

This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



ELSEVIER

Available online at www.sciencedirect.com
 ScienceDirect

Journal of Physiology - Paris 101 (2007) 64–77

Journal of
Physiology
Paris

www.elsevier.com/locate/jphysparis

Bayesian modeling of dynamic motion integration

Anna Montagnini ^{a,*}, Pascal Mamassian ^b, Laurent Perrinet ^a, Eric Castet ^a,
 Guillaume S. Masson ^a

^a *Institut de Neurosciences Cognitives de la Méditerranée, UMR 6193 CNRS – Université de la Méditerranée, 31 chemin Joseph Aiguier, 13402 Marseille, France*

^b *Laboratoire Psychologie de la Perception, FRE 2929, CNRS – Université Paris Descartes, 45 rue des Saints-Pères, 75006 Paris, France*

Abstract

The quality of the representation of an object's motion is limited by the noise in the sensory input as well as by an intrinsic ambiguity due to the spatial limitation of the visual motion analyzers (*aperture problem*). Perceptual and oculomotor data demonstrate that motion processing of extended objects is initially dominated by the local 1D motion cues, related to the object's edges and orthogonal to them, whereas 2D information, related to terminators (or edge-endings), takes progressively over and leads to the final correct representation of global motion. A Bayesian framework accounting for the sensory noise and general expectancies for object velocities has proven successful in explaining several experimental findings concerning early motion processing [Weiss, Y., Adelson, E., 1998. Slow and smooth: a Bayesian theory for the combination of local motion signals in human vision. MIT Technical report, A.I. Memo 1624]. In particular, these models provide a qualitative account for the initial bias induced by the 1D motion cue. However, a complete functional model, encompassing the dynamical evolution of object motion perception, including the integration of different motion cues, is still lacking. Here we outline several experimental observations concerning human smooth pursuit of moving objects and more particularly the time course of its initiation phase, which reflects the ongoing motion integration process. In addition, we propose a recursive extension of the Bayesian model, motivated and constrained by our oculomotor data, to describe the dynamical integration of 1D and 2D motion information. We compare the model predictions for object motion tracking with human oculomotor recordings.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Object motion; Aperture problem; Smooth pursuit eye movement; Bayesian model; Recursive inference; Temporal evolution

1. Introduction

Efficient object motion processing is achieved in humans and non-human primates by integrating multiple noisy local motion signals. It is appropriate to distinguish two types of local motion signals, ambiguous and non-ambiguous ones. Motion signals from elongated uni-dimensional (1D) contours are ambiguous when analyzed through a spatially limited aperture (see Fig. 1), similar to the receptive field of many neurons in the motion-sensitive middle-temporal (MT) cortical area (Albright, 1984). The ambiguity relies

on the fact that the motion of the contour in the tangential direction is unknown, so that the observed movement is consistent with a family of possible motion directions and velocities (Fig. 1). In contrast, motion signals from local 2D features (e.g. terminators) are non-ambiguous, and psychophysical (Lorenceau and Shiffrar, 1992) and physiological (Pack and Born, 2001) studies have demonstrated that these signals can be used to reliably solve the aperture problem. However, the integration of 1D and 2D information is time-demanding and very short presentations of moving objects may give rise to characteristic perceptual errors (Lorenceau et al., 1993) that are biased in the direction orthogonal to the contour. For example, Lorenceau and Shiffrar (1992) found that tilted lines (+20° anti-clockwise with respect to the vertical) moving to the right and down

* Corresponding author.

E-mail address: Anna.Montagnini@incm.cnrs-mrs.fr (A. Montagnini).




Fig. 1. The aperture problem. A simple moving stimulus such as a tilted line carries ambiguous information about motion when it is observed through a small aperture, like in the right panel. In this example, a tilted line which is moving horizontally tends to be perceived as moving orthogonally to its edge.

are perceived as moving upward if presented very briefly at a low contrast. For longer presentations, the perceptual bias tends to be reduced and eventually eliminated. In parallel, electrophysiological recordings have shown that direction selectivity for motion-sensitive neurons in MT changes across time, over a typical time interval of ~ 60 ms (Pack and Born, 2001). MT neurons respond mostly to the direction orthogonal to object's motion, whereas later they encode the actual object's motion.

Recent studies on smooth pursuit eye movements (SPEM) do also provide an account of early motion processing which parallels the former findings in psychophysical experiments (see also Section 2). When human subjects or monkeys are required to visually track a moving object carrying different 1D and 2D information, eye-velocity traces are transiently biased, at pursuit initiation, toward the 1D-cued direction, i.e. orthogonally to the object's contour (Masson and Stone, 2002; Wallace et al., 2005; Born et al., 2006). Later, the edge-orthogonal 1D-bias (or *tracking error*) is progressively eliminated and eye-velocity converges to the object's global motion. Typically, the resolution of motion signal ambiguity is achieved within the first 300–400 ms after the presentation of the moving stimulus.

Beside the intrinsic ambiguity resulting from the local edge direction of motion, visual motion is also affected by the noise embedded in the sensory input *per se*. These two sources of uncertainty can be well integrated within a Bayesian framework (Weiss et al., 2002; Weiss and Fleet, 2002; Stocker and Simoncelli, 2006; Perrinet et al., 2005) where the perceived motion is the solution of a statistical inference problem. In these models, the information from local 1D and 2D motions can be represented by their likelihood functions and these functions can be derived for simple objects with the help of a few reasonable assumptions (Weiss and Fleet, 2002). Bayesian models also allow the inclusion of prior constraints and the most common assumption used in motion models is a preference for slow speeds. The effects of priors are especially salient when signal uncertainty is high. One way to increase the uncertainty of a visual stimulus is to reduce its contrast, and in these cases, perceived velocity is indeed underestimated (Thomp-

son, 1982), thereby providing some experimental support for the slowness prior. Interestingly, Priebe and Lisberger, 2004 have demonstrated that increasing the spatial frequency of the moving stimuli leads to qualitatively similar results than a decrease of contrast (or, more generally, an increase of visual noise), namely to the underestimation of perceived motion speed.

Up to now, Bayesian motion models have been applied to qualitatively predict, for instance, the initial bias toward 1D motion signals observed experimentally and its dependence on sensory noise (Weiss et al., 2002). We propose here to develop this theoretical framework in order to model smooth pursuit eye movements when tracking moving objects that carry multiple local cues. In particular we will focus on the dynamical evolution of the tracking error which reflects, in our opinion, the main characteristics of the underlying dynamical motion integration process. Our dynamic model is composed of a Bayesian kernel and an updating rule. The Bayesian kernel is fairly traditional, combining prior knowledge on speed with the current estimate to produce a robust inference of velocity. The updating rule revises the prior with time, thereby reflecting all past evidence about particular velocities. We propose that prior knowledge represents initially a default assumption independently of any stimulus that is then recursively updated by using the previous posterior probability as the current prior. The recursive injection of posterior distribution boosts the spread of information about the object's global shape, favoring the disambiguation of 1D by 2D cues. We also propose to both constrain and validate this model by means of experimental recordings of smooth pursuit eye movements.

2. Dynamic motion integration: an oculomotor account

Humans and monkeys are perfectly able to visually track the center of a moving extended object. The general purpose of these smooth voluntary eye movements is the stabilization of the image of the moving object on the fovea. Tracking accuracy during the steady-state movement is very high regardless of the orientation of the object's edges with respect to motion direction.

However, before the steady-state movement is achieved, significant biases can be observed. When the orientation of a moving line is not orthogonal to motion direction, the oculomotor trajectory initially deviates from the global motion direction, moving in a direction biased toward the perpendicular to the edge. For an horizontally moving tilted line (see Fig. 2), for instance, a non-zero transient vertical component is consistently observed, both in human and non-human primate observers (Wallace et al., 2005; Born et al., 2006). This *tracking error* typically peaks to a maximum at around 100 ms after pursuit onset and then it decreases to zero within the next 100–200 ms. The peak of the tracking error occurs before the beginning of the closed-loop phase, namely when feedback information about the eye motion is not yet available. This suggests that

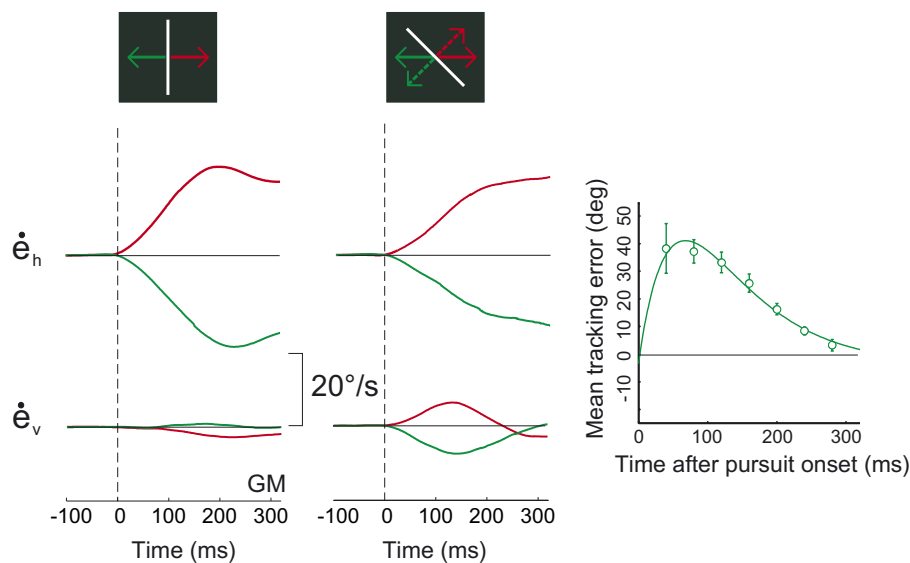


Fig. 2. The aperture problem and smooth pursuit eye movements. Top: cartoon of stimuli composed of non-conflicting (vertical line) and conflicting (tilted line) information from 1D and 2D cues. Solid (dashed) arrows represent actual (initially perceived) motion. Bottom Left and Middle: Horizontal and vertical eye velocity when tracking the non-conflicting and the conflicting stimulus. Bottom Right: Mean tracking error with respect to object trajectory. This figure is modified, with permission of authors, from (Wallace et al., 2005).

the dynamic correction of the tracking error is primarily a visual phenomenon.

The horizontal and vertical eye-velocity profiles, as well as the tracking error during pursuit initiation of vertical and tilted lines moving rightwards at 20 deg/s are shown in Fig. 2 for a representative human subject. For the vertical line, edge-related (1D) and terminator-related (2D) cues are in agreement, both indicating horizontal motion. In contrast, for the tilted line 1D and 2D cues carry conflicting information, leading to the transient tracking error.

In the example in Fig. 2 the peak tracking error reaches almost 40°, which is actually quite rare when using high-contrast stimuli. Most often, SPEM tracking error starts with a value bigger than 0 and below the 45° predicted by the 1D-cue and then it decreases to 0. For the Ocular Following Reflex at ultra-short latency, Masson and colleagues (Masson et al., 2000) have shown that the initial response, initiated about 85 ms after stimulus onset is directed in the exact 1D-cued direction and only ~20 ms later the 2D information starts to become relevant leading to the correction of the tracking error. One explanation of this discrepancy between different types of eye movements is that, because of their longer latency, the smooth pursuit system has the time to integrate some information before movement initiation, thereby confounding very small offset differences in the processing of different motion components.

Overall, the time course of the ocular tracking error is in good agreement with the psychophysical results on the motion perception bias (Lorenceanu and Shiffrar, 1992; Lorenceanu et al., 1993), again supporting the idea that the transient oculomotor bias originates from local visual motion processing. When compared to the temporal evolution of motion direction tuning curves observed in MT cells (Pack and Born, 2001), the dynamics of the tracking error

correction appears to be slower. This fact has to be interpreted as the evidence that the oculomotor response corresponds to a low-pass filtered account of the visual dynamics, because of the properties of the oculomotor plant (Robinson et al., 1986).

