

Bayesian Integration in Force Estimation

Konrad P. Körding,¹ Shih-pi Ku,² and Daniel M. Wolpert¹

¹Institute of Neurology, Sobell Department of Movement Neuroscience, University College London, London WC1N 3BG, United Kingdom; and ²Max-Planck Institute for Biological Cybernetics, Department of Physiology of Cognitive Processes, 72076 Tübingen, Germany

Submitted 19 March 2004; accepted in final form 14 May 2004

Körding, Konrad P., Shih-pi Ku, and Daniel M. Wolpert. Bayesian integration in force estimation. *J Neurophysiol* 92: 3161–3165, 2004. First published June 9, 2004; 10.1152/jn.00275.2004. When we interact with objects in the world, the forces we exert are finely tuned to the dynamics of the situation. As our sensors do not provide perfect knowledge about the environment, a key problem is how to estimate the appropriate forces. Two sources of information can be used to generate such an estimate: sensory inputs about the object and knowledge about previously experienced objects, termed prior information. Bayesian integration defines the way in which these two sources of information should be combined to produce an optimal estimate. To investigate whether subjects use such a strategy in force estimation, we designed a novel sensorimotor estimation task. We controlled the distribution of forces experienced over the course of an experiment thereby defining the prior. We show that subjects integrate sensory information with their prior experience to generate an estimate. Moreover, subjects could learn different prior distributions. These results suggest that the CNS uses Bayesian models when estimating force requirements.

INTRODUCTION

Sensorimotor tasks in the real world are inherently uncertain, making it necessary to estimate variables that are important to task requirements. Consider picking up a full glass of Pimms.¹ To apply the appropriate level of force, we need to know, among other parameters such as the frictional property of the surface, how heavy the glass is. Computationally we are faced with the task of estimating its weight. What information could we use for this estimation? We can generate an estimate based on sensory information from vision and, if we had recently lifted the glass, from tactile inputs. By combining information from these different modalities, the accuracy of the estimate can be improved (Ernst and Banks 2002; Hillis et al. 2002; Jacobs et al. 1991; Van Beers et al. 2002). However, such an estimate can be further improved by using prior knowledge about the possible distribution of the weight of the glass. Bayesian theory (Jaynes 1986; MacKay 2003) tells us how to combine this a priori information about the distribution of the weight with the evidence provided by sensory feedback to obtain an optimal estimate. This combination process requires prior knowledge, how probable each possible weight is, and knowledge of the uncertainty inherent in the estimate provided by our sensory inputs.

Several studies have examined the use of Bayesian integration in vision. In the visual system, the assumption that subjects use prior information can predict a range of visual illusions

based on Bayesian processing (Fleming et al. 2003; Kersten and Yuille 2003; van Ee et al. 2003). These studies assume a prior, for example, that the velocities of objects in the world have a Gaussian distribution with zero mean and can predict a range if visual illusions (Weiss et al. 2002). Recently we have shown that people can learn to use novel visually defined priors and visual feedback to improve their performance in the context of sensory uncertainty (Körding and Wolpert 2004). However, it has been suggested that people cannot learn the distribution of forces experienced but instead learn the average over a set of recent trials (Scheidt et al. 2001). Here we examine whether subjects could learn to use prior knowledge when estimating force with two differences from the Scheidt et al. (2001) study. First, subjects had more extensive training on the task. Second, the task was constructed so that sensory feedback from the initial part of each trial could be used to determine the appropriate forces to apply on the second half of the trial.

We used a virtual reality setup with a force field generated by a robot arm to test if subjects combine prior knowledge with sensory feedback to estimate force levels. During each trial, we controlled the force experienced and measured subjects' estimate of the force. We thereby controlled the distributions of forces experienced and show that this influences performance in the way predicted by Bayesian integration. We furthermore show that people can learn different distributions one after the other.

METHODS

On each trial, subjects experience two force pulses and were required to accurately oppose the effects of the second pulse. For each trial, the amplitude of the force pulses was randomly drawn from a probability distribution $p(F_{\text{true}})$. To perform accurately, subjects needed to estimate the magnitude of the first force pulse so as to apply the appropriate counteracting force during the second pulse. We analyzed the systematic errors that subjects make when estimating the force and how the probability distribution influenced these errors.

