VU Research Portal

Sensorimotor priors in non-stationary environments

Narain, D.; van Beers, R.J.; Smeets, J.B.J.; Brenner, E.

published in

Journal of Neurophysiology
2013

DOI (link to publisher)

[10.1152/jn.00605.2012](https://doi.org/10.1152/jn.00605.2012)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Narain, D., van Beers, R. J., Smeets, J. B. J., & Brenner, E. (2013). Sensorimotor priors in non-stationary environments. *Journal of Neurophysiology*, 109, 1259-1267. <https://doi.org/10.1152/jn.00605.2012>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Sensorimotor priors in nonstationary environments

Devika Narain, Robert J. van Beers, Jeroen B. J. Smeets, and Eli Brenner

MOVE Research Institute Amsterdam, Faculty of Human Movement Sciences, VU University, Amsterdam, The Netherlands

Submitted 13 July 2012; accepted in final form 6 December 2012

Narain D, van Beers RJ, Smeets JB, Brenner E. Sensorimotor priors in nonstationary environments. *J Neurophysiol* 109: 1259–1267, 2013. First published December 12, 2012; doi:10.1152/jn.00605.2012.—In the course of its interaction with the world, the human nervous system must constantly estimate various variables in the surrounding environment. Past research indicates that environmental variables may be represented as probabilistic distributions of a priori information (priors). Priors for environmental variables that do not change much over time have been widely studied. Little is known, however, about how priors develop in environments with nonstationary statistics. We examine whether humans change their reliance on the prior based on recent changes in environmental variance. Through experimentation, we obtain an online estimate of the human sensorimotor prior (prediction) and then compare it to similar online predictions made by various nonadaptive and adaptive models. Simulations show that models that rapidly adapt to nonstationary components in the environments predict the stimuli better than models that do not take the changing statistics of the environment into consideration. We found that adaptive models best predict participants' responses in most cases. However, we find no support for the idea that this is a consequence of increased reliance on recent experience just after the occurrence of a systematic change in the environment.

dynamic prior; Kalman filter; nonstationary environment; sensorimotor control

IN THE SPORT OF CRICKET, IT IS common practice for captains to introduce a slow-pitching bowler after the opponent batsman has faced a long spell of fast bowling. The assumption is that batsmen learn the statistics of the bowling distribution and rely on these to make predictions for the next ball. However, such learned statistics are rendered uninformative when the bowler is changed. The batsman should therefore best lower his reliance on past experience until he has learned the statistics of the new bowling distribution from fresh observations.

Past research indicates that learned variables may be stochastically represented as a priori distributions or priors (Knill and Richards 1996; Weiss et al. 2002; Körding and Wolpert 2004; Adams et al. 2004; Berniker et al. 2010; Ernst 2007; Knill 2007; Turnham et al. 2011; Verstynen and Sabes 2011). In this study, we are interested in priors that represent the prediction of the nervous system when engaged in an online estimation problem, like in the aforementioned cricket example. If the statistics of the environmental variable in question change over time, we expect to find the extent to which one relies on the prior pertaining to this variable to change as well. It is worth noting that the prediction or prior we refer to here is not a single value but assumed to be a probability distribution. The estimate of the prior refers to the mean of this distribution and the reliance that we refer to is the inverse of its variance. Unlike most previous work in this area, we assume that the distribution of the prior changes over time.

In the event of abrupt systematic changes, temporarily reducing one's reliance on the prior will result in reduced overall error. We assume that people determine the statistics of their environment based on previous observations, and we therefore design an experimental environment such that optimal performance is attained by relying on only a few past trials to determine the variance of the prediction. We test human participants' performance in this environment in which the average required response sometimes changes abruptly. Throughout the experiment, participants encounter interspersed trials in which sensory information about the stimuli is present in that trial or is absent. When sensory information is absent on a trial, the response represents a noisy readout of the mean of the participant's a priori prediction (prior) at that time. We thus obtain an online estimate of the mean of the sensorimotor prior and proceed to fit various models to understand how the system's dependence up priors is determined. If participants do not update their reliance on the prior after the occurrence of abrupt changes that cause a change in the mean of the observed environmental variable, we expect models like the Kalman filter (Kalman 1960) to closely predict such behavior. The Kalman filter in its standard implementation assumes that the variance of the prior and of the observation do not change over time and is therefore considered stationary. However, if humans were to use such a model, it is not clear how the brain would determine the value of its parameters. We propose that humans may be using a model that can estimate its own parameters recursively through an evaluation of the observed environment. Such a model that recursively estimates its own parameters based on the requirements of the observed environment is referred to as adaptive.

