is delayed in a manual tracking task comes from models of cerebellar function such as that proposed in the Smith predictor hypothesis (Miall et al. 1993). Central to this theory is the idea that a “delay component” in the predictive model is set to be the same as the delay in the external feedback loop (Smith 1959). Hence the sensory feedback can be compared in temporal synchrony with the predicted feedback, and the difference in the two signals is simply the error in the prediction and can be added directly into the next movement. In their paper Miall et al. (1993) demonstrated, using a computer simulation, that a mismatch of

A.J.M. Foulkes · R.C. Miall (✉)
University Laboratory of Physiology, Parks Road, Oxford,
OX1 3PT, UK
e-mail: rcm@physiol.ox.ac.uk
Tel: +44-1865-282162, Fax: +44-1865-272469

the internal and external delays causes instability in the feedforward control mechanism. This would result in greater error in the subjects' tracking. Warwick (1949), Pew et al. (1967) and Miall et al. (1985) have all demonstrated that the errors in the subjects' tracking are worse with greater feedback delays. A decrease in this error with continued exposure to the feedback delay would therefore suggest that subjects are adapting to the delay. Miall et al. (1993) showed that reducing the open loop gain, which produced sluggish tracking, was one adaptive strategy. Adapting the internal delay was much more effective, however, allowing high open-loop gain tracking, despite the increased feedback delay.

Measurement of tracking error alone is rather crude and may not capture very well the more subtle aspects of the subjects' behaviour such as intermittent tracking responses. Intermittency is a characteristic feature of manual tracking (Craig 1947), and is heavily dependent on the visual feedback delay (Pew et al. 1967; Miall et al. 1985). A simple way to look at intermittency is to take the power spectrum of the velocity record of the subject's tracking. The instability predicted by Miall et al. (1993) when the internal and external feedback delays are mismatched would manifest itself as an increased magnitude of intermittency at certain frequencies. Adaptation to the feedback delay would therefore be indicated by a gradual decrease in the magnitude of the intermittency.

It is expected that this decrease in the magnitude of the intermittency would be accompanied by a gradual increase in the frequency of intermittency. Miall et al. (1985) also found that the spectral components due to intermittency moved to lower frequencies as the feedback delay increased, so a gradual increase in the frequency of intermittency as subjects are exposed to a consistent feedback delay would imply that the subjects are adapting specifically to that delay.

A further idea of what changes might be expected during adaptation to visual feedback delays comes from Miall's (1996) study of the effects of using different target speeds on subjects' tracking performance. Miall looked at the spectral composition of subjects' velocity records when tracking fast and slow targets and concluded that subjects could modify the dynamics of their internal model for tracking tasks, depending on the speed of the target, allowing them better, smoother tracking. Another way of looking at adaptation to visual feedback delays, therefore, might be to look at subjects' responses to sharp perturbations in the target waveform. If the subjects are adapting to the delay one might expect that their responses to such target jumps, initially slow, would become brisker over time, suggesting that there is a retuning of the subject's response dynamics, in line with the adaptive changes predicted by the Smith predictor model.

The aims of this study, therefore, are to demonstrate the adaptability of the internal processes used in manual tracking by observing a clear improvement in tracking performance during exposure to a visual feedback delay,

and to gain some insight into the mechanism of adaptation by analysing the frequency composition of the subjects' velocity records during adaptation, and by looking at the subjects' responses to sharp, unexpected jumps in the target waveform, over this period.

Materials and methods

Subjects

The subject group comprised 18 subjects, divided into three experimental groups, 6 subjects in each, with an equal number of men and women, age range 20–42 years. The subjects were, with one exception, new to the task; no gross differences in their initial responses were seen. All subjects gave their informed consent prior to their inclusion in the study.

Experimental set-up

Each subject sat approximately 50 cm in front of a 38-cm computer monitor (with 640×399 pixel resolution) on which a target was displayed as a small, solid white square (4 mm × 4 mm). The position of the target was controlled by the experimental computer and followed a 2-dimensional pseudo-random path. The pseudo-random waveform for each axis was generated independently by the sum of five nonharmonic sinusoids (0.06, 0.11, 0.13, 0.25 and 0.33 Hz). At the beginning of each new trial, the amplitudes of these component sinusoids were randomised $\pm 10\%$ (from default equal amplitudes) and their relative phases randomised completely. The pseudo-random path for a typical 15-s period is shown in Fig. 1. The target could move across most of the screen area and had a peak velocity of 14°s^{-1} . The subject was instructed to "track the target as accurately as possible" by using a hand-held joystick to control the x - y position of a small cursor, an open white square of similar dimensions to the target, also displayed on the monitor. The joystick was a light-weight, low-friction unsprung model (Radio Spares 162–984, with the self-centring spring removed). It was 8 cm in length and needed to be moved through $\pm 28^{\circ}$ to follow the target, or 7.5 cm at its tip. There were no dynamics between the joystick and the monitor cursor display. The positions of the target and joystick were digitally sampled at 70 Hz with 12-bit resolution and saved onto disk after each trial.