After providing written informed consent, 11 right-handed subjects (6 male, 5 female, aged 22–40) participated in this study. A local ethics committee approved the experimental protocols. While seated, subjects held a robotic manipulandum (Phantom Haptic Interface 3.0, Sensible Devices) that allowed us to control the forces they experienced. A virtual-reality system was used that prevented subjects seeing their hand and allowed us to present visual images into the plane of the movement (for full details of the setup, see Goodbody and Wolpert 1998).

Each trial started with the manipulandum producing no force, and subjects moved the manipulandum to a visual starting point that was

Address for reprint requests and other correspondence: K. P. Körding, Institute of Neurology, Sobell Dept. of Movement Neuroscience, UCL London, UK (E-mail: konrad@koerding.de).

¹A delicious English gin-based drink usually served mixed with lemonade and chopped up cucumber and fruit.

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "advertisement" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

24 cm left of midline, ~18 cm ahead and 45 cm below the subject’s eyes. During this phase, subjects saw a white sphere of 0.25 cm diam at the position of their hand. Once on the starting point, subjects saw a line representing the position of their hand along the x axis (transverse axis), but no feedback was provided about the position along the other axes. The hand was moved by the robot to the right along a transverse axis at a constant speed. The overall movement of 40 cm took 1,600 ms. To achieve a constant speed in the x direction, the robot applied a spring-like force that was proportional to the deviation from the current desired position with a spring constant of 600 N/m (this was ramped up linearly over the 1st 6 cm). Subjects were instructed to move their hand along the straight line so as not to resist the robot. Their task was to experience a first force pulse and counteract a second so that their cursor passed as close as possible to the final target (Fig. 1A).

Over the course of this movement, subjects sequentially experienced three force fields acting in the saggital direction (Fig. 1B). First, a smoothly varying force pulse was applied that varied with the distance, x measured in centimeters, that the hand had moved from the starting location along the transverse axis. This force $F_y = F_{\text{true}}\{1 + \cos[\pi(x - 12)/4]\}$ was applied when x was between 8 and 16 cm (corresponding to 320–640 ms from the start of the trial). As this force could perturb the hand off the straight line between starting location and target, to ensure the second force pulse was always experience from a similar initial state, a second force was applied to bring subjects back to the straight line at the midpoint of the movement. To achieve this, a spring-like force $F_y = -(x + 4)y$ was applied that increased in strength as the transverse distance moved, x , increased from 20 and 28 cm. This force acted like a funnel to bring the hand back to the horizontal line. Finally, a second force pulse was

applied when the horizontal displacement x was between 32 and 40 cm (1,280–1,600 ms)

$$F_y = F_{\text{true}}(1 + \cos(\pi(x - 36)/4)) + F_{\text{pert}}$$

where F_{pert} is a small distortion discussed in the following text. The target was located on the horizontal line at 36 cm (1,440 ms). This corresponds to half way through the second short pulse (160 ms into the pulse) to minimize correction from haptic feedback during this pulse. As the hand passed, the target they were shown for 200 ms the position of their hand as a white sphere. The discrepancy between the hand and target positions at this point is called Δy . After they had passed the 40-cm point, the movement finished and the trial ended. Subjects were instructed that on each trial the two force pulses would be identical so that they could use the size of the first force pulse to estimate and compensate for the second. To achieve this, they had to compensate for the second force pulse. Before starting the experiment, subjects were familiarized with the apparatus and the experiment.

The magnitude of the forces F_{true} was randomly drawn each trial from a Gaussian distribution $p(F_{\text{true}})$ with a mean μ_{prior} of 2 N and a SD that depended on the experimental condition. On a given day, each subject either experienced a “narrow distribution” in which the SD, σ_{prior} , was 0.5 N or a “wide distribution” in which the SD, σ_{prior} , was 1 N. All subjects performed the experiment on three different days. They experienced one distribution for the first 2 day and then the second distribution for the 3rd day. On each day, the experimental session lasted ~2 h with between five and seven blocks each day with each block containing 200 trials. Six subjects experienced the wide distribution first and five the narrow distribution first.