While the Kalman filter has been used to explain behavior in a variety of sensorimotor tasks, in some cases it has failed to predict observed human learning rates leading to the suggestion that humans may not always assume the statistics of the environment to be stationary (Burge et al. 2009). Other work has explored this issue using frameworks other than the Kalman filter, but these models have only successfully been applied to environments where no systematic changes occur in the mean of the environmental variable and only the variance is altered (Berniker et al. 2010). In the present work, we combine elements from these works and suggest a modification for the Kalman filter that would enable it to increase its reliance on newly observed information from the environment when systematic changes occur. Such an adaptive Kalman filter would enable faster responses to systematic changes compared with the stationary Kalman model. For the stationary Kalman filter, the mean value of the prediction is influenced by the mean of the previous trial's prediction but the variance of this prediction is invariant over time. In the adaptive Kalman models we propose, we additionally make the variance of the prediction depend on the variance of the observed environmental process over a window or exponent of past trials.

Address for reprint requests and other correspondence: D. Narain, Faculty of Human Movement Sciences, VU Univ., Van der Boechorststraat 9, 1081BT Amsterdam, The Netherlands (e-mail: d.narain@vu.nl).

In addition to these models, we test two adaptive models (Adams and MacKay 2007; Berniker et al. 2010) proposed in literature, which both address the problem of predicting variables in dynamic environments but do not use the Kalman framework. What interests us most in this study is whether adaptive models can provide a more convincing approximation to measured human behavior than a stationary model.

MATERIALS AND METHODS

Experimental paradigm. The Ethical Committee of the Faculty of Human Movement Sciences, VU University approved the program to which this study belongs. Eight healthy right-handed adult volunteers (6 naïve, 5 female) consented to participate in the study. There were no systematic differences in the data that distinguished the naïve and nonnaïve participants. Participants were seated in the dark at a setup with their hands resting on a horizontal tablet surface (Wacom Digitizer UD-1825-A, 45.7×63.5 cm) upon which they saw computer generated images (refresh rate: 85 Hz; resolution: $1,024 \times 768$ pixels) (Brenner and Smeets 2009). A stylus, lit at the tip by an LED, was used to make responses. Participants were asked to use the stylus to intercept a moving visual target (Gaussian blob with 10% peak contrast and 10-mm SD on a 12.6 cd/m^2 gray background) at a fixed location (point of interception) marked with a cross mark (Fig. 1A). This location was situated 150 mm further than the starting point along the midline of the participant's body, so all movements were in the sagittal direction moving away from the participant. The target was designed such that it would be detectable to all subjects but nevertheless would be quite difficult to see. The target always appeared to the left of the starting point and traversed a radial arc of 77.3° with radius 150 mm and the subject's starting point as the radial

center. The participant initiated the trial by moving to a start position. The color of the start position indicated the type of trial; if blue, the target would be visible during the trial (visible-target trial), and if orange, the target would move invisibly (invisible-target trial).

Participants were encouraged to compete with fellow participants based on the scores they obtained and were made aware of their ranking among other participants. Invisible target trials carried twice as many points as visible target trials, if intercepted. At the end of each trial, a graded score was awarded with a maximum of 100 points if the response was within 2 mm of the target center (along the arc), and fewer points were awarded as the error grew larger. The awarded points and cumulative score were displayed. Additionally, for visible-target trials, the position of the target when the stylus reached the point of interception was shown to make the error in the participant's timing explicit.

Within a trial, the target always moved at a constant angular velocity, but the velocity varied across trials. The time to interception of the stimulus was sampled from different distributions. To make the transition between blocks difficult to detect, the distributions were selected such that they overlapped (Fig. 1B). For the first 50 trials, the stimulus time to interception of the target was drawn independently from an almost noise-free Gaussian distribution (mean: 900 ms, SD: 1 ms). Thereafter, the experiment consisted of 5 blocks of 100 trials in which the stimulus time to interception was independently drawn from one of two Gaussian distributions (mean: 800 and 1,000 ms, SD: 50 ms). Each block consisted of trials drawn from one of these two distributions. Every third trial within a block was an invisible-target trial.