Two recent studies (Osborne et al., 2005, 2007) have shown that most of SPEM velocity variability, at initiation, can be explained in terms of uncertainty about the sensory input. Although the exact match between perceptual and oculomotor variability is still a debated issue (Gegenfurtner et al., 2003; Priebe and Lisberger, 2004), we will assume that, at least at smooth pursuit initiation, oculomotor variability does mainly reflect sensory uncertainty (see Section 3.1). A purely oculomotor component of pursuit variability is likely to affect the smooth pursuit data: Osborne et al. (2005, 2007) have measured what they call the *background* noise and have shown that this component is clearly distinguishable from the sensory-driven variability and that the two components are simply linearly summed during visual pursuit. Similarly to their evaluation of background noise, in the present study, we evaluated the standard error of the mean horizontal eye velocity in a 40-ms time window centered at the time of presentation of the moving stimulus. We did it for all four stimuli we used, a circular blob, a very long vertical line, a 17°-tilted line and a 12° vertical line. In all cases was the estimated background noise about an order of magnitude smaller than the velocity SEM computed during the steady-state pursuit phase: overall mean SEM velocity was respectively 0.09 and 0.94 deg/s, the difference being statistically highly significant (paired *t*-test, $p < 10^{-4}$). For this reason we decided to disregard here the background motor component of velocity variability.

Finally, we will assume, at this first stage, that oculomotor variability affects independently the horizontal and

vertical components of smooth pursuit eye movements (in line with the assumption that sensory variability is considered independent at different points in the velocity space, see Section 3.2.1).

2.1. Effects of the physical properties of the moving stimulus

The tracking error has been shown to parametrically depend on several characteristics of the moving stimulus. When luminance contrast is reduced, the tracking error increases in size, reaching, at its maximum, values very close to the edge–orthogonal direction. In addition, the time duration of the transient bias is bigger (Wallace et al., 2005). The tracking error tends also to increase weakly with stimulus speed, but in this case its time course remains roughly unchanged. When the moving stimulus is presented in the visual periphery, the smooth pursuit of the eccentric stimulus presents only a very reduced tracking error. Stimulus size (line length) boosts the transient bias (Born et al., 2006), whereas the introduction of more terminators (e.g. braking the contours in smaller dashes) reduces both the size and the duration of the orthogonal transient bias (Wallace et al., 2005). These findings replicate in many aspects older analogous studies in visual psychophysics (Castet et al., 1993) that had investigated the dependence of the perceived illusory motion induced by the aperture problem upon several physical properties of the stimuli. This set of experimental observations provides a composite and straightforward benchmark for our model of the dynamics of motion processing based on sensory uncertainty.

2.2. Effect of motion predictability: no transfer across trials

A striking aspect of the initial bias for pursuit is that it is highly reproducible and seems to be immune to cognitive influence such as shape cueing (Wallace et al., 2005). We have recently tested the robustness of the tracking error to motion predictability by manipulating, across experimental blocks, the probability of occurrence of a given stimulus orientation and/or motion direction (Montagnini et al., 2006). Interestingly, no reduction of the average tracking error at pursuit initiation was observed, across many trials, when motion predictability was very high (90% probability) or complete (100% probability). In apparent contradiction with this finding, a robust anticipatory pursuit in the true global motion direction was observed, whose size was proportional to the strength of predictability. This fact proves that an efficient predictive signal about correct object motion was indeed available to the oculomotor system, although it was clearly not used at the initial stage of motion cue integration to reduce the tracking error. Taken together, these findings suggest that predictive information can differently affect distinct phases of smooth pursuit eye movements. Moreover, and more specifically to the present purpose of a Bayesian model of motion perception, these findings constrain the limits and

the nature of the prior that is the crucial dynamic variable of our model. In the current study, we use the term prior to refer to the accumulated evidence that an observer assigns to each point in velocity space. This probability distribution is defined at any time since motion onset, but it does not transfer from trial to trial. In other words, for any trial, at time 0 our prior P_0 corresponds to the traditional definition of a prior, i.e. a stimulus-independent statistical belief about motion. Based on our past work (Montagnini et al., 2006), we assume that previous trial experience does not influence the initial prior on velocity P_0 although it might influence motion expectations at a different level to explain the anticipatory eye movements. As explained in detail in the next section (3.2), the prior in our model is not taken to be fixed, and its recursive update is fundamental to correct the tracking error within the first 300 ms of motion.

3. General ideas and methods

3.1. Oculomotor recordings: constraining and testing the model

3.1.1. Smooth pursuit recording and analysis

In the first set of oculomotor experiments, we recorded smooth pursuit eye movements from three human subjects (two authors of the paper and one naïve subject) while they were tracking one of two objects. The first object was a circular Gaussian spot and the second a line whose length could be approximated as infinite, in the sense that terminators were very far in the periphery and therefore their influence was presumably very limited. These stimuli moved with various motion directions and speeds (see Fig. 3). Vertical and horizontal position of the right eye was recorded by means of the scleral eye coil technique (Collewijn et al., 1975). Data were sampled at 1 kHz, low-pass filtered (DC – 130 Hz) and digitized with 16 bit resolution. Eye-position data were linearized off-line and smoothed using a spline algorithm before differentiation to obtain eye-velocity profiles. In this first experiment, stimuli were either a central blob with Gaussian luminosity profile (standard deviation $\sim 0.2^\circ$ of visual angle) or a long line (length: 48° same Gaussian profile for its width). Peak luminance was always 60 cd/m^2 . In all cases the moving targets were highly visible against the black background. We used three different target speed values: 5, 10 and 15 deg/s ; target motion was always horizontal and the direction (right or left) was randomly chosen at each trial. The second set of oculomotor experiments provided a database to quantitatively test the model's prediction concerning the tracking error's dynamics. Stimuli were always a single line, either tilted 45° relative to horizontal, (length: 17° of visual angle) or vertical (length 12°), the latter providing a reliable control stimulus for possible idiosyncratic movement bias. We collected between 100 and 150 trials for each condition (4 stimulus types \times 2 directions \times 3 speed values) from each subject. Off-line inspection of the individual eye-velocity curves allowed us to eliminate




Fig. 3. Left: Moving stimuli carrying mostly 2D (central spot, top) or 1D (very long line, bottom) information. Right: Example of horizontal velocity traces at pursuit initiation recorded when one subject was tracking the long vertical line. The trial-by-trial variability of horizontal velocity around the time of peak acceleration is used to estimate the variance of the 2D and 1D likelihood functions.

aberrant trials (overall less than 5%), like those including eye blinks. In addition, an automatic procedure was used to detect *catch-up* saccades within the pursuit traces and to cut out the corresponding part of the velocity traces, before the evaluation of the sample mean and variance and any other analysis.

All statistical analyses were repeated in two conditions, i.e. with the traces aligned with respect to stimulus onset and with the traces aligned with respect to pursuit onset. The reason to do this was to control and eventually eliminate the component of variability in the eye velocity which is due to the jitter in pursuit latency and possibly to uncertainty in time estimation (as pointed out by Osborne et al. (2005)). As a matter of fact, pursuit latency was quite stable across conditions in our experiments (averaging between 100 and 120 ms) and we could not detect any major difference in our results depending on the time-locking choice. Therefore, since our model does not make any specific hypothesis about latency variability, in this paper we chose to present only results obtained with the eye movement traces aligned on pursuit onset.

3.1.2. Using pursuit data to estimate Bayesian uncertainty

Previous works (Hürimann et al., 1992; Stocker and Simoncelli, 2006) have attempted to validate and constrain the theoretical Bayesian framework for motion perception (Weiss et al., 2002) by using psychophysical experimental data. Here, we propose to use smooth pursuit eye movements data for both phases, i.e. in order to estimate the crucial model parameter (thereby constraining the model at the quantitative level) and, later, to test its predictions (see Section 4).

We assume that the variability of pursuit velocity at initiation mostly reflects the variability of the estimated target

velocity. This assumption allows us to infer the variance of the posterior velocity distribution in our model from the measurement of pursuit data variability. In turn, the inferred variance for the posterior can be used to infer the variance of the Bayesian prior and likelihood functions. In particular, using the specific stimuli of the first experiment (small circular Gaussian spot and very long line), we can infer the spread of the likelihoods for the 2D and 1D visual sources of information. The mathematical details of this derivation are explicated in Section 3.2.1.

For each pursuit trial, we used the mean eye velocity across a 40 ms time window centered around the time of peak acceleration (i.e. between 20 and 60 ms after pursuit onset or 120–160 ms after stimulus onset) as the relevant estimate for the initial posterior distribution (see Fig. 3, right panel). This choice was motivated by the fact that pursuit dynamics at the peak acceleration can be considered as approximating the impulse-response function to a step in the target velocity (the peak eye acceleration being proportional to target velocity for smooth pursuit eye movements (Lisberger et al., 1987)). We computed the average and variance across the trial population of this impulse-response velocity. The derived quantities were used as estimates of the means and variances of the posterior distributions (see Section 3.2.1).

3.2. Model: An adaptive Bayesian framework to explain dynamic motion integration

Our modeling effort is an attempt to reproduce the characteristic tracking errors witnessed during smooth pursuit of a moving tilted line. Mathematical details are reported below (Section 3.2.1). We assume here that the problem is solved at the level of velocity space, thereby




Fig. 4. Statistical inference in the velocity space. Top: the prior and the two independent 1D- and 2D-likelihood functions (for a tilted line moving rightwards at 5 deg/s) are multiplied to obtain the posterior velocity distribution (bottom). The inferred image motion is estimated as the velocity corresponding to the posterior maximum (MAP). Probability density functions are colour-coded, such that dark red corresponds to the highest probability and dark blue to the lowest one.

considering solved the critical stage of extracting velocities from the image, for which we follow the approach proposed by Weiss and colleagues (Weiss and Adelson, 1998). We also follow some basic assumptions of previous Bayesian static models of motion processing (see Fig. 4), namely that

- at time $t = 0$ (stimulus motion onset) the prior is Gaussian, centered on $(v_x, v_y) = (0, 0)$, with variance σ_0^2 ;
- the likelihood functions computed at any given point of the moving object are independent. As a consequence, 1D (edge related) and 2D (terminator related) likelihood functions can be multiplied to compute the object likelihood;
- the sensory noise affecting the likelihood functions is Gaussian, with variance σ_1^2 and σ_2^2 for the edge and terminator likelihood respectively.

The posterior distribution of perceived motion is related to prior and likelihood functions by means of Bayes' rule (as represented in Fig. 4). Image motion is inferred by means of a *Maximum a Posteriori* estimate (MAP).

Our aim here being to model the time course of the pursuit tracking error, we implemented a recursive method to update the prior using past information.¹ This is somewhat equivalent to a dynamic filtering procedure (Kalman, 1958). Indeed the role of the prior distribution in our recurrent probabilistic model is very similar to the idea that the prior information in the Kalman filter includes all the knowledge accumulated across time, before the integration of the current piece of information.

An interesting model of a cortical recurrent neural network implementing dynamic Bayesian inference has been proposed by Rao (2004). It is important to underline here that although our model is critically grounded on the

¹ A (rapidly) evolving prior is somewhat contradictory with the established Bayesian definition of prior (but see Mamassian and Landy, 2001). Beyond the possibly misleading name, our aim here is only to suggest a formal way to update the statistical inference which is at the basis of motion perception. The updating rule could probably be different – for instance affecting the likelihood instead of the prior – but the crucial point is that it has to progressively account for an increasing weight of 2D cues.

recurrent update of information, it does not provide any indication about the nature of the underlying cortical architecture. For instance, one possibility is that feedback connections from MT to V1 could play a crucial role in the accumulation of evidence leading to the correction of the tracking error, as suggested by Bayerl and Neumann (2004). Alternatively, the update of information could even occur at the early stages of integration of the visual input, as suggested by Weiss and Adelson (1998). In practice, in our model at each discrete time step (in arbitrary units) the previous posterior function is used as the current prior function (see arrow scheme on the left of Fig. 4). Fig. 5(top) shows a cartoon of the simulated evolution of the posterior distribution for a moving tilted line from motion onset to the final matching with the object velocity. The inferred motion direction is given at any time by the velocity vector corresponding to the *maximum a posteriori* (MAP) of the probability distribution. Fig. 5(bottom) represents the simulated evolution of such inferred direction expressed as tracking-error. The initial tracking is clearly biased in the direction orthogonal to the line (1D motion cue), and it is progressively corrected to match the true object motion (2D motion), as it was experimentally observed in human smooth pursuit (see Fig. 2).

Note that during smooth pursuit the retinal velocity converges to zero: the stabilization of the moving visual image on the retina is indeed the purpose of this kind of eye movements. Classical models of oculomotor control (Robinson et al., 1986) postulate the existence of an *efferent copy* of the oculomotor command, in the closed-loop phase, which is fed back to the input-stage of the smooth

pursuit control system. The information carried by the efferent copy is summed to the retinal visual input, thereby preventing the eyes from interrupting their tracking motion once the image is stabilized and the retinal velocity is around zero. In our model we assume that in the closed-loop phase the efferent copy is combined with the retinal input with no delay and no additional oculomotor noise, which is equivalent to consider visual information in an absolute allocentric reference frame instead of the retino-centric frame.

A more general prediction of our recurrent model is that posterior variance, which is assumed to be faithfully represented by the experimental smooth pursuit velocity, will decrease across time. This fact can be easily observed in the upper panels of Fig. 5 and will be made more explicit mathematically later (see Eqs. (9) and (10)). This prediction holds regardless of the moving stimulus considered. Therefore, in order to test this prediction we pooled data from all sets of recordings (i.e. using the four types of visual stimuli described in Section 3.1.1), for all subjects and conditions, and we compared the standard error of the mean for eye-velocity computed at the beginning of the pursuit steady state (~ 240 ms after stimulus onset) and at a later phase of the steady state (~ 340 ms after stimulus onset). Mean velocity SEM was respectively 0.94 and 0.35 deg/s, this difference being highly significant (as confirmed by a one-tailed paired *t*-test, $t(69) = 3.27$, $p < 10^{-3}$), and suggesting that indeed velocity variability is consistently reduced across time. Unfortunately our experimental trials lasted only 500 ms after stimulus onset and therefore they did not allow us to analyse the dynamic decrease of velocity variance on a longer interval.

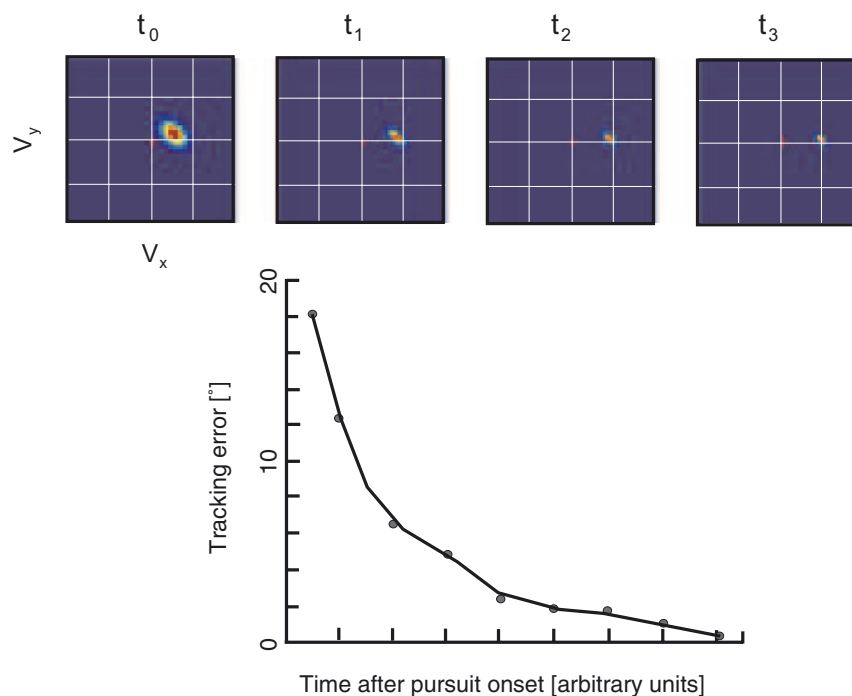


Fig. 5. Top: Evolution of the posterior probability: four arbitrary time steps are shown from left to right. Bottom: Simulated tracking error.

The model can now be tested against quantitative predictions about the time course of SPEM tracking error for different conditions of the sensory input. For instance, all the experimental effects described in Section 2.1 can virtually be reproduced by the model. In principle the motion integration dynamics for more complex objects can also be predicted, as a combination of several independent 1D and 2D components.

3.2.1. Mathematical details of the model

Let us assume that the likelihood function is composed of two parts, one corresponding to the information coming from the middle of the line, the other to the information coming from the ends of the line (terminators). We assume that these two pieces of information are processed independently and that their uncertainty function is well described by a normal distribution. With no loss of generality, we will assume that the line is moving horizontally at speed v_0 . The likelihood L_1 for the middle of the line is subject to the aperture problem, i.e. the velocity component tangent to the orientation of the line is undetermined. For each coordinate pair (v_x, v_y) in velocity space, we have:

$$L_1(v_x, v_y) = \frac{1}{Z} \exp \left(-\frac{((v_x - v_0) \cos(\theta) + v_y \sin(\theta))^2}{2\sigma_1^2} \right), \quad (1)$$

where Z is the partition function (or normalizing constant: we will use the same notation Z for different distributions along the paper), θ is the orientation of the line (0 is vertical), and σ_1 is the standard deviation of the uncertainty on speed of the line motion along the normal to the line. The likelihood L_2 for the terminators is not subject to the aperture problem and is therefore a noisy unbiased estimate of the true velocity:

$$L_2(v_x, v_y) = \frac{1}{Z} \exp \left(-\frac{(v_x - v_0)^2 + v_y^2}{2\sigma_2^2} \right), \quad (2)$$

where σ_2 is the standard deviation of the uncertainty on speed of the line motion at the terminators. Because of our independence assumption, the overall likelihood for the line motion is simply:

$$L(v_x, v_y) = L_1(v_x, v_y)L_2(v_x, v_y). \quad (3)$$

At the onset of the stimulus, we assume not only that we do not have any bias to perceive any particular direction of motion but also that motion speed is slow. In other words, the initial prior P_0 of velocity is centered on the origin of the velocity space (i.e. the most expected speed value is 0) and it is normally distributed, with variance σ_0^2 about the origin:

$$P_0(v_x, v_y) = \frac{1}{Z} \exp \left(-\frac{v_x^2 + v_y^2}{2\sigma_0^2} \right). \quad (4)$$

Likelihood and prior information can be combined thanks to Bayes' rule to produce the posterior distribution Q_0 :

$$Q_0(v_x, v_y) = L(v_x, v_y)P_0(v_x, v_y). \quad (5)$$

If we adopt as decision rule the maximum-a-posteriori (MAP), the inferred velocity at the onset of the stimulus is

$$(\hat{v}_x, \hat{v}_y) = \arg \max_{v_x, v_y} Q_0(v_x, v_y). \quad (6)$$

It should be noted however that in the case of Gaussian distributions the argument of the maximum-a-posteriori is equivalent to the mean of the distribution.

The posterior probability distribution function includes all the information extracted from the stimulus and from prior knowledge. Therefore, it makes sense to pass on this knowledge at the next iterative stage when new information is extracted from the stimulus. By then, the original prior is out-of-date and the best strategy is to replace it with the past posterior. Iteratively from time $(t - 1)$ to time t , we thus obtain:

$$P_t(v_x, v_y) = Q_{t-1}(v_x, v_y). \quad (7)$$

Finally, the posterior distribution at any time t is given again by the Bayes' rule:

$$Q_t(v_x, v_y) = L(v_x, v_y)P_t(v_x, v_y). \quad (8)$$

3.2.1.1. Analytical formulation. Because all the densities involved are normally distributed, their product is Gaussian and we can solve Eq. (8) analytically (Perrinet et al., 2006), with respect to the two velocity components $v_x(t)$ and $v_y(t)$ (not shown). More simply, by projecting the equation along two axes (tangent and perpendicular to the line respectively), there exist two simple coupled equations (one according to the orthogonal 'orth', one according to the tangent 'tang') allowing to explicit the mean and variance of the distribution on the left-hand side of Eq. (8) as a function of the mean and variance of the distributions on the right-hand side:

$$\begin{cases} \sigma_{Q, \text{orth}}^{-2} = (\sigma_1^{-2} + \sigma_2^{-2})(t/\tau) + \sigma_0^{-2}, \\ \sigma_{Q, \text{orth}}^{-2}(v_x \cos(\theta) + v_y \sin(\theta)) = (\sigma_1^{-2} + \sigma_2^{-2})(t/\tau)v_0 \cos(\theta), \end{cases} \quad (9)$$

$$\begin{cases} \sigma_{Q, \text{tang}}^{-2} = \sigma_2^{-2}(t/\tau) + \sigma_0^{-2}, \\ \sigma_{Q, \text{tang}}^{-2}(v_x \cos(\theta) - v_y \sin(\theta)) = \sigma_2^{-2}(t/\tau)v_0 \sin(\theta), \end{cases} \quad (10)$$

where τ is a time constant that corresponds to the duration of one iterative step in the discrete description of our model in Eq. (7) and (8). The constant τ serves also the purpose to keep coherent dimensions on both sides of the equations, namely the dimension of an inverse square velocity (first equation of each set) or that of an inverse velocity (second equation).

Finally, one can extract (v_x, v_y) as function of time from this set of two linear equations.

3.2.1.2. Parameter estimation to constrain the model. The variance terms σ_0^2 , σ_1^2 and σ_2^2 of Eqs. (1), (2) and (4) are free parameters in our model. We estimate them by applying Bayes' rule, similar to Eq. (5), to the pure 1D (long vertical

line) and *pure* 2D (central Gaussian blob) motion stimuli of our first experiment:

$$Q_{0,i}(v_x, v_y) = L_i(v_x, v_y)P_0(v_x, v_y), \quad (11)$$

with $i = 1$ or 2 depending on whether we displayed the 1D or 2D stimulus, respectively. Again, given that both the likelihoods and the prior are normal distributions, their product is also proportional to a normal distribution. Following the same reasoning leading to Eqs. (9) and (10), we can write:

$$\begin{cases} \mu_{Q_{0,i}} \sigma_{Q_{0,i}}^{-2} = \mu_i \sigma_i^{-2} + \mu_0 \sigma_0^{-2}, \\ \sigma_{Q_{0,i}}^{-2} = \sigma_i^{-2} + \sigma_0^{-2}. \end{cases} \quad (12)$$

The values $\mu_{Q_{0,i}}$ and $\sigma_{Q_{0,i}}^{-2}$ have been estimated from the oculomotor recordings for the 1D and 2D-stimulus respectively (see Section 3.1). The likelihood mean value μ_i is assumed to coincide with the stimulus speed v_0 (see Eqs. (1) and (2)), whereas the prior mean μ_0 is assumed to be 0 initially (see Eq. (4)). Therefore, the only unknowns left are σ_i^{-2} and σ_0^{-2} . The above pair of equations is a set of two linear equations with these two unknowns, thus leading to a straightforward unique solution. However, if we consider the two pairs of equations (for $i = 1, 2$), these equations are overdetermined because there are altogether only three unknowns σ_1^{-2} , σ_2^{-2} and σ_0^{-2} . The variance of the prior σ_0^2 is indeed assumed to be independent of the stimulus type (at least initially). In order to validate the self-consistency of the model we checked how well the model parameters estimated from one set of oculomotor measurements are capable of predicting the other set of measurements. In practice, the double set of Eqs. (12), for $i = 1, 2$, has a solution if and only if the following equality is valid (we have implicitly eliminated the term with μ_0 , since the prior mean is assumed to be 0 initially):

$$\sigma_{Q_{0,1}}^{-2} = \sigma_{Q_{0,2}}^{-2} \frac{\mu_2 - \mu_{Q_{0,2}}}{\mu_2} \frac{\mu_1}{\mu_1 - \mu_{Q_{0,1}}}. \quad (13)$$

This equality shows that the 1D-posterior variance (left-hand term) can be estimated by using the terms on the right-hand side, which are available from experimental measures or from the model assumptions. At the same time the 1D-posterior variance is itself an experimental observable (from the set of measurements with the 1D stimulus). We therefore compared, for each subject, model predictions and experimental measures of the 1D-posterior variance in all experimental conditions. In order to evaluate the goodness of fit for our model's predictions, we computed the coefficient of determination R^2 , defined as

$$R^2 = 1 - \frac{\text{SSE}}{\text{SST}}, \quad (14)$$

where SSE is the sum of squared errors (distance between predicted and measured datapoints) and SST is the total sum of squares for the measured variables. We obtained a coefficient of determination R^2 equal to 0.56, 0.33 and

0.41 for the three subjects respectively, these values being well within the desired range [0–1].

After this validation, in order to reduce measurement noise, we chose to estimate σ_0^{-2} as the average between the two independent estimates of this quantity obtained by solving the two sets of equations for $i = 1$ and 2 . The estimated Bayesian variables are shown in Fig. 6 for all three subjects and all target speed and direction conditions.

4. Results

Fig. 6 presents, for each subject and target motion direction, the estimated variance of the prior and the two independent likelihood distributions as a function of the target speed. It is important to underline that these are estimates of hidden variables which are supposed to characterise the internal inferential processes underlying motion integration. Because these variables are fully constrained by experimental data, they may provide a first general validation of the model. Fig. 6 deserves two comments: first, the estimated prior variance (green curves in Fig. 6) turns out to be roughly constant – as expected for a prior – across a threefold increase in target speed. Second, our estimates for the 1D and 2D likelihood variance (red and black curves, respectively) for different speed provide evidence of a monotonic increase of the variance with target speed. Note that this functional relationship was not included in the assumptions of the original theoretical model of velocity perception (Weiss et al., 2002), although it appears to be perfectly reasonable in light of Weber–Fechner's law.

An interesting attempt to quantitatively constrain the Bayesian framework for motion processing with psychophysical data has recently been published (Stocker and Simoncelli, 2006). Motivated by the weakness of the classical model in reproducing trial by trial variability in motion perception, Stocker and Simoncelli used experimental data in a speed discrimination task to infer the internal noise characteristics. One of their main findings was that the internal noise (i.e. the likelihood variance) is approximately proportional to the stimulus speed, in agreement with our results.

Our Bayesian recurrent model of the tracking error dynamics was implemented numerically in a Matlab (MathWorks) routine that took the parameters of the Bayesian prior and likelihood functions as input arguments. The discrete evolution of the posterior distributions, as well as of the MAP-defined current velocity estimate, were provided as output. At each time step we could derive the estimated tracking error (like for the example in the bottom panel of Fig. 5), as the inverse tangent of the ratio between vertical and horizontal velocity estimate.² For each subject, direction and speed, we have estimated the

² Given the use of an arbitrary time unit in the model simulations, all predictions about the time course of motion processing have to be rescaled along the temporal dimension.

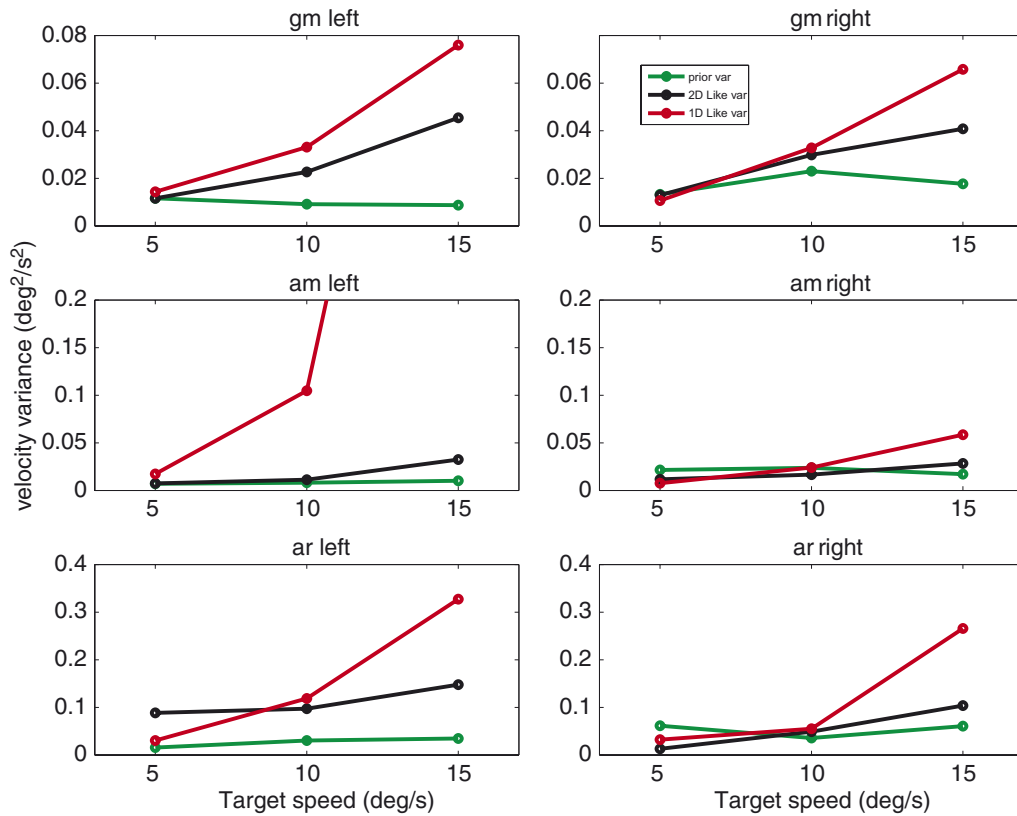


Fig. 6. Estimates of the prior and Likelihood variance for all subjects (rows), target motion directions (columns) and motion speed. Note that for subject AM ocular tracking of the long line with the highest speed was excessively noisy and dominated by catch-up saccades, making the variance estimate unreliable.

input parameters (see Section 3.2.1) and have run the simulations. In Fig. 7 the experimental average eye velocity and the model predicted estimated velocity are displayed (both for the horizontal and vertical component) for one subject. Qualitatively similar results were obtained for the other subjects (not shown).

We decided to introduce one free parameter (variable across subjects and motion directions but fixed for the three speed values), namely a gain parameter dividing the line likelihood variance. In practice, in Eq. (1), we substituted the term σ_1^2 with σ_1^2/g , leading to:

$$L_1(v_x, v_y) = \frac{1}{Z} \times \exp \left(- \frac{((v_x - v_0) \cos(\theta) + v_y \sin(\theta))^2}{2\sigma_1^2/g} \right). \quad (15)$$

The best value of the gain parameter g was determined as the one minimizing the prediction error for the peak of the vertical component. In practice, g was varied between 0.5 and 10 in steps of 0.5, while the sum (across speed conditions) of the squared difference between experimental and model-predicted peak vertical velocity was computed (and minimized). The need of a scaling parameter for the likelihood variance of the line is reasonable if

one thinks that the 1D-likelihood parameters were estimated for a very long line, necessarily implying a much larger energy in the motion signal when compared with the short line used in the tilted-line experiment. Interestingly, the optimal gain parameter was found to be similar for all subjects and directions and it ranged between 3 and 5, which turned out to be close to the length ratio between the two lines (~ 3). It is worth mentioning that only one 2D likelihood function was used in the simulation, even though one could argue that the short tilted line stimulus contains two non-ambiguous sources of information (one terminator at each end). While the short line contains two terminators, these 2D cues are more eccentric than the central circular Gaussian spot we used in experiment one, so the weakening effect of eccentricity arguably compensates the benefit of having two terminators.³

There are two apparent reasons for the mismatch between model predictions and experimental eye-velocity curves, as exemplified in Fig. 8. First, our model at this stage is able to predict the perceived velocity which is used

³ We tried experimentally to ask people to track two eccentric blobs simulating the two terminators of the short line stimulus but failed to obtain genuine smooth pursuit behaviours, presumably because there was no stimulus in the fovea. For this reason, in experiment one, we gave subjects a central circular spot to track instead.

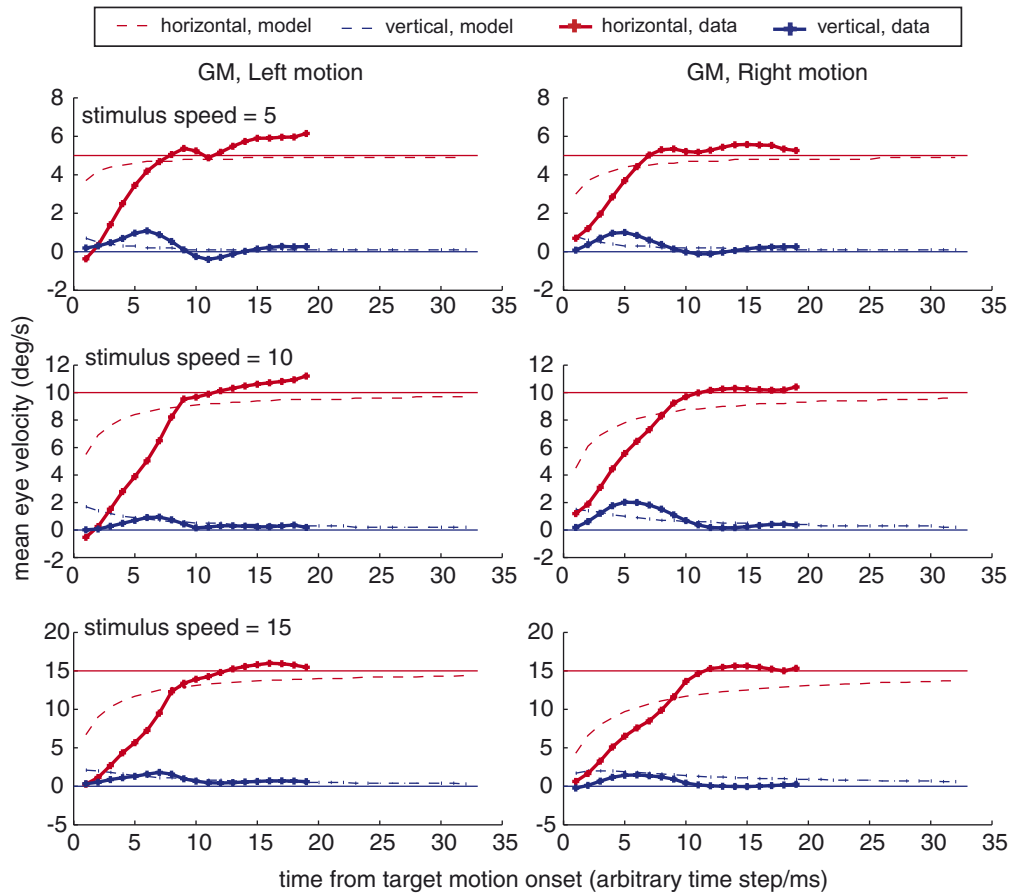


Fig. 7. Experimental eye movement velocity traces and model-predicted estimated velocities. The units on the time axis have to be regarded as arbitrary steps (for the model-predicted curves) or as 20 ms time shifts of the sliding average window for the experimental data. Time 0 represents pursuit onset. We have changed the sign of velocity curves corresponding to leftward motion (left panels) for an easier comparison with rightward motion curves.




Fig. 8. Mismatch between experimental oculomotor curves and model predictions: two simple factors explain the main differences (observable in the left panel). First, a shift along the time-axis accounts for the offset between data and model due to the fact that the model does not account for the time needed to accelerate from zero to the final pursuit velