The CNS updates its context estimate in the absence of feedback

Philipp Vetter[CA] and Daniel M. Wolpert

Sobell Department of Neurophysiology, Institute of Neurology, University College London, Queen Square, London WC1N 3BG, UK

[CA]Corresponding Author

Human motor behaviour is remarkably accurate and appropriate even though our bodies and the objects we interact with change over time. To achieve such performance the motor system has to tailor the motor commands to the current context, that is the properties of objects in the world and the prevailing environmental conditions. The current context can be estimated by integrating two sources of information, sensory feedback and knowledge about how the context is likely to have changed from the previous estimate. Here we show that in the absence of sensory feedback about the context the second process is able to extrapolate the likely evolution of the context without requiring awareness that the context is changing. *NeuroReport* 11:3783–3786 © 2000 Lippincott Williams & Wilkins.

Key words: Context estimation; Psychomotor performance; Tracking; Virtual reality; Visual feedback

INTRODUCTION

Our ability to generate accurate and appropriate motor behavior relies on tailoring our motor commands to the prevailing movement context. This context embodies parameters of both our own motor system, such as the level of muscle fatigue, and the outside world, such as the weight of a bottle to be lifted. As the consequence of a given motor command depends on the current context, the CNS has to estimate this context so that the motor commands can be appropriately adjusted to achieve accurate control.

A general framework relating the context evolution and sensory feedback is shown in Figure 1. At a given point in time, t, there exists a true context C_t. However, the CNS does not have direct access to the true context, but can gain information about it through sensory feedback, S_t. As sensory transduction is noisy, sensory feedback depends probabilistically on the current context with a distribution given by $P(S_t|C_t)$. Also, the context is not static but tends to evolve in a structured way. For example the weight of a bottle tends to decrease as we pour from it, and muscles fatigue with use. Here, we represent this evolution as a probabilistic transition which relates the current context to the previous context, $P(C_t|C_{t-1})$. Although we frame this in a probabilistic manner, it also encompasses the case where the context evolves deterministically.

Therefore, in context estimation the CNS's task is to estimate the context from a (noisy) sequence of sensory feedback $S_0...S_t$. Within this framework this task can be solved optimally by internally modeling both the probability of the sensory information received for each possible context and the likely transitions between contexts [1]. Previous studies have shown that the motor system can

Fig. 1. Context evolution with sensory feedback as a hidden Markov process. The top row shows the evolution of the true context C, for example the weight of a bottle as we pour from it, which is hidden from the CNS, but must be estimated for accurate control. The bottom row, the sensory feedback S observed by the CNS enables it to infer the context. At each time t the current context is determined by the previous context with a fixed transition probability given by $P(C_t|C_{t-1})$. As sensory transduction is a noisy process the sensory feedback has a probability distribution that depends on the current context $P(S_t|C_t)$.

adapt to multiple contexts, such as altered relationships between actual and perceived finger position [2,3] and is able to estimate the appropriate context based on the state of the body [4–10]. In this case sensory feedback about the state of the body provides the main information as to the current context, i.e. using $P(S_t|C_t)$.

Recently, we have extended these studies by showing that subjects model both the likelihood of sensory information $P(S_t|C_t)$ and also the likely evolution of the context $P(C_t|C_{t-1})$ [1]. The latter was demonstrated by exposing subjects to two discrete contexts and transitioning between them randomly. Here, we are interested in how the CNS models the likely evolution of the context, when the
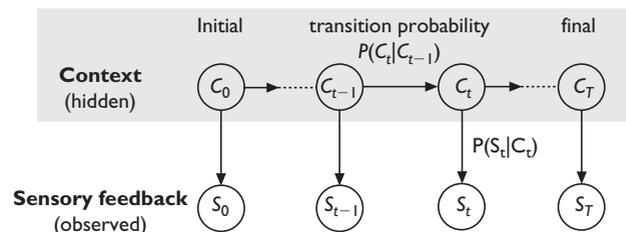
relationship between perceived and actual finger position, that is the context, changes continuously and deterministically over time. Subjects tracked a target which moved in a horizontal circle (Fig. 2a). Using virtual reality we introduced a time-varying sinusoidal discrepancy between the actual and displayed height of the finger. At various times during tracking the visual display of the finger was extinguished in order to assess how the actual height of the finger changed, as a measure of the subjects' estimate of the current context.