Analysis of responses. Responses that deviated by >2.5 mm from the indicated point of interception were discarded ($<1\%$). The time at which the stylus passed the point of interception was corrected for a

Fig. 1. A: time courses of the two types of trials. The marked point of interception is the cross mark. Dashed lines illustrate the path of the visible target and solid lines that of the stylus. B: example of target distribution over the course of the experiment. Stimulus times for visible-target trials (blue circles) were drawn from 1 of 2 Gaussian distributions (in gray). Invisible-target trials (orange squares) were treated (for feedback on performance) as if they arrived at the mean times of these distributions. For each subject, stimulus time-to-interception values were independently drawn from the same distribution. For the first 50 trials, we only plot 1 in 10 visible trials (blue circles) for clarity. The first 50 trials were drawn from a near-zero variance (SD: 1 ms) Gaussian distribution (not shown above).

constant delay in the tablet hardware (estimated to be 62 ms; Brenner and Smeets 2009; Di Luca 2011). The temporal error was determined as the time difference between when the target and stylus passed the point of interception.

Generative models. Our purpose here is to investigate whether adaptive models fare better than commonly utilized stationary models in explaining online data of the human prior in a nonstationary environment. To this end, we compared several types of algorithms and frameworks.

We began with one of the most prevalent computational frameworks for online state estimation of the sensorimotor system, the Kalman filter (Korenberg and Ghahramani 2002; Wolpert and Ghahramani 2000; Baddeley et al. 2003; Burge et al. 2008; Wei and Körding 2010; Izawa and Shadmehr 2011; Wolpert 1997; van Beers 2012). The Kalman filter combines the information from a noisy observation with its own a priori prediction based on past observation and its knowledge of the process dynamics (Kalman 1960). The aforementioned literature on applying the Kalman filter to sensorimotor paradigms assumes that the observation of the variable, and the prediction from the sensorimotor system, are combined in a definite proportion (Kalman gain). The Kalman gain is a function of the variance of the environmental process and the variance of the sensory measurement (or observation) of the process. It is often assumed that neither of these variances change over time. The measurement variance relates to inherent noise in a sensor apparatus, and it is thus reasonable to assume that it remains stationary.

If the nervous system does not alter its gain much over time, we expect that subjects will behave in accordance with the standard (or stationary) Kalman filter. On the other hand, if the nervous system is sensitive to local changes in the statistics of the environment, we expect adaptive models with gains dependent on the recent process to provide a better fit to the data. An adaptive model recursively estimates its parameters to make better predictions for future observations from the environment. We develop adaptive Kalman filters that recursively estimate the variance of the environmental process based on past observations and adjust the proportion (gain) by which they rely on their prediction and observation accordingly. These adaptive models estimate the environment's variance over a span of past observations. We propose two different adaptive models that differ in the dynamics of how they consider past information. The window model considers a discrete window of past trials to estimate process variance uniformly whereas the exponent model estimates the process variance by weighting all past trials exponentially. We designed the environment such that the theoretically optimum window size of the window model or the time constant of the exponential model is only a few past trials. If humans are optimal with respect to adjusting their gain based on variance in a limited number of past observations, we expect to find a short time window (or time constant) resulting in rapid adaptation. Using a longer time window, although not optimal (see Fig. 3), may provide robustness by having a less fluctuating gain. This would also elicit a decreased responsiveness towards systematic changes in the environment.

We also tested two models that make predictions in dynamic environments but do not utilize the Kalman filter framework. The change-point model (Adams and MacKay 2007) is a recursive Bayesian algorithm that estimates the probability that a change in the mean of the environmental process has occurred and even infers the changed statistics. The process that generated our stimuli is a change-point process, i.e., it consists of abrupt changes (change points) in the statistics of a sequentially presented variable (Barry and Hartigan 1993), and since the change-point model is specifically designed to predict and follow change-point processes, we expect it to perform better than the other models described here.

The accumulator model (Berniker et al. 2010) is also a recursive Bayesian algorithm that estimates the mean of the state as well as its variance. In this study, we only have the resources to experimentally estimate the mean of the prior, so we shall discuss the accumulator

model in this light alone. The accumulator algorithm estimates the current state as the mean of all past observations. This model therefore represents the theoretical optimum for an environment that is noisy and stationary but will not be able to follow systematic changes of the mean in a nonstationary environment.