Measuring the estimated forces

To investigate the force estimation process, we wish to know the magnitude of the force that subjects estimated for the first force pulse, which we term $F_{\text{estimated}}$. We cannot directly measure this force estimate but can measure the positional error, Δy , during the second force pulse, which is related to the force estimation error. If subjects perfectly estimate the first force, they should be able to compensate for the second force pulse, and the positional error should be zero. Any inaccuracy in the estimate will lead to a positional error. The magnitude of this positional error should be related to the size of the force estimation error. To determine the relationship between force estimation errors and positional errors, we added a small perturbation F_{pert} to the size of the second force pulse magnitude, which was drawn from a zero-mean Gaussian distribution with SD of 0.2 N for the wide distribution and 0.1 N for the narrow distribution. As this force perturbation was added to the force drawn from the prior distribution, it had the same time course and shape as the second force pulse. The variance of this additional perturbation was chosen so as to be small compared with the overall force pulse. On postexperimental questioning, subjects were unaware of any difference in magnitude between the first and second force pulses. As this perturbation was unpredictable and the duration of the second force pulse is short (160 ms) so that subjects cannot compensate for errors that arise during the pulse, we can treat the arm as a passive virtual mass during this phase. As the time is short, the distances and speeds in the y direction are very small, and we can also assume that frictional forces that are proportional to velocities are small. Errors in force compensation should thus translate proportionally into errors in position with a proportionality constant c . The basic assumption is that there exists a linear relation between F_{pert} and spatial errors (such as captured by a mass-spring-damper system). By examining the positional error induced by the additional perturbation, we can relate positional errors to force errors. We model the error as $\Delta y = c(F_{\text{true}} - F_{\text{estimated}} + F_{\text{pert}})$ and use a least-squares fitting procedure to estimate c and $F_{\text{estimated}}$ from the measurement data. This allows us to obtain $F_{\text{estimated}}$ from the positional errors made by the subjects.




FIG. 1. Experimental design. A: the visual feedback is sketched. B: the perturbation forces are sketched. C: a typical hand trajectory is shown. The arrows describe the acting forces. D: the influence of the force difference ΔF on the endpoint error Δy is shown for 1 subject experiencing the narrow distribution (left) and for all subjects (right). E: the fit slope is shown for subjects that experienced the wide distribution (black) and those that experienced the narrow distribution (gray).

Computational models

Several computational models predict different optimal strategies and thus make different predictions how the estimation error $\Delta F = F_{\text{true}} - F_{\text{estimated}}$ should depend on the actually experienced force F_{true} (Fig. 2A):

Model 1: Naïve compensation

Subjects could ignore the prior distribution and fully compensate for the perceived force. In this case, they would just produce whichever force they sensed during the first force pulse. Assuming no bias in sensory perception, the error should on average always be zero, independent on the force: $E[\Delta F] = F_{\text{true}} - E[F_{\text{estimated}}] = 0$.

Model 2: Full Bayesian compensation

If subjects use an optimal Bayesian strategy, we can calculate the optimal estimated force ($F_{\text{estimated}}$) of the current trial given a perceived force ($F_{\text{perceived}}$) and the prior distribution of the forces $p(F_{\text{true}})$. If the perceived force has a sensory uncertainty character-

ized by a SD of σ_{feedback} , then the optimal strategy is (see Körding and Wolpert 2004)

$$F_{\text{estimated}} = \frac{\sigma_{\text{feedback}}^2}{\sigma_{\text{prior}}^2 + \sigma_{\text{feedback}}^2} \mu_{\text{prior}} + \frac{\sigma_{\text{prior}}^2}{\sigma_{\text{prior}}^2 + \sigma_{\text{feedback}}^2} F_{\text{perceived}}$$