Tracking task

Each subject performed seven control trials, followed by a test period of 40 trials, started directly after the control trials. Each trial ran for 2 min, with a pause between trials of 10–30 s. Throughout every 5th test trial, instantaneous, unexpected jumps in target position, at least half the width of the screen in amplitude, were added periodically into the target waveform. These trials were called

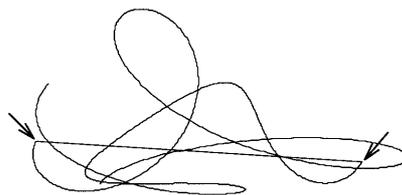


Fig. 1 The pseudo-random target waveform for a typical 15-s period, including a target perturbation, indicated by arrows. The x - and y -axes represent the x - and y -dimensions of the monitor screen

“perturbed” trials and were included to assess the subject’s response to the sudden velocity impulse, so that their “impulse response function” could be monitored as they tracked the target. Note that the impulse response function itself was never calculated, but that the averaged velocity records of the subject’s response to the perturbation were used as its equivalent in the time domain. There were about ten target jumps in each perturbed trial, and the time between jumps was randomised, though never less than 10 s. Figure 1 includes the trace of a typical target jump, indicated by arrows.

Subjects were randomly allocated to one of three groups. During the 40 test period trials, a delay was introduced between the subject moving the joystick and the corresponding movement of the cursor on the screen. The first group of six subjects were exposed to a 200-ms delay, the second to a 300-ms delay and the last group to no delay, which acted as a control group. To avoid the need to use cumbersome phrases such as “200-ms feedback delay group subjects” or worse, subjects will be referred to throughout as simply 0-ms, 200-ms or 300-ms subjects. The test period was broken into blocks of 15–20 trials, and subjects did one block of trials each day, spread out over 2–3 days, to avoid tiredness and boredom.

Immediately after the 40 test trials, all subjects performed two additional trials, the first a perturbed trial without any feedback delay, the second an “unperturbed” trial with a 400-ms delay. The experimental procedures used had full prior approval from the local ethics committee.

Data analysis

The data collected were the position on the screen of the target and the angular position of the joystick, calibrated in screen pixels. Four lines of analysis were taken. First, an error score was calculated for each trial as the mean tangential distance between the target and joystick positions over the trial, excluding the 1st s of data. This served as our basic measure of the subjects’ tracking performance.

Second, the target and joystick data were cross-correlated, and the mean correlation coefficient and lag calculated. Third, a perturbation response record was calculated for each perturbed trial: a 2-s block of joystick positional data was taken after each jump during the trial, differentiated and the tangential velocity of the joystick calculated. This was then normalised to the amplitude of the target jump. The mean of all these blocks of normalised velocity data gave the perturbation response record for that trial. This record was used to represent the subject’s impulse response function. The assumption has been made that in averaging the velocity data both within each trial, and later across subjects, the potential differences in the dynamics of movements in different directions due to biomechanical factors can be ignored.

Finally, a mean power spectrum was obtained for the unperturbed trial immediately following each perturbed trial (i.e. one unperturbed trial in every five trials). The velocity data was used rather than the position data to emphasise the high-frequency response components at the expense of the lower frequency target components. For each trial, a second-order 10-Hz Butterworth filter was applied to the velocity data from eleven 10-s segments of joystick positional data, missing out the 1st s of the trial. Each 10-s segment was then smoothed with a Hanning filter. Power spectra were determined for each of the 10-s segments using a fast Fourier transform (700 samples per 10-s segment, padded with zeros to 1024), and then averaged to give an overall power spectrum for the trial. A “centre of gravity” (CoG) measurement was used to characterise the spectrum to allow comparison with spectra from other trials. For this, the spectrum between 0.4 and 2.5 Hz was selected, and its centre of gravity calculated to give a mean frequency coordinate (Eq. 1) and a mean power coordinate (Eq. 2). The advantage of this approach is that it reduced the broad, noisy spectrum down to two parameters of interest (see Results). These two parameters were used to gain further insight into the changes in the intermittency in the subjects’ tracking.

$$\text{mean freq.} = \frac{\sum_{i=1}^{26} f_i \cdot P_i}{\sum_{i=1}^{26} P_i} \quad (1)$$

$$\text{mean power} = \frac{\sum_{i=1}^{26} f_i \cdot P_i}{\sum_{i=1}^{26} f_i} \quad (2)$$

where f_i is the frequency and P_i the power of the 26 data points within the range 0.4–2.5 Hz.

Because the frequencies are uniformly sampled, Eq. 2 can be reduced to:

$$\text{mean power} = \frac{\sum_{i=1}^{26} P_i}{26} \quad (3)$$

Results

General observations

Typical tracking data from unperturbed trials early in the test period are displayed in Fig. 2A–C. It can be seen that the magnitude of the intermittency in subjects’ tracking is greatest for the 300-ms subject (Fig. 2C), less in the 200-ms subject (Fig. 2B), and least in the 0-ms subject (Fig. 2A). This qualitative change is in line with the observations of Miall and his co-workers (1985). A comparison of Fig. 2B with Fig. 2D shows that the magnitude of the intermittency in this 200-ms subject’s tracking reduced considerably with practice, an observation that was common to all of the 200-ms and 300-ms subjects. A general improvement with time was also apparent in the subjects’ error scores: all subjects in all groups had improved significantly by the end of the test period.

Another brief qualitative observation on subjects’ responses to target jumps is that those with a feedback delay consistently overshoot the target jump, while those without a feedback delay did not overshoot very often, and that the overshoot in the 200-ms and 300-ms subjects’ tracking lessened over the test period.

Improvement in tracking performance

We first calculated each subject’s error score, defined as the mean tangential distance between the target and joystick positions over the whole trial. While this measure is dependent upon several factors, it has still proved a useful measure, showing clearly that subject performance on this task improved with practice.