General framework for Kalman filter algorithms. The Kalman filter is an online state estimation algorithm that generates a prediction of the current state based on the most recent observation, knowledge of the process dynamics, and its own previous state. The proportion in which the observation contributes to the estimate is the Kalman gain (K) of the system and depends on the variance of the prediction $[(\sigma_k^-)^2]$ and the variance in the measurement or observation (σ_m^2) of the environment. The variance of the prediction in turn depends on the variance of the process (σ_p^2) and the variance of the previous state estimate (σ_{k-1}^2). For all the adaptive models, the gain of the system (K) changes over time. The standard Kalman framework customized for the assumptions underlying our study is described by the following set of equations:

Prediction of state of the world.

$$\hat{x}_k^- = \hat{x}_{k-1}$$

$$(\sigma_k^-)^2 = \sigma_{k-1}^2 + \sigma_p^2$$

Update of state estimate given a measurement of the real world/environmental process.

$$K_k = \frac{(\sigma_k^-)^2}{(\sigma_k^-)^2 + \sigma_m^2}$$

$$\hat{x}_k = (1 - K_k)\hat{x}_k^- + K_k z_k$$

$$\sigma_k^2 = (1 - K_k)(\sigma_k^-)^2$$

Here the state estimate ($\hat{x}_k$), which is the estimate of the state (x_k) given noisy information at time k , is updated over time. Before the next observation (z_k) is made at the next discrete time-step, a prediction ($\hat{x}_k^-$) is generated based on knowledge of the dynamics of the environment. The prediction ($\hat{x}_k^-$) in these models is the variable we seek to estimate experimentally from human responses on the invisible-target trials of the experimental study. The variance in the prediction $[(\sigma_k^-)^2]$ and the measurement (σ_m^2) determine in what proportion (gain) the prediction and observation are combined.

Stationary model. We use a standard implementation of the Kalman filter (Kalman 1960) with stationary process and measurement variances, a prevalent assumption in perceptual and motor control literature (Baddeley et al. 2003; Izawa and Shadmehr 2011; Korenberg and Ghahramani 2002; Wolpert 1997; Wolpert and Ghahramani 2000, van Beers 2012). The variance of the measurement is the only free parameter in this model. The experimental process standard deviation is provided to the model ($\sigma_p = 50 \text{ ms } \forall k$). In our implementation, the only free parameter for the Stationary Model is the measurement variance (σ_m^2).

Window model. If there are large systematic changes in the environment, the variance of the process will increase just after a change occurred (Fig. 2). An efficient prediction system will use this increased variance as an indication that its prediction has been rendered unreliable and that under the circumstances, observations are more reliable. This is characteristic of an adaptive control system and is the kind of mechanism that allows the controller to momentarily increase its gain before reaching steady state again. Following this line of reasoning, we modify the Kalman filter algorithm such that the process variance (σ_p^2) is dynamically estimated from a window of past observations as the system experiences the environment over time. This approach lends the model the ability to adjust its gain based on the immediate demands of the environment (Fig. 2). The model has two free parameters: size of the window (T) and variance of the measurement (σ_m^2).

Fig. 2. Illustration of the window Kalman filter with the state estimation process being supplemented by the reestimation of process variance (green line) over a moving window (green shading) of the observed process (dark gray line). The 2 parameters of the model are the window size of T past trials (green shading) and the variance in measurement (purple) and the prior (yellow) determine the prediction (blue). The exponential Kalman filter operates with similar principles, except instead of calculating process variance in a window, it obtains an exponentially weighted estimate of the variance over all past trials with time constant τ .

Recalculation of real-world process statistics based on window of T past trials.

$$\bar{\mu}_k = \frac{1}{T} \sum_{i=k-T}^{k-1} z_i$$

$$\sigma_{P_k}^2 = \frac{1}{T-1} \sum_{i=k-T}^{k-1} (z_i - \bar{\mu}_k)^2$$

Exponent model. It is also plausible that the manner of estimating the past process variance could follow an exponential dynamic, thereby allowing more recent information to be weighted heavily while considering the more distant past to a lesser extent (Scheidt et al. 2001, Baddeley et al. 2003). We modify the Kalman filter algorithm to estimate the variability in the recently experienced process based on an exponentially decreasing weighting of past observations. The time constant (τ) of the exponent and the variance of the measurement noise are free parameters.

Recalculation of real-world process statistics based on exponential weighting with time constant τ over past trials.

$$w_i = \frac{e^{-(k-i)/\tau}}{\sum_{i=1}^{k-1} e^{-(k-i)/\tau}}$$

$$\bar{\mu}_k = \sum_{i=1}^{k-1} w_i z_i$$

$$\sigma_{P_k}^2 = \sum_{i=1}^{k-1} w_i (z_i - \bar{\mu}_k)^2$$

Change-point model. The experimental environment used in our study is nonstationary and can be formally described as a change point process (Barry and Hartigan 1993). In view of this consideration, the change-point model (Adams and MacKay 2007), which provides online inference and detection of changed statistics, is the most appropriate model for predicting the next stimulus. This online Bayes-