Figure 2b–e shows four a priori hypotheses which we considered. The solid line shows the time-varying context and the dotted line shows the predicted finger height when the feedback was extinguished at four possible points in the cycle. First, subjects may just stick with the context they experienced last (Fig. 2b). Another possibility, is that subjects tend towards the average context (Fig. 2c). A more sophisticated mechanism would involve extrapolating the rate of change of the context when visual feedback is withheld (Fig. 2d). Finally, the CNS may be able to extract the rate of change of context and its higher order derivatives and/or decay towards the mean context (Fig. 2e).

MATERIALS AND METHODS

Apparatus: Subjects made unconstrained movements to visual targets under stereoscopic visual feedback (for full details of the system see [1]). The three-dimensional virtual visual feedback setup allows visual feedback of the finger location to be overlayed onto the actual finger location while preventing subjects from seeing their arm. The target was displayed as a red 3.5 mm radius sphere and finger location as a blue 5 mm sphere whose position was updated at 50 Hz. Computer-controlled discrepancies could be introduced between the actual finger position, sampled with an Optotrak 3020 at 50 Hz, and the displayed finger

position. The virtual visual feedback system also allows the timing of visual feedback to be controlled by either displaying or hiding the virtual finger.

Paradigm: Sixteen right-handed subjects who participated in experiment 1 ($n = 8$) and experiment 2 ($n = 8$) gave their informed consent and were naive to the purpose of the experiment. Subjects initiated trials by moving their finger to the target. As soon as the the finger had come within 1 cm and had slowed to < 1.5 cm/s, the target started moving around a circle in the horizontal plane (8 cm radius, centre 24 cm in front of and 41 cm below subject's eyes) at a velocity of 8 cm/s. To improve accuracy, this circle was displayed throughout the experiment. The starting position of the target was pseudorandomly chosen at one of six equally spaced locations on the circle. Subjects were asked to track the target, regardless of whether or not they received visual finger position feedback. Subjects' finger feedback was offset vertically by $h \sin(\pi t/4)$, where h is the maximum perturbation and t is the time in seconds. The duration of a perturbation cycle is thus 8 s, which differs from the 6.3 s taken to go around the circle. The perturbation cycle was divided into four 2 s segments, at phases 0, $\pi/2$, π and $3\pi/2$. Subjects always received full visual finger feedback for at least five segments and then feedback was withheld for the duration of one segment. On each trial, feedback was withheld during each of the segments twice in a pseudorandom order, which differed across trials and subjects. Note that the spatial location at which feedback was removed provided no information about the phase of the perturbation cycle.

The experiment consisted of six trials with $h = 1.5$ cm and six trials without a perturbation ($h = 0$ cm) in a counter-balanced order across subjects. A rest period was interspersed after six trials.
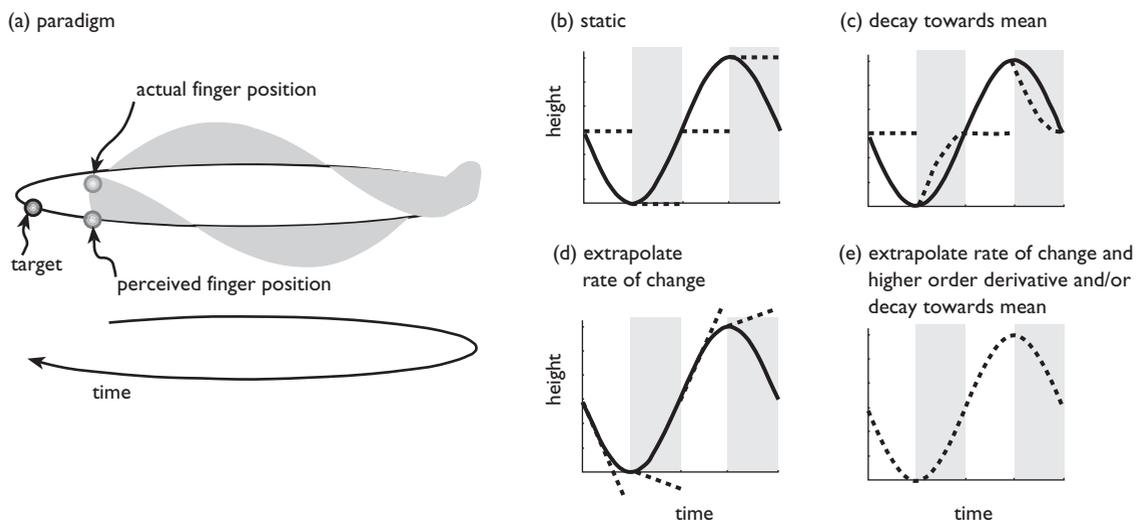