Figure 3 shows the mean group data for error scores. The main graph is for unperturbed trials; the inset shows data only from perturbed trials. The perturbations caused the mean errors to be higher in the latter trials, but the overall pattern is close to that of the unperturbed trials. One might expect in a learning/adaptation task such as this that an exponential function would model the fall in errors best (Schmidt 1991). However, linear and exponential regressions showed that neither function gave a much better fit than the other (Table 1). A linear function

Fig. 2A–D Typical records of both target and subject tracking for a 0-ms subject (A), a 200-ms subject (B), and a 300-ms subject (C), all from trials early in the test period, and for the same 200-ms subject from a trial late in the test period (D). The motion recorded along one axis during 2D tracking. The principal observations are that both the error and the magnitude of the intermittency in the subjects' tracking are least for the 0-ms subject and greatest for the 300-ms subject, with the 200-ms subject intermediate. Also note that the error and the magnitude of the intermittency have dropped after extended exposure to the feedback delay (compare B and D)

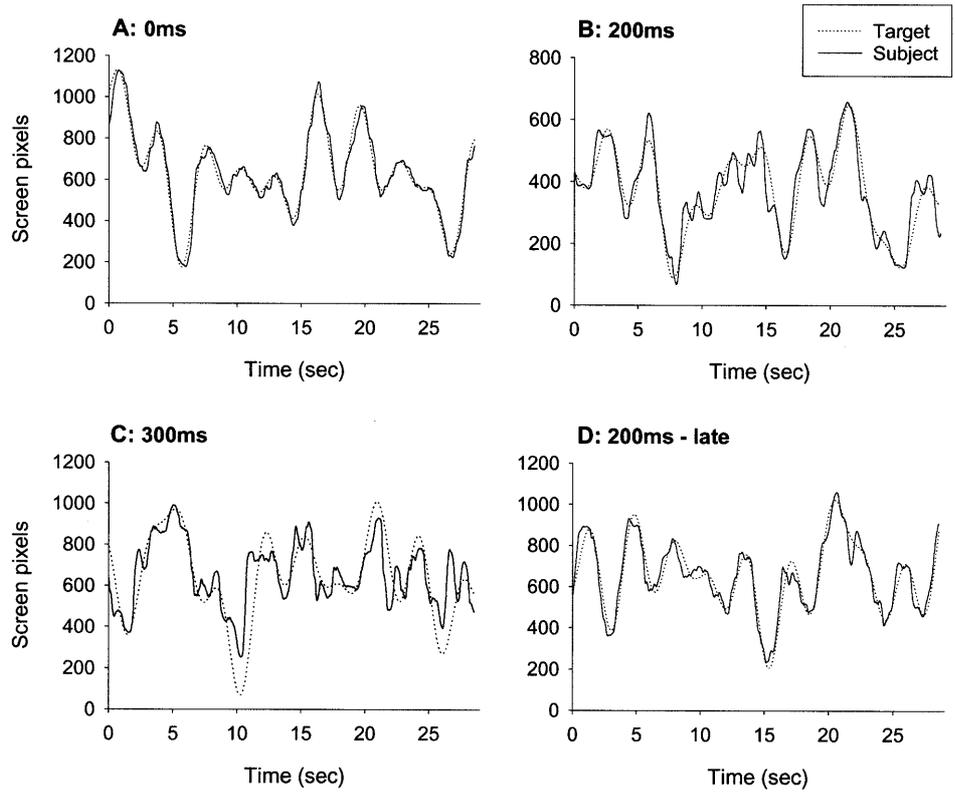
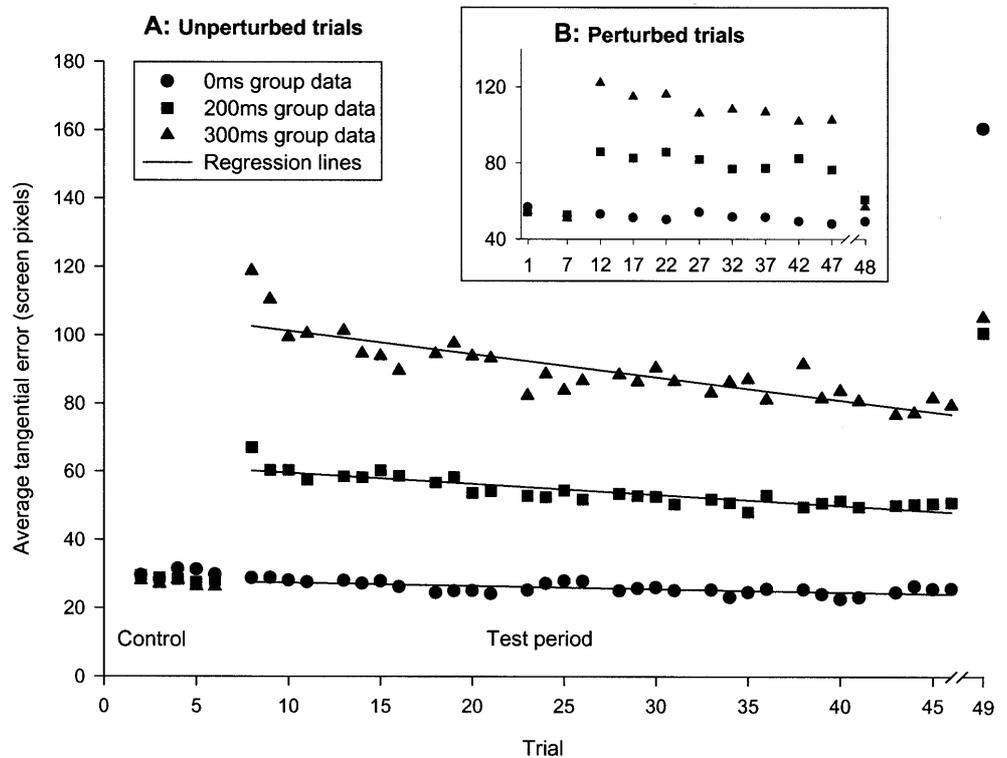


Fig. 3A,B Mean error scores for the three subject groups. A Data from unperturbed trials; B data from perturbed trials. All subjects performed similarly in the control trials, whereas there are clear differences between the groups when the feedback delay is introduced (test period). The 200-ms and 300-ms group data have significant gradients. Trial 48 was a perturbed trial with no feedback delay; trial 49 was an unperturbed trial with a 400-ms delay



has therefore been used for all further regression statistics using the error score data, because it simplifies greatly the statistics for intergroup comparisons.