If we assume that the expected value of the perceived force is unbiased and therefore the same as the true force, this predicts that on average the error would be

$$\begin{aligned} E[\Delta F] &= F_{\text{true}} - E\left[\frac{\sigma_{\text{feedback}}^2}{\sigma_{\text{prior}}^2 + \sigma_{\text{feedback}}^2} \mu_{\text{prior}} + \frac{\sigma_{\text{prior}}^2}{\sigma_{\text{prior}}^2 + \sigma_{\text{feedback}}^2} F_{\text{perceived}}\right] \\ &= F_{\text{true}} - \frac{\sigma_{\text{feedback}}^2}{\sigma_{\text{prior}}^2 + \sigma_{\text{feedback}}^2} \mu_{\text{prior}} - \frac{\sigma_{\text{prior}}^2}{\sigma_{\text{prior}}^2 + \sigma_{\text{feedback}}^2} \underbrace{E[F_{\text{perceived}}]}_{= F_{\text{true}}} \\ &= \frac{\sigma_{\text{feedback}}^2}{\sigma_{\text{prior}}^2 + \sigma_{\text{feedback}}^2} (F_{\text{true}} - \mu_{\text{prior}}) \end{aligned}$$

Therefore the error is proportional to the difference of the actual force and the mean of the prior distribution. Furthermore, the slope should

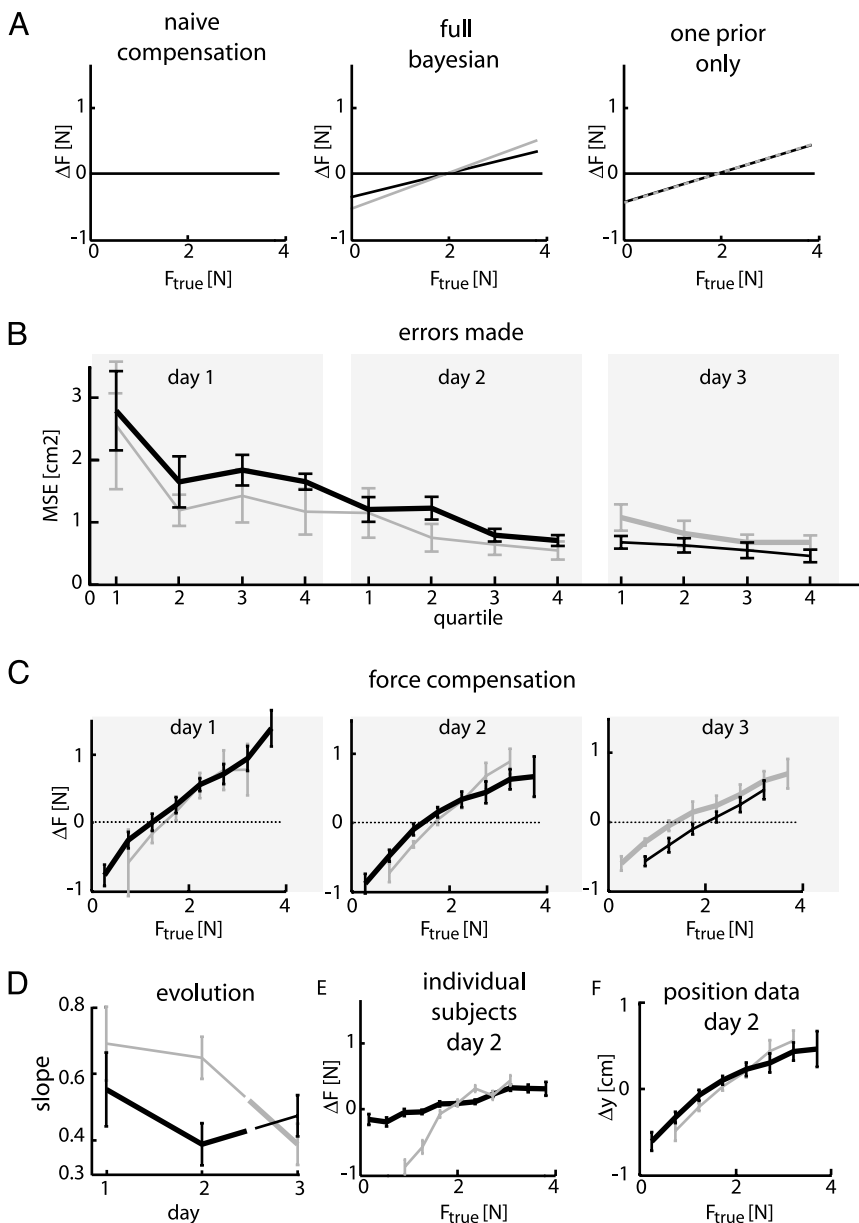


FIG. 2. Computational models and discerning data. Black denotes subjects that started with a distribution of SD = 1 N and gray stands for subjects that started with SD = 0.5 N. All error bars denote the SE. A: the predictions from the 3 computational models are sketched. B: the MSE of the estimated force is plotted against time. C: the average force estimation error is plotted against the true force for the population of subjects. D: the average slope of the curves shown in C is plotted as a function of the day. E: for 2 different subjects, the average bias when estimating the force is plotted against the true force. F: the endpoint error averaged over subjects is plotted against the true force.

increase with decreasing variability of the prior. Although subjects systematically deviate from the target in this strategy, by reducing variance, this always leads to a smaller mean squared error (MSE) than the naïve compensation strategy (Körding and Wolpert 2004).

Model 3: Bayesian integration with one fixed prior

Subjects could just use one prior that is either genetically pre-defined or acquired. In this case, the error would be proportional to the mean of the distribution, but the slope would be constant and not depend on the distribution that was actually used. There should thus be no difference between average performance when subjects experience the narrow and wide distributions.

RESULTS

Subjects were required to perform a force-matching task in which they experienced a force pulse in the first half of each trial the amplitude of which was drawn randomly from a prior distribution and had to compensate for the same force pulse during the second half of the trial. They received feedback on how well they compensated for the force pulse. Each subject was tested with two Gaussian prior distributions of forces that differed only in their SD.

Over the course of the experiment, subjects reduced their mean squared positional error at the end of the movement (Fig. 2B). The errors rapidly decreased over the first day and approached an asymptote by the end of the second day. On the third day, subjects experienced a different prior distribution. For the group who first experienced the narrow distribution (Fig. 2B, gray), the errors rose on introduction of the wide distribution. This effect was expected as the forces now have a larger range. Subjects from this group subsequently decreased their error on the third day. The overall decrease in positional error shows that the subjects are able to improve their performance in the face of uncertainty about the level of force.

To compare subjects behavior to the models (Fig. 2A), we need to determine the force estimation error (ΔF) that subjects make as a function of the true force (F_{true}). To do so, we measure the proportionality constant c between the small additional force F_{pert} imposed on the second force pulse and the positional error Δy using a least-squares fitting procedure (see METHODS). A linear relation can be seen when plotting the average Δy against F_{pert} averaged over all trials both for a typical subject (Fig. 1D, left) and also over the population (Fig. 1D, right). To check that subjects did not change the dynamics of their arm, we plot the proportionality factor for both groups for the second and third days of the experiment (Fig. 1E). The value stays approximately constant over the course of training. Using this proportionality constant we can convert positional errors into force estimation errors.

We plot the average force estimation error ($\Delta F = F_{\text{true}} - F_{\text{estimated}}$) against the presented force F_{true} (Fig. 2C) for both groups and for each of the 3 days. The force estimation error is approximately proportional to the presented force as predicted by the optimal Bayesian strategy (Fig. 2A, middle). There is a bias seen at the mean of the prior, which is not predicted by Bayesian statistics alone. Subjects seem to systematically underestimate the true force. Such conservative force production could arise from a process that not only tries to minimize the spatial errors but also reduces the necessary effort. Such a trade-off can be captured by a loss function that includes a

spatial and a force term. Such phenomena can be modeled within the more general framework of Bayesian decision theory. Furthermore, the slope in the case of the wide distribution (thick lines) is smaller after extensive learning (day 2, Fig. 2D) as predicted by the model where priors are tuned to the distribution. Performing linear regression for each subject's data on day 2 gives slopes that range from 0.16 to 0.89 and as a group are significantly greater than zero ($P < 0.0001$, t -test), showing that subjects do not perform naive compensation (slope = 0).