Fig. 2. (**a**) Paradigm. The task is to track the target moving around the circle, while the relationship between the perceived and actual finger position is altered sinusoidally over time. Note that the sinusoidal frequency is exaggerated for clarity. (**b**–**e**) Different hypotheses about tracking behaviour when feedback is withheld. Solid lines show the ideal behaviour in the presence of feedback, while the dashed lines indicate the performance expected when feedback is suddenly at four points of the cycle. The grey bars indicate the duration of each quarter cycle. (b) In the simple static case, the CNS expects the context to remain constant, as what it last experienced. (c) Here the CNS is sensitive to the mean context, to which it converges gradually. (d) The CNS extracts and extrapolates the rate of change of the context. (e) The CNS extracts the rate of change of context as well as higher order derivatives thereof and/or decays towards the mean.

Analysis: We analysed subject's finger height, as a measure of the subject's current context estimate [10]. For each trial, the time series of finger height was split into four quarter cycle segments, each aligned to the phase (0, $\pi/2$, π and $3\pi/2$) of the perturbation cycle. For each subject, segment and condition (finger seen/hidden × perturbation on/off), we averaged the time series over repetitions. Finally, a group average for each condition, along with a 95% confidence interval, was calculated by averaging these data ($n = 8$) over subjects.

Based on the results of the first experiment, we conducted a second experiment that was identical, except that the cycle was divided into two 4s segments, starting at phases 0 or π, respectively.

RESULTS

All subjects found the task easy and none noticed anything unusual on questioning. On further questioning they did not realize that a discrepancy had been introduced between their actual and displayed finger position. In the absence of a perturbation, subjects' tracking was accurate (Fig. 3a), and when feedback was withheld, there was a small, but significant tendency for the finger height to increase (Fig. 3b). In the presence of the sinusoidal perturbation, subjects tracked the target with a small lag (Fig. 3c) 460 ms (gain 0.80). When feedback was withheld for a quarter of a cycle (Fig. 3d) subjects' performance was similar, albeit more noisy, than with feedback (lag 660 ms; gain 0.76). Comparing this with the a priori hypotheses, it is clear that subjects can at least extrapolate the rate of change of context (compare Fig. 2b–e with Fig. 3d). This can be concluded because when feedback was turned off at the beginning or halfway through the cycle, when there was no net perturbation, subjects continued to point downward or upward, respectively. Since feedback decay towards the mean when feedback is withheld 1/4 or 3/4 through the cycle, this suggests that the CNS is also sensitive to the mean context and/or higher derivatives of the rate of change in the context.

Two cues could drive such learning. First, the CNS may extract information from the continuously varying context to learn how to extrapolate from this estimate, when feedback is withheld. However, a different cue could also drive learning in our experiment, since there was always a difference of 1.5 cm between the perturbation at the time when feedback was extinguished and when feedback was turned back on again. This difference could potentially serve as a cue to drive learning. To examine which of these cues was more salient, we conducted a second experiment that was identical to the first, except that visual finger feedback was only turned off and back on again when there was no net perturbation, that is at the beginning or halfway through the cycle. Therefore, if the CNS's default strategy is not to extrapolate the context (Fig. 2b) or to simply decay towards the mean (Fig. 2c), then there is no visual error between expected and perceived finger feedback to induce a change of strategy. As in the first experiment, subjects' performance was accurate in the condition when there was no perturbation (Fig. 4c,d), while compensating for the sinusoidal perturbation in the presence of feedback (Fig. 4a). More interestingly, Fig. 4b shows that even when there is no visual error signal, subjects still extrapolate the rate of change of context. This suggests that their learning was based on the continuous feedback rather than the possible discrepancy in perturbation between the removal and reappearance of feedback. Note that Fig. 4b also shows how the initial direction taken when feedback is turned off reverses, similar to when feedback is available.