The data from the first trial in each new condition (trials 1, 8, 48 and 49) were subdivided into three sequential

segments, each third having its own error score calculated, in an attempt to bring out any exponential tendency in the data. However, no such trend was observed.

All three groups performed similarly in the control trials, whether perturbed (Fig. 3B) or unperturbed (Fig.

Table 1 r^2 -values for linear and exponential regressions of group mean error score data, fitted over unperturbed trials 8–46; see Fig. 3

	Linear	Exponential
0 ms	0.405	0.396
200 ms	0.766	0.782
300 ms	0.721	0.748

3A). The gradients of the regression lines for the 200-ms and 300-ms groups, fitted over unperturbed test trials 8–46, had highly significant gradients (both $P < 0.0001$), whilst the gradient for the 0-ms group's data was nearly significant ($P = 0.0848$). The 95% confidence intervals for these gradients demonstrate that they are all significantly different. Student's t -tests comparing the first 5 unperturbed trials of the test period against the last 5, for each subject group, confirm this reduction in error ($P < 0.001$; see Table 2). t -tests also show that the control performances (trials 1–7) were similar across the three subject groups ($P > 0.104$; see Table 3). There were no significant differences between the 0-ms subjects' performances in this control period and their performances in the first 5 trials of the test period, whereas there clearly were for the 200-ms and 300-ms subjects (Fig. 3). All statistical results for intragroup and intergroup comparisons are shown in Tables 2 and 3, respectively.

The last data points in Fig. 3A,B show the group mean error scores for the final 0-ms and 400-ms delay trials (trials 48 and 49, respectively). Note that trial 48 was a perturbed trial (Fig. 3B), so the error scores are raised relative to error scores for trials plotted in Fig. 3A, which are all unperturbed. There was no difference in the performance of the 0-ms subjects between the last perturbed trial of the test period (trial 47) and their performance in trial 48, as expected. However, both the 200-ms and 300-ms groups performed better than in pre-

vious perturbed trials (Fig. 3B), but worse than the control perturbed trials. All three groups performed very poorly in trial 49, with a 400-ms feedback delay (Fig. 3A). The interesting point is that the group whose performance was most different between trial 46 (the last unperturbed test trial) and trial 49 was the 0-ms group, whilst the group whose performance was least different was the 300-ms group. The change in performance of the 200-ms group was intermediate.

Cross-correlation analysis

The mean correlation coefficients calculated between the cursor and target records were initially very high (mean 0.987) for the unperturbed control trials and fell to 0.933 and 0.859, respectively when the 200-ms and 300-ms delays were imposed. There was a gradual return toward the initial values, such that the final unperturbed trial coefficients reached 0.974 and 0.945, respectively. These changes mirrored closely the changes observed for the tracking errors (Fig. 3).

The mean lag of the cursor motion with respect to the target was initially very small (mean 28.6 ms) and jumped to a lag of 84 ms and 88 ms when the 200- and 300-ms delays, respectively, were introduced. Again, there was a gradual change during the adaptation period, so that the lags on the final unperturbed trial were 91-ms and 124-ms, respectively. These figures represent close to half the imposed delay (45% and 41%, respectively), in other words, on average the joystick led the target motion, while the cursor motion lagged the target. Hence, calculation of the tracking errors between the target and the delayed cursor gave values almost identical to the tracking errors between target and joystick position, shown in Fig. 3. While the changes in lag were consistent throughout the adaptation period, they were not significant for either 200-ms or 300-ms groups.

Table 2 Student's t -test statistics for intragroup comparisons of error scores across trials. Control trials are those before the delay was imposed; *first* and *last* refer to the first five and the last five unperturbed trials of the test period; see Fig. 3

	t			P		
	0 ms	200 ms	300 ms	0 ms	200 ms	300 ms
Control vs first	0.829	-15.440	-23.407	0.410	<0.001	<0.001
Control vs last	2.268	-15.811	-18.405	0.027	<0.001	<0.001
First vs last	1.492	5.363	6.367	0.141	<0.001	<0.001

Table 3 Student's t -test statistics for intergroup comparisons of error scores. See Table 2 for details

	t			P		
	0 ms vs 200 ms	0 ms vs 300 ms	200 ms vs 300 ms	0 ms vs 200 ms	0 ms vs 300 ms	200 ms vs 300 ms
Control	0.841	1.650	0.976	0.404	0.104	0.333
First	-13.959	-21.569	-12.196	<0.001	<0.001	<0.001
Last	-14.996	-17.524	-10.186	<0.001	<0.001	<0.001

Measurement of response changes

Impulse response characterisation

The logic behind the perturbed trials was to be able to look at the dynamics of the responses of subjects as they adapted. Subjects' responses to the target jumps were brisk movements of the joystick, seen as peaks in the velocity traces starting 200–300 ms after the target jumped. It was predicted that as subjects adapted to the feedback delay the peaks in the velocity trace would become higher and appear earlier, corresponding to a faster response to each target jump. However, no such trends were seen in these data: the responses to a target jump appeared to be quite inconsistent, showing a large variance in onset time and magnitude. Talking to subjects after the experiments, it seems that various strategies were employed to return the joystick cursor onto the target after a jump, and that individual subjects did not use them consistently, even within a single trial. No further analyses of these perturbed trials will be presented.

Spectral analysis

Another measure of intermittency is the power within the 0.4- to 2.5-Hz frequency range in the power spectrum of individual trials, averaged across the eleven 10-s segments of data from each trial. The CoG of this spectrum was used to quantify changes in this band of frequencies.

Figure 4 demonstrates the sort of variation seen in power spectra from unperturbed trials between the three subject groups. The high power below about 0.4 Hz is a reflection in the tracking of the target's component frequencies. Frequencies below this were discounted in the CoG measurements of the spectra, as were frequencies above 2.5 Hz, since it was considered that intermittency in the subjects' tracking that was dependent upon visual feedback was not represented outside these frequency limits. It can be seen in Fig. 4 that the magnitude of the intermittency in the subjects' tracking, as indicated by the power in the frequency peak at about 0.8 Hz, is least for the 0-ms subjects (Fig. 4A), greater for the 200-ms subjects (Fig. 4B), and greatest for the 300-ms subjects (Fig. 4C), and that, comparing Fig. 4B and D, the magnitude of the intermittency in the 200-ms subjects' tracking reduced with practice. No clear change in the peak frequency was evident in these data.

These observations are supported by statistical analysis of the differences in the CoGs of these power spectra between the three subject groups and of the changes in the CoGs within each group with practice. There were no significant changes over the course of the experiment in mean frequency. Figure 5 shows the changes in mean power, which represents the magnitude of the intermittency in the original joystick trace, for the three subject groups. There were significant falls in mean power for the 200-ms and 300-ms subjects' data ($P \leq 0.0127$), and it

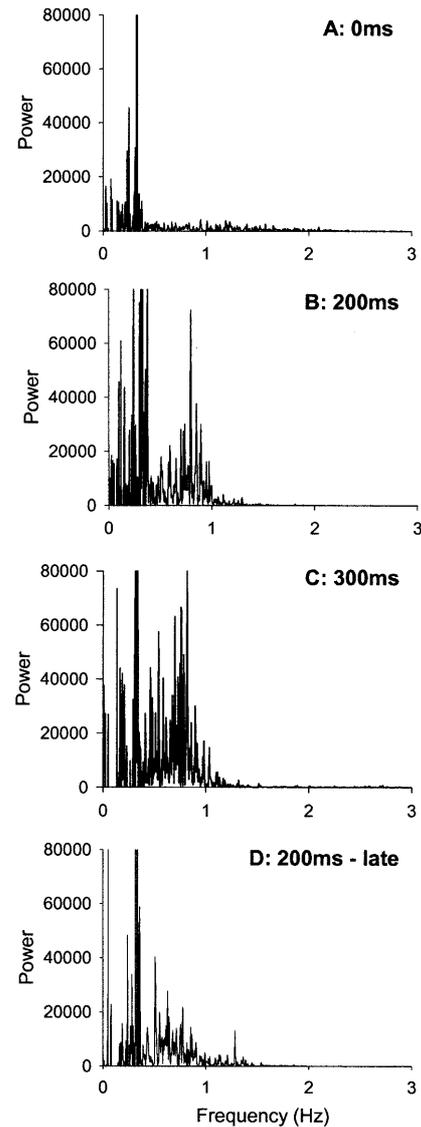


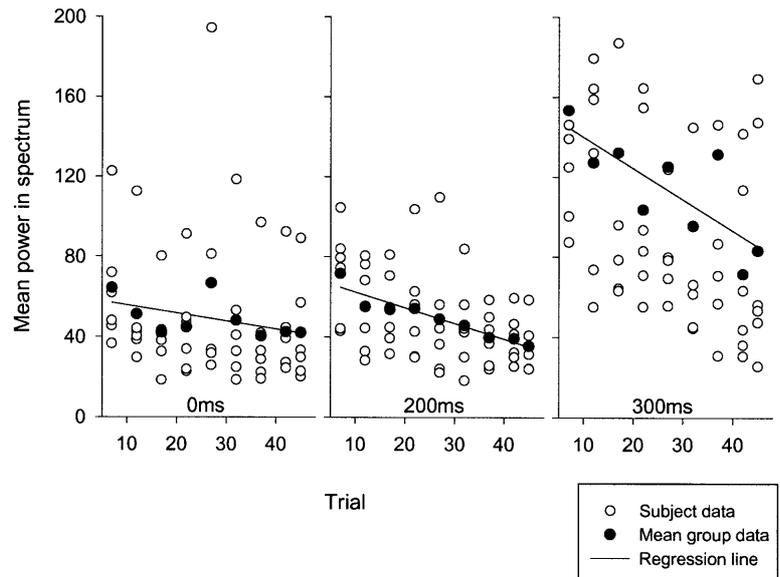
Fig. 4A–D Typical mean power spectra for a 0-ms subject (A), a 200-ms subject (B), and a 300-ms subject (C), all for trials early in the test period, and for the same 200-ms subject for a trial late in the test period (D). There is a clear increase in power in the frequency peak, around 0.8 Hz, with increased feedback delay, and a clear decrease in power as the subject adapts. The high power below about 0.4 Hz, in all four spectra, is the reflection in the tracking of the target's component frequencies

is clear from the figure that the gradients of these regression lines are related to the imposed delay.