If we assume that people use Bayesian statistics, we can use the slopes of the curves in Fig. 2D to analyze the level of uncertainty that subjects have in their feedback using the equations derived for the Bayesian model. We assume that subjects have correct knowledge about the distribution and put in the true prior. The fitted σ_{feedback} is 0.81 ± 0.10 (SE) N is the same for both groups ($P > 0.12$, paired t -test). We can then use this estimate along with the equation defining the optimal strategy to predict the ratio of slopes to be 0.55, which is close to the measured value of 0.60 on day 2.

The slope of the relationship between estimation error and actual force changed over time (Fig. 2D). On the first day, there was no significant difference in slope between the two groups. A comparison of the slopes for the second day of the experiment showed that the slopes were significantly different ($P < 0.04$, $n = 5, 6$, t -test) with the group who first experienced the narrow distribution having a steeper slope than the group who first experienced the wide distribution. Data for one subject from each group on day 2 is shown in Fig. 2E. This difference can also be seen in the subjects' raw positional data taken from day 2 (Fig. 2F). This shows that subjects can tune their performance appropriately for the different priors. In addition, subjects relied more on the prior, as indicated by a steeper slope, when it was more informative (narrow distribution), as predicted by the Bayesian strategy. On the third day, when distributions were swapped, the subjects that were faced with a narrower distribution learned to depend more on the feedback and significantly reduced the slope ($P < 0.03$, paired t -test). However, there was no significant change in slope for the group in which the distribution narrowed on day 3. However, on the third day the slopes of the two populations crossed over. This change over time showed that subjects that learned a narrow distribution can change their behavior in response to a wider distribution.

Previous studies examining Bayesian learning have focused on the visual system, either assuming a prior (e.g., Weiss et al. 2002) or applying a novel prior (Körding and Wolpert 2004). Here we have shown that learning such a prior is not limited to visual tasks. When a novel prior over experienced forces is introduced, subjects make use of this prior information to optimize their performance in a way consistent with Bayesian integration. Moreover, their learning was specific to the distribution experienced.

ACKNOWLEDGMENTS

We thank the Gatsby Computational Unit for inspiring discussions and J. Ingram for technical support.

GRANTS

This work was supported by the Wellcome Trust, McDonnell Foundation, and Human Frontiers Science Programme.

REFERENCES

- Ernst MO and Banks MS.** Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415: 429–433, 2002.
- Fleming RW, Dror RO, and Adelson EH.** Real-world illumination and the perception of surface reflectance properties. *J Vis* 3: 347–368, 2003.
- Goodbody SJ and Wolpert DM.** Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79: 1825–1838, 1998.
- Hillis JM, Ernst MO, Banks MS, and Landy MS.** Combining sensory information: mandatory fusion within, but not between, senses. *Science* 298: 1627–1630, 2002.
- Jacobs RA, Jordan MI, Nowlan SJ, and Hinton GE.** Adaptive mixtures of local experts. *Neural Computat* 3: 79–87, 1991.
- Jaynes ET.** Bayesian methods: general background. In: *Maximum-Entropy and Bayesian Methods in Applied Statistics*, edited by Justice JH. Cambridge, UK: Cambridge Univ. Press, 1986, p. 1.
- Kersten D and Yuille A.** Bayesian models of object perception. *Curr Opin Neurobiol* 13: 150–158, 2003.
- Körding KP, and Wolpert DM.** Bayesian integration in sensorimotor learning. *Nature* 427: 244–247, 2004.
- MacKay DJC.** *Information Theory, Inference, and Learning Algorithms* Cambridge, UK: Cambridge Univ. Press, 2003.
- Scheidt RA, Dingwell JB, and Mussa-Ivaldi FA.** Learning to move amid uncertainty. *J Neurophysiol* 86: 971–985, 2001.
- Van Beers R, Baraduc P, and Wolpert D.** Role of uncertainty in sensorimotor control. *Transact Roy Soc* 357: 1137–1145, 2002.
- van Ee R, Adams WJ, and Mamassian P.** Bayesian modeling of cue interaction: Bistability in stereoscopic slant perception. *J Opt Soc Am A* 20: 1398–1406, 2003.
- Weiss Y, Simoncelli EP, and Adelson EH.** Motion illusions as optimal percepts. *Nat Neurosci* 5: 598–604, 2002.