ian algorithm differs in its implementation from the Kalman filter. At every time step, based on previous observation and a parameter (hazard rate ζ) indicating the rate at which change-points occur (we assume this to be a constant), it infers the suitability of continuing with existing statistics of the environment or detecting a change point. Once a change has occurred, the number of trials until which the environment is considered stationary is called a "run." The run length is estimated online and at each change-point it is reset to zero. Here, the prediction of the state at a given point ($\hat{x}_k$) depends on the marginal probability of the state given the joint probability of all past observations within the run (r_{t-1}), and existing run length ($x'_{1:t-1}$) given all past observed stimuli observed thus far ($x'_{1:t-1}$) (see Adams and MacKay 2007 for details). We assumed that the generative distribution between change-points was a Gaussian with an unknown mean and a standard deviation of 50 ms. The only free parameter for the change-point model is the hazard rate ζ .

$$\hat{x}_k = \sum_{r_t} p(x_k | x'_{1:t-1}, r_{t-1}) p(r_{t-1} | x'_{1:t-1})$$

Accumulator model. The accumulator model is an online Bayesian estimation algorithm that generates predictions for dynamic environments. This too differs from the Kalman filter framework. The final state of the model converges towards the mean of the entire history of the observed process. The model computes the mean of the prior by estimating over all past observations (z_i) with equal weight given to the entire history (for details see Berniker et al. 2010). For our purposes therefore, there are no free parameters in the accumulator model.

Update of parameters (mean of state of the world) over discrete time.

$$\hat{x}_k = \frac{1}{k-1} \sum_{i=1}^{k-1} z_i$$

Optimal predictions for stimuli. To quantify the expected differences among the models, we first determined for which parameters these generative models would optimally predict the invisible stimuli of the sets

used in the experiment using least-squares optimization (Fig. 3). We then used corrected Akaike Information Criteria (AICc) to report the relative likelihood of each model, which is a comparison of how well each of the five models can predict the invisible-target stimuli given the visible-target stimuli (see Burnham and Anderson 2002 for details). We use the information-theoretic approach to compare our models, and therefore, concepts like significance levels, error bars, and P values do not apply. Information-theoretic approaches (like AICc), unlike hypothesis testing, allow us to compare multiple models simultaneously with respect to each other while accounting for the number of parameters of each (complexity of the model).

We found that, like the runs on the stimuli sets, the change-point model outperforms all other models (Fig. 3). This was expected since our stimuli follow a change-point process by design. The optimal hazard rate (ξ) of the change-point model was approximately five trials. Of the Kalman filter-based models, the window model performs the best, with a window size of about four trials. The optimal time constant for the exponent model is about four trials. Figure 3B shows an example of how the Kalman gain changes over time for the three models for one arbitrarily chosen stimulus set. For the Kalman filter-based models, the update for visible-target trials was performed in accordance with the filter update equations specified earlier. It is worth noting that in invisible-target trials, no state update of the mean of the state was performed. The absence of sensory information is equivalent to setting σ_m to infinity, or the gain to zero, on invisible-target trials (not plotted to preserve clarity in figures). The gain does not entirely recover in the immediately succeeding trial, which results in a jagged profile of the gain function (as seen in Fig. 3B, *inset*). All three Kalman-based models, when optimized to the stimuli, yield roughly the same values for measurement variance. Since we only estimate the mean of the environmental state estimate, there are no parameters to optimize for the accumulator model.

Optimal prediction of human behavior. The methods for fitting models to the participants' data were similar to those for the optimization with respect to stimuli, except now the parameter optimization was performed with respect to the participants' responses rather than

to the stimuli. Participants were provided with exactly the same stimulus sets as we tested the models on (Fig. 3). In the least-squares optimization, the root-mean squared error values for model comparisons were computed on differences between model prediction on invisible-target trials and participant responses on the same trials. All the Kalman-based models have measurement variance as a free parameter. Initial values for the process variance are the actual experimental process statistics that are specified in the same way for every model (σ_P). The adaptive models additionally had either a window size or a time constant as a free parameter, while the change-point model had one free parameter, the hazard rate (ξ).

The model equations are based on subject's actual observations (z_k) on a single trial. The exact values of the observations are unknown: we know that they resemble the stimuli but are corrupted with measurement noise. We performed Monte Carlo simulations in which random numbers were drawn to simulate the effects of measurement noise on all models. These simulations showed that since the measurements are independently drawn, when the predictions are averaged over a large number of such simulations, the prediction for each model was equivalent to the prediction when the model was simulated with zero measurement variance. We therefore ran our final simulations using the actual stimuli (x_k) rather than the observations (z_k) in the above equations. We did, of course, incorporate the variance of the measurement noise in the update equations.