Discussion

The aim of this study was to measure the adaptation in subjects' tracking responses to consistent exposure to a visual feedback delay and to assess the changes in their tracking behaviour over this period. The principal findings are that subjects show a clear adaptation to visual feedback delays, and that this can be seen as decreases both in the error scores of each subject and in the mean

Fig. 5 Mean power in the spectra calculated from unperturbed trials for the three subject groups. All three subject groups showed a significant reduction in mean power, quantified using the centres of gravity of the spectra, during exposure to a feedback delay. The gradients of the regression lines are clearly related to the imposed feedback delay



power in the power spectra of these trials over time. However, neither consistent changes in the peak frequencies of these spectra nor any consistent trend in subjects' responses to target jumps were found.

The error score data for the three subject groups are much as expected. They show a very clear pattern of improvement in the two groups with feedback delays, suggesting that there was a basic adaptation to these delays, in terms of better tracking behaviour. The rate of adaptation appears to be roughly proportional to the magnitude of the delay: the 300-ms group has the steeper gradient of the two test groups. The regression line through the 0-ms group error scores has a gradient that is almost significant, even though there was no delay to adapt to. Clearly practice improved the subjects' performance, independent from adaptation to the delays. This is presumably because, even though this task is very elementary, certain aspects of it can be learnt given sufficient experience, such as the basic properties of the target movement. Experience of how the target moves and, to some extent, when it is going to turn, especially as it gets towards the edge of the screen, can be acquired. It is widely recognised that almost any task, of any nature, has some aspect that can be learnt (Keele 1982).

A point to be noted concerning these error score data is that all three groups started with the same baseline score in the control trials. This implies that the groups' initial tracking performances were not greatly different and validates the intergroup comparisons.

A cross-correlation analysis was also performed on the tracking data to get an estimate of both the similarity of shape and the lag between target and cursor traces. The improvement in performance indicated by the increasing correlation coefficient was close to that indicated by the mean tangential errors. The lag between the traces for the 200-ms and 300-ms groups jumped up, from an initial mean value of 28 ms to 84 ms and 88 ms, respectively. One might have expected that the lag would

have jumped to a value close to the full feedback delay, but clearly subjects were able to predict ahead and thus compensate for about half this amount. Such subject prediction is visible in Fig. 2B–D, and has been previously demonstrated in a manual tracking task by Vercher and Gauthier (1992). By the end of the test period, the observed lag had increased to 91 ms and 124 ms, respectively. This may have been because, as subjects adapted and the magnitude of the intermittency in their tracking decreased, they did not have to predict so far ahead to maintain optimal tracking. However, the change in lag during the test period was not significant and cannot on its own account for the reduction in tracking error. It therefore seems that the subjects adapted their response latency but were also able to track more smoothly, reducing tracking error and increasing the correlation coefficient. Partial prediction of this sort of target function is commonly observed (Miall et al. 1985; Vercher and Gauthier 1992), even though the target waveform is pseudo-random. Human subjects are able to make short-term predictions about the smooth motion of the target, based on knowledge of its mean velocity, maximal range (clearly cued by the edges of the computer monitor) and perhaps also its frequency structure (Viviani and Mounoud 1990; Miall 1996).

The time course of the observed adaptation to feedback delays is very slow in comparison with those observed in studies of adaptation to other spatial visuomotor disturbances. In a study by Deuschl et al. (1996), subjects required about 50 single-movement trials to adapt to changes in the gain of their visual feedback. Similarly, Thach et al. (1992) showed that adaptation to a lateral displacement of the subject's line of vision using prisms was nearly complete after just 24 single-movement trials. From the present study, it would appear, through extrapolation from the regression lines of the 200-ms and 300-ms group data, that approximately 160 trials (about 5 h exposure) would be necessary for a similar degree of adap-

tation to these visual feedback delays. It should be noted that each trial in the present study lasted 2 min and consisted of many sub-movements, each of which could provide feedback information to assist adaptation, in contrast to the trials in the studies by Deuschl et al. (1996) and Thach et al. (1992), which were single ballistic movements. It seems, therefore, that adaptation to feedback delays has a different time course to the others mentioned above, and so may be implemented through a different mechanism. This longer adaptation period concurs with the findings of a study by Deno et al. (1989), in which monkeys adapted to an 80-ms delay in an ocular pursuit tracking task. Adaptation was found to be two-thirds complete after 3 days of 3-h training sessions. However, it is not clear what visual experience, if any, these monkeys had outside the exposure periods. In the present study it is likely that adaptation would remain relatively unaffected by interleaved behaviours performed without the additional feedback delay, because our task involved the use of a joystick in a specific context and in a specific behaviour, with which subjects were familiar, but which they were not using outside the experiment.

The reasons for including the perturbed trials have already been addressed in the Materials and methods and Results sections. However, there are two additional points worth making about these results. First, individual responses were very variable and, given the noisy data obtained, it is not so surprising that there was no clear trend. This problem might be overcome in future either by using very much larger groups of subjects, allowing more intersubject averaging, or by pre-training the subjects to reduce the variability of their performance.

Second, there was an unexpected problem due to the fact that they appeared to use various strategies to overcome the difficulty of following the target when it jumped. If they had used these strategies consistently, a trend of some sort in the velocity data would be expected. Instead we gained the impression that they were experimenting with various different strategies. It is clear that the instructions given were too open to interpretation by the subjects, so that in future experiments a stricter instruction such as “when the target jumps, follow it as quickly as possible” is needed.

The results from the spectral analysis of the unperturbed trials revealed a strong downward trend in the mean power of the spectra, but no changes in the mean frequency. Following Miall's (1996) suggestion that subjects are able to modify the dynamics of their internal model for tracking tasks, one might expect that the frequency of the spectral peak would increase as the subject adapted to the feedback delay and the dynamics of the internal model for the task were retuned, allowing brisker and more accurate tracking. The main difficulty was defining the spectral peaks. Miall (1996) used successfully a smoothing process to reveal the main peaks. However, in noisy spectral data such as those in the present study (Fig. 4), large isolated spikes in the spectrum would have a big influence on the smoothing process.

The smoothing would have to be very heavy to get well-defined peaks, resulting in a coarse damping of the overall shape of the spectrum. A smoothing process is therefore not ideal for these data. Furthermore, averaging over 11 segments of data per subject already provides some smoothing.

Instead of further smoothing the spectra, it was decided that a CoG measurement would be used, since it suffers less than the smoothing process from the biasing effect of large isolated peaks. However, no dramatic frequency shift was evident in the data, they were simply too noisy to see small changes such as might have been expected.

The question then arises as to why the data presented here were so much noisier than those, for instance, by Miall (1996). There are few obvious differences between the tasks in the two experiments that one might expect to produce this sort of disparity between the data. Miall (1996) used a one-dimensional (1D) target waveform, whereas a 2D waveform was used here; quite why this would make a difference to the data is not obvious, although it is possible that introducing a higher dimension of target movement could enrich the frequency spectrum of the joystick data, thus increasing its noise. Another difference is that Miall (1996) used experienced subjects, providing more reproducible data, so that there was less noise in the spectra, whereas in this study all but one of the subjects were naive. Finally, it may also be that using shorter feedback delays, which would cause less tracking difficulties, would provide more consistent, and therefore less noisy, data.

Returning to the data on the mean power of the spectra, which do show a significant downward trend, there is a simple conclusion to be drawn. The magnitude of the intermittency in the tracking decreased during exposure to the feedback delay, in full agreement with the error score data. The gradients of the regression lines through the mean power data are not significantly different between groups. However, all the gradients are themselves statistically significant, and, importantly, the 300-ms line is the steepest, followed by the 200-ms line, the shallowest gradient being that of the 0-ms line, reinforcing the idea that the rate of adaptation is roughly proportional to the magnitude of the feedback delay.

Our results provide clear evidence of adaptation to the feedback delays, but only limited evidence that this is achieved by a retuning of the system dynamics. The shift in response lag, cancelling about half the imposed delay, was accompanied by changes that led to more accurate, less intermittent, tracking. However, neither the spectral analysis nor the data from the perturbed trials provide strong indications of how this was achieved.

One general question that might be asked of the adaptation seen here is whether it is due to a readjustment of the dynamics of a single, adaptive internal model to a changed delay in the feedback loop, or due to the development of a second, new internal model tuned to the new feedback delay. A simple way to analyse this problem would be to use an experimental technique proposed by

Held and Gottlieb (1958), consisting of pre-adaptation, adaptation, and post-adaptation periods. If the performance of the subject in the first few trials of the post-adaptation period differs significantly from the performance in the pre-adaptation trials then it can be concluded that there has been a readjustment of the original internal model. If the pre- and post-adaptation performances do not differ significantly then it can be concluded that a second, new, internal model has developed.

The development of a second, new internal model has been demonstrated in prism adaptation experiments, but over a period of many weeks, not days (Martin et al. 1996b). It is therefore unlikely that the development of a new internal model is the major cause of the increase in performance observed here, where the course of the experiment was only 2 h of adaptation within a 48- to 72-h period. A single post-adaptation trial without a feedback delay (trial 48) is not alone sufficient evidence to draw a firm conclusion on this matter. However, as a qualitative observation, the post-adaptation performance was not as good as the pre-adaptation performance, pointing to a re-tuning of the dynamics of a single adaptive internal model for this tracking task. Furthermore, the fact that the change in performance between the last unperturbed trial of the test period and trial 49 (400-ms delay) was least for the 300-ms subjects and greatest for the 0-ms subjects, with the 200-ms subjects intermediate between the two, supports the suggestion that some re-tuning of the dynamics of an adaptive internal model for this task had taken place.

Finally, the significance of these results in the broader context of the various predictive internal models that have been proposed to account for certain characteristics of human movement must be considered. Such predictive models are necessary to explain how human movements, with durations of only 200–300 ms, can be too fast to be controlled by visual feedback. However, it is also known that normal movements are not purely feedback-controlled either, because disturbing the feedback signals during a movement leads to corrections before the movement is complete: to know that modification of an ongoing movement is needed implies knowledge of its expected outcome.

It has thus been proposed that the cerebellum could act as an inverse model on a feedforward pathway (Ito 1970; Neilson et al. 1988; Kawato and Gomi 1992), or a forward model on a feedback pathway (Miall 1989). However, many of these theories of internal models fail to fully account for the mismatch in time delays between the internal model of the controlled system and the actual feedback from that system. What is required by the controller is an estimate of the delay before the actual feedback will be received, so that it can distinguish between the expected reafferent feedback from its own performance and the sensations caused by errors in performance or from unexpected changes in the external environment. In showing that humans can adapt to changes in the timing of visual feedback, these results are supportive of the hypothesis that the internal processes used

for tracking tasks have a delay component in them. An example of this kind of model is the Smith predictor model (Smith 1959; Miall et al. 1993).

In conclusion, a clear improvement in tracking performance was observed over the course of the test period. This general improvement was accompanied by a decrease in power at the peak frequency in the power spectra of these trials, providing evidence consistent with the concept of a delay component in the internal processes used in manual tracking, as has been proposed in models such as the Smith predictor. Coupled with a subsequent deterioration in performance in a post-adaptation trial, in which there was no feedback delay, these results suggest that the adaptation is a re-tuning of an existing predictive model, rather than the development of a second, new model.

Acknowledgements This work was supported by the Wellcome Trust (R.C.M.) and the MRC (A.J.M.F.).

References

- Bock O (1992) Adaptation of aimed arm movements to sensorimotor discordance: evidence for direction independent gain control. *Behav Brain Res* 51:41–50
- Bock O (1993) Early stages of load compensation in human aimed arm movements. *Behav Brain Res* 55:61–68
- Craik KJW (1947) Theory of the human operator in control systems. 1. The operator as an engineering system. *Br J Psychol* 38:56–61
- Deno DC, Keller EL, Crandall WF (1989) Dynamic neural network organisation of the visual pursuit system. *IEEE Trans Biomed Eng* 36:85–91
- Deuschl G, Toro C, Zeffiro T, Massaquoi S, Hallett M (1996) Adaptation motor learning of arm movements in patients with cerebellar disease. *J Neurol Neurosurg Psychiatry* 60:515–519
- Held R, Gottlieb N (1958) Technique for studying adaptation to disarranged hand-eye coordination. *Percept Mot Skills* 8:83–86
- Ito M (1970) Neurophysiological aspects of the cerebellar motor control system. *Int J Neurol* 7:162–176
- Kane SA, Thach WT (1989) Palatal myoclonus and function of the inferior olive: are they related? In: Strata P (ed) *The olivocerebellar system in motor control*. Springer, Berlin Heidelberg New York, pp 427–460
- Kawato M, Gomi H (1992) A computational model of four regions of the cerebellum based on feedback-error-learning. *Biol Cybern* 68:95–103
- Keating JG, Thach WT (1990) Cerebellar motor learning: quantitation of movement adaptation and performance in rhesus monkeys and humans implicates cortex as the site of adaptation. *Soc Neurosci Abstr* 16:762
- Keele SW (1982) Component analysis and concepts of skill. In: Scott Kelso JA (ed) *Human motor behaviour: an introduction*. Erlbaum, London, pp 143–159
- Kerr GK, Miall RC, Stein JF (1993) Visuo-motor adaptation during inactivation of the cerebellar nuclei: a preliminary report. *Hum Mov Sci* 12:71–83
- Lackner JR, DiZio P (1994) Rapid adaptation to Coriolis force perturbations of arm trajectories. *J Neurophysiol* 72:299–313
- Martin TA, Keating JG, Goodkin HP, Bastian AJ, Thach WT (1996a) Throwing while looking through prisms. I. Focal olivocerebellar lesions impair adaptation. *Brain* 119:1183–1198
- Martin TA, Keating JG, Goodkin HP, Bastian AJ, Thach WT (1996b) Throwing while looking through prisms. II. Specificity and storage of multiple gaze-throw calibrations. *Brain* 119:1199–1211

- Miall RC (1989) Internal representation of human arm movements in visually guided tracking (abstract). *J Physiol (Lond)* 420: 61P
- Miall RC (1996) Task-dependent changes in visual feedback control: a frequency analysis of human manual tracking. *J Mot Behav* 28:125–135
- Miall RC, Weir DJ, Stein JF (1985) Visuomotor tracking with delayed visual feedback. *Neuroscience* 16:511–520
- Miall RC, Weir DJ, Wolpert DM, Stein JF (1993) Is the cerebellum a Smith predictor? *J Mot Behav* 25:203–216
- Neilson PD, Neilson MD, O'Dwyer NJ (1988) Internal models and intermittency: a theoretical account of human tracking behaviour. *Biol Cybern* 58:101–112
- Pew RW, Duffendack JC, Fensch LK (1967) Temporal limitations in human motor control. Psychonomic Society, Human Performance Centre, University of Michigan
- Schmidt RA (1991) Motor learning and performance: from principles to practice. Human Kinetics, Champaign, IL
- Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. *J Neurosci* 17:409–419
- Shadmehr R, Mussa-Ivaldi FA (1993) Motor adaptation to a virtual mechanical environment. Proc 1993 MIT CBCL Learning Day at Endicott, pp 11–12
- Shadmehr R, Brashers-Krug T, Mussa-Ivaldi FA (1995) Interference in learning internal models of inverse dynamics in humans. In: Tesauro G, Touretzky DS, Leen TK (eds) *Advances in neural information processing systems*, vol 7. MIT, Cambridge, MA, pp 1117–1124
- Smith OJM (1959) A controller to overcome dead time. *ISA J* 6:28–33
- Thach WT, Goodkin HP, Keating JG (1991) Inferior olive disease in man prevents learning of novel eye-hand synergies. *Soc Neurosci Abstr* 17:1380
- Thach WT, Goodkin HP, Keating JG (1992) The cerebellum and the adaptive coordination of movement. *Annu Rev Neurosci* 15:403–442
- Vercher J-L, Gauthier GM (1992) Oculo-manual coordination control: ocular and manual tracking of visual targets with delayed visual feedback of the hand motion. *Exp Brain Res* 90:599–609
- Viviani P, Mounoud P (1990) Perceptuomotor compatibility in pursuit tracking of two-dimensional movements. *J Mot Behav* 22:408–443
- Warwick MJ (1949) Effect of transmission-type control lags on tracking accuracy. USAF Air Material Command, Tech Rep 5916. Wright-Patterson Air Force Base, Dayton, Ohio