DISCUSSION

Our results show that subjects actively compensate for a time-varying sinusoidal perturbation in the absence of visual feedback without being aware of the perturbation. The compensation strategy used suggests, that the CNS is
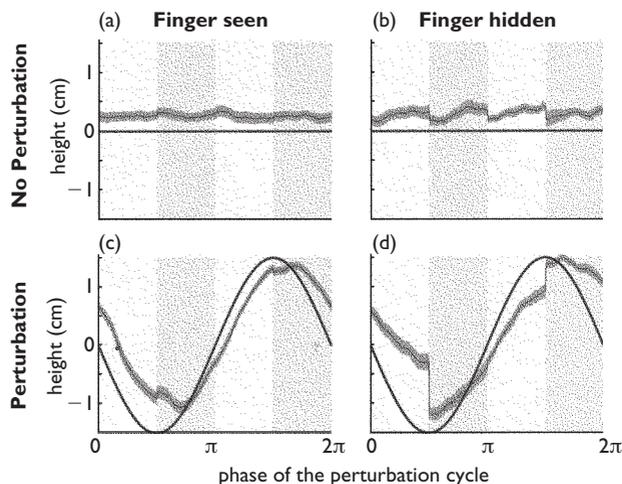

Fig. 3. Tracking performance (grey) and ideal tracking performance (heavy black). The data are averaged over subjects ± 95% confidence limits for each of the four conditions and each quarter of the perturbation cycle. The duration of each quarter cycle is indicated by the grey bars.
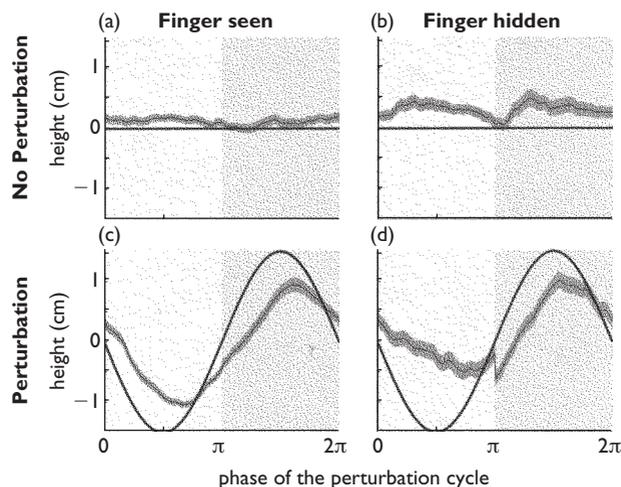

Fig. 4. Tracking performance (grey) and ideal tracking performance (heavy black). The data are averaged over subjects ± 95% confidence limits for each of the four conditions and each half of the perturbation cycle. The grey bar indicates the duration of a half cycle.

able to extract the rate of change of the context. Furthermore, the CNS may be able to estimate higher order derivatives of the rate of change in context or be sensitive to the mean context. To investigate whether this behaviour is due to extrapolation of the current context estimate or due to learning driven by the difference in perturbation amplitude between when feedback is extinguished and when it is turned back on again, we conducted a second experiment, in which the perturbation at feedback removal and reappearance were identical. Subjects still updated their estimate similarly, indicating that performance is based on extrapolation.

In terms of the context estimation framework of Fig. 1, our results suggest that the CNS generates an estimate of how the context is likely to evolve over time $P(C_t | C_{t-1})$. This implies that the CNS can construct and use an internal model: a brain process simulating the behavior of both the body and the outside world [11–15]. Previous studies have shown that internal models play a role in maintaining accurate control in the presence of sensory feedback delays [11], generating anticipatory responses [16–18] and in distinguishing our own actions from externally produced stimuli [19,20]. Here we propose that an internal model exists to predict how the context is likely to evolve over time. Consistent with this proposition, a recent study [1] found evidence that the transition probability is sensitive to the average context. In the experiment the context, the relationship between actual and perceived finger position, changed randomly between two states, and subjects behaviour in the absence of feedback decayed towards the mean. Conditt and Mussa-Ivaldi [21] exposed subjects to time-dependent force fields. From the lack of generalization, they concluded that the CNS has no explicit representation of time. These findings are not necessarily incompatible with ours. First, their experiment manipulated dynamic aspects of the context, while we used kinematic perturbations. It has been shown that internal models for dynamics and kinematics can be learned independently [22].

Conditt and Mussa-Ivaldi [21] used bell-shaped force perturbations that lasted 333 ms, while an equivalent half-cycle perturbation in our experiments was roughly 12 times slower. Our results can also be seen on the background of a large tracking literature. It emerges that compensatory tracking (as in our experiment) is more difficult than pursuit tracking, and is also more difficult to extrapolate in the absence of visual feedback (for a review see [23]). It is therefore quite surprising that subject's performance was so good, given that they were unaware of the perturbation, and tracking was in three dimensions. Extrapolation of context estimates has also been explored

in smooth pursuit eye movements, which interestingly cannot be elicited voluntarily. Barnes et al. [24] showed that subjects can sustain sinusoidal smooth pursuit eye movements in the absence of target motion, by stabilizing the target on the fovea and giving subjects appropriate attentional cues. While the authors interpret this as evidence for an short term memory trace of eye movements, the results are equally compatible with the framework of Fig. 1.

CONCLUSION
We have shown that the CNS may explicitly represent time in order to actively update its context estimate. This process does not require awareness, suggesting that the CNS may be continuously adapting motor behaviour to changing contexts without tapping attentional resources.

REFERENCES
1. Vetter P and Wolpert D. *J Neurophysiol* **84**, 1026–1034 (2000).
2. Kravitz JH and Yaffe F. *Percept Psychophys* **12**, 305–308 (1972).
3. Welch R, Bridgeman B, Anand S and Browman K. *Percept Psychophys* **54**, 195–204 (1993).
4. Cohn J, DiZio P and Lackner J. *J Neurophysiol* **83**, 3230–3240 (2000).
5. Kohler I. *Proc Aust Acad Sci* **227** (1950).
6. Hay JC and Pick HL. *J Exp Psychol HPP* **72**, 640–648 (1966).
7. Shelhamer M, Robinson DA and Tan HS. *Ann NY Acad Sci* **656**, 889–891 (1991).
8. Baker JF, Perlmutter SI, Peterson BW *et al. Exp Brain Res* **69**, 220–224 (1987).
9. Gandolfo F, Mussa-Ivaldi F and Bizzi E. *Proc Natl Acad Sci* **USA**, 93, 3843–3846 (1996).
10. Ghahramani Z and Wolpert DM. *Nature* **386**, 392–395 (1997).
11. Ito M. *The Cerebellum and Neural Control.* New York: Raven Press; 1984.
12. Kawato M, Furawaka K and Suzuki R. *Biol Cybernet* **56**, 1–17 (1987).
13. Wolpert DM, Ghahramani Z and Jordan MI. *Science* **269**, 1880–1882 (1995).
14. Wolpert DM. *Trends Cogn Sci* **1**, 209–216 (1997).
15. Merfeld D, Zupan L and Peterka R. *Nature* **398**, 615–618 (1999).
16. Forssberg H, Kinoshita H, Eliasson AC *et al. Exp Brain Res* **90**, 393–398 (1992).
17. Gordon AM, Westling G, Cole KJ and Johansson RS. *J Neurophysiol* **69**, 1789–1796 (1993).
18. Jenmalm P and Johansson RS. *J Neurosci* **17**, 4486–4499 (1997).
19. Jeannerod M. *The Cognitive Neuroscience of Action.* Oxford: Blackwell; 1997.
20. Blakemore S, Wolpert D and Frith C. *Nature Neurosci* **1**, 635–640 (1998).
21. Conditt M and Mussa-Ivaldi F. *Proc Natl Acad Sci USA* **96**, 11625–11630 (1999).
22. Krakauer J.W. *Ghilardi M and Ghez C. Nature Neurosci* **2**, 1026–1031 (1999).
23. Poulton EC. *Tracking Skill and Manual Control.* London: Academic Press; 1974.
24. Barnes GR, Goodbody S and Collins S. *Exp Brain Res* **106**, 301–317 (1995).