RESULTS

The invisible-target trials provide us with a noisy online estimate of the participant's current prediction (prior), which is the state variable we sought to experimentally estimate. In the visible-target trials, the prediction is combined with a sensory likelihood, which yields an estimate of the posterior. We average these estimates across subjects to give the reader some indication of how well human participants were able to follow changes in the given environment (Fig. 4). We find that both for responses on visible and invisible-target trials, participants conform to the

Fig. 3. *A*: relative performance of each model for the given stimulus sets. The optimal parameters for each model are given. Error bars represent the SE in relative probability across stimulus sets. *B*: example from a single simulation of the time course of gain changes for the 3 Kalman filters (with optimal parameters). *C*: example of all 5 models run on a single stimuli set.

Fig. 4. Average subject responses over the course of the experiment for visible target trials (light blue) and invisible-target trials (orange) and the average of all stimulus sets (dark gray). Error bars (orange) represent standard error of the invisible-target trials across all subjects.

changes in the environment. In considering the responses to invisible-target trials, we find that participants are generally slower to adapt to the steps in the environment compared with the responses on visible-target trials. This was to be expected, since online-corrections are likely to occur in the presence of continuously observed visual stimuli and therefore responses are likely to be more accurate for visible-target trials.

In Fig. 5, we illustrate the simulations of all models for an individual dataset (*participant 6*). Given the knowledge that the stimuli environment is a change-point process, we know from earlier calculations (Fig. 3A) that the ideal model for our environment is the change-point model. Participants, unlike the model have no knowledge that the global structure of the environment is a change-point process. Participants therefore could not have known in advance that the optimal approach to follow our stimuli would coincide with the change-point model. We designed our stimuli such that they have very poor information about the change-points in the environment. It is therefore not surprising that the participant does not respond to step changes as fast as the best-fitting change-point model predicts since they had no foreknowledge that our stimuli followed a change-point process. Nevertheless, we wanted to use the change-point model to provide us with an upper performance bound for ideal behavior should the participant acquire information about change-points. The results show that this was not the case.

The accumulator model (light gray), as its dynamics dictate, weighs all past information with equal consideration and produces a response that converges to the mean of all involved distributions. The responses of the three Kalman filter-based models lie intermediate to those of the change-point and the accumulator models. It is difficult to visually distinguish the efficacy of the three Kalman models given the difference in number of parameters and

similarity of output, which is why we compared the models using information theoretic approaches.

In the Introduction, we reasoned that rapid adaptation to variations in an environment with nonstationary statistics requires switching to a higher gain following an abrupt change in the environment. Such behavior would be consistent with that of the adaptive Kalman models with a small window size or a short time constant. As we state earlier, the optimal window size for our stimuli is small (~ 4 trials) and the optimal time-constant short (~ 4 trials). We therefore first ran optimization routines on the two adaptive Kalman filter models (window and exponent) for window sizes and time constants of < 50 trials. Simultaneous model comparisons among all five models were performed using the corrected AICc (Burnham and Anderson 2002). The relative likelihood we display in Fig. 6C gives an indication of how well each model performs in explaining the participant responses with the optimal parameters that model has to offer. The relative likelihood is an information theoretic measure that gives the probability of correctness of models relative to each other and should not be confused with hypothesis testing.

The change-point model provides a poor description of the data even though the hazard rate (3–4 trials) closely matches the optimal values with respect to the stimuli (5 trials). The three models based on the Kalman filter perform evidently better than either the accumulator or change-point models for each participant. Figure 6A displays how the Kalman gain of the Kalman-based models changes over the course of the experiment. The gains for the adaptive models initially plunge to zero because they correctly estimate the process variance in the first 50 trials to be almost zero (variance: 1 ms^2). The Kalman gain becomes larger after the process variance in-

Fig. 5. A: simulation of the 5 models the parameters of which were optimized for a single participant. Participant responses for visible (blue circles) and invisible (orange circles) target trials are also plotted in the background.

Fig. 6. *A*: Kalman gains for all 3 Kalman models for the best short window sizes and time constants (≤ 50 trials), stationary model (red), window model (green), and exponent model (blue). Parameters for window size (T) and for time constant (τ) are indicated. Participant numbers are indicated at *top right* each subpanel. *B*: Kalman gains for models with parameters at global minima. *C*: relative probabilities (from Akaike Information Criteria) for all 5 models for short window sizes and time constants. *D*: relative probabilities for global minima for window size and time constant.

creases and fluctuates based on the variance within the window or exponent span of past trials. In the Introduction, we mentioned how such an increase in gain after the onset of a step change may help the system adapt faster by decreasing reliance on past knowledge and giving more weight to new observations, and this is what we observe. The caveat here is the size of the window and of the exponential time constant that for many participants were substantially longer (Fig. 6A) than the optimal value we expected after simulations (Fig. 3). These results provide no support for the idea that the behavior of the participants is best described by reliance on past information over a small window size or a short time constant. We also found it puzzling that many optimal values for the adaptive models lay at the limit of our search space, which indicated that the global minima values may be even larger.

We therefore also searched larger parameter domains (up to 400 trials). For five of the eight subjects, global minima were found at higher values of window size and time constants, which are much larger than the theoretically optimal value of about four trials (Fig. 6B). The window size and time constant values vary across participants, but the gain of the stationary model is quite consistent (Fig. 6B). With the parameters of the global minima,

the model comparisons support the adaptive models for most participants (Fig. 6D), but the gains for participants with large windows and time constants approximately stabilized to the stationary gain value (Fig. 6B). In those cases, the gain no longer increases following a step, as it does for the subjects with small window sizes or time constants, so the potential benefit in performance of the adaptive models (in terms of following the stimuli) was not realized in accordance to our expectation.

We plotted cumulative absolute errors between the Kalman-based models and participant data over trials to check whether the advantage for the adaptive models arose from the first 50 trials of the experiment where there is almost no noise in the stimuli. This was not the case. The adaptive models often started out with higher errors compared with the stationary model and converged to a lower overall error over the course of the experiment, suggesting a pervasive and widespread advantage for the adaptive models over the course of the experiment.

DISCUSSION

With this study, we sought to investigate whether the nervous system dynamically changes the extent to which it relies

on past information on encountering abrupt changes in the environment. Our study was partly motivated by the observation that for many everyday sensorimotor tasks, humans seem to adapt very rapidly to an entirely new distribution of an environmental variable. One possible manner to achieve such rapid adaptation in nonstationary environments would be to cease ones' reliance upon prior knowledge in favor of a greater reliance on recently observed information about the process. If the variance of the process were estimated over a small window of past information, a temporary increase of the gain of the system would take place after such a change occurred. Such an increase in gain would lead to faster adjustments and thereby better performance. Under the assumption that the process variance is recursively estimated by the observer, we designed an experimental environment in which the best possible performance for our stimuli is furnished by a gain based on a window of only a few past observations.

We designed an experiment to estimate the prior that the participant was using at a high temporal resolution during the course of the experiment. Berniker et al. (2010) were the first to infer the human sensorimotor prior from data and since then others have done the same (Turnham et al. 2011). We, however, did not infer the prior but obtained a direct measurement of its mean in the absence of any sensory evidence whatsoever. This estimate of the mean of the prior is directly comparable to predictions (state estimates) generated by online models. We therefore fit models to the online estimate of the before determine which of adaptive or stationary prediction mechanisms are at play in human performance.

We knew from simulations that the change-point model was the best candidate model to predict the stimuli and the window model was the next-best solution. We found, however, that for most participants the adaptive window Kalman model performed better than all other models. The fitted values of the window to participants' data were much larger than the optimal value obtained from simulations. This could be interpreted as participants being more conservative in adapting the gain of their responses to abrupt environmental changes. While a small window over past information would yield faster adaptation after a step change in the environment, a large window provides a more accurate estimation of the statistics of the distribution during intervals with no change. It may appear from this that participants preferred performing better during a block rather than adapting rapidly to a step change, resulting in longer window sizes and time constants than those that are optimal overall. On the other hand, they seem to compensate for the invariance of their Kalman gain by having an overall higher gain than would be optimal (compare gain plots in Figs. 3 and 6).

The accumulator model (Berniker et al. 2010) and the change-point model (Adam and MacKay 2007) represent opposite poles of the above described control tendencies. The accumulator model is well suited to accurately estimate environmental process statistics over all past information, leading to a steady regression in its response towards the mean of all past observations. The change-point model is well suited to deal with abrupt changes to new distributions. Neither the change-point nor the accumulator model performs well in describing human behavior in our experiment.

In conclusion, we sought to provide a mechanism that could explain how the brain may estimate environmental parameters from past observations. The Kalman filter has long been proposed

as a solution to online prediction challenges faced by the brain, however, its parameters are often determined offline and are stationary. We propose some adaptive modifications to the Kalman filter to address how the brain may be able to adjust its gain online by monitoring recent changes in the environmental variance. While our efforts show that such adaptive tendencies better explain data for most participants, the range of past trials over which humans seem to be monitoring environmental variance is much larger than optimally predicted by our model. We must conclude that while adaptive behavior is present for many participants, the adaptive Kalman models we propose do not capture the observed sensorimotor behavior in the manner that we expected.

ACKNOWLEDGMENTS

We thank Wei Ji Ma, Marc Ernst, and Loes van Dam for insightful discussions. We also thank Max Di Luca for lending us calibration equipment.

GRANTS

Funding for this research was provided by the European Community's Seventh Framework Programme FP7/2007-2013 (Grant Agreement No. 214728-2), and by the VICI Grant (453-08-004) from the Netherlands Organization for Scientific Research Grants.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

AUTHOR CONTRIBUTIONS

Author contributions: D.N., R.J.v.B., J.B.J.S., and E.B. conception and design of research; D.N. performed experiments; D.N. analyzed data; D.N., R.J.v.B., J.B.J.S., and E.B. interpreted results of experiments; D.N. prepared figures; D.N. drafted manuscript; D.N., R.J.v.B., J.B.J.S., and E.B. edited and revised manuscript; D.N., R.J.v.B., J.B.J.S., and E.B. approved final version of manuscript.

REFERENCES

- Adams RP, MacKay DJ. *Bayesian Online Changepoint Detection* (Technical Report). Cambridge, UK: Univ. Cambridge, 2007.
- Adams WJ, Graf EW, Ernst MO. Experience can change the "light-from-above" prior. *Nat Neurosci* 7: 1057–1058, 2004.
- Baddeley R, Ingram H, Miall R. System identification applied to a visuomotor task: near-optimal human performance in a noisy changing task. *J Neurosci* 23: 3066–3075, 2003.
- Barry D, Hartigan JA. A Bayesian analysis of change point problems. *J Am Stat Assoc* 88: 309–319, 1993.
- Berniker M, Voss M, Körding KP. Learning priors for Bayesian computations in the nervous system. *PLoS One* 5: e12686, 2010.
- Brenner E, Smeets J. Modifying one's hand's trajectory when a moving target's orientation changes. *Exp Brain Res* 196: 375–383, 2009.
- Burge J, Ernst MO, Banks M. The statistical determinants of adaptation rate in human reaching. *J Vis* 8: 20: 1–19, 2008.
- Burnham KP, Anderson DR. *Model Selection and Multimodel Inference: A Practical-Theoretic Approach*. New York: Springer, 2002.
- Di Luca M. New method to measure end-to-end delay of virtual reality. *Presence Teleoperators Virtual Environ* 6: 569–584, 2011.
- Ernst MO. Learning to integrate arbitrary signals from vision and touch. *J Vis* 7: 1–14, 2007.
- Izawa J, Shadmehr R. Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comput Biol* 7: e1002012, 2011.
- Kalman RE. A new approach to linear filtering and prediction problems. *J Basic Eng* 82: 35–45, 1960.
- Knill D. Learning Bayesian priors for depth perception. *J Vis* 7: 13, 2007.
- Knill DC, Richards W. *Perception as Bayesian Inference*. Cambridge, UK: Cambridge University Press, 1996.
- Körding KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature* 427: 244–247, 2004.
- Korenberg AT, Ghahramani Z. A Bayesian view of motor adaptation. *Curr Psychol Cogn* 21: 537–564, 2002.

- Scheidt RA, Dingwell JB, Mussa-Ivaldi FA.** Learning to move amid uncertainty. *J Neurophysiol* 86: 971–985, 2001.
- Turnham EJA, Braun DA, Wolpert DM.** Inferring visuomotor priors for sensorimotor learning. *PLoS Comput Biol* 7: e1001112, 2011.
- van Beers RJ.** How does our motor system determine its learning rate? *PLoS One* 7: e49373, 2012.
- Verstynen T, Sabes PN.** How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *J Neurosci* 31: 10050–10059, 2011.
- Wei K, Körding K.** Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Front Comput Neurosci* 4: 878–881, 2010.
- Weiss Y, Simoncelli EP, Adelson EH.** Motion illusions as optimal percepts. *Nat Neurosci* 5: 598–604, 2002.
- Wolpert DM.** Computational approaches to motor control. *Trends Cogn Sci* 1: 209–216, 1997.
- Wolpert DM, Ghahramani Z.** Computational principles of movement neuroscience. *Nat Neurosci* 3: 1212–1217, 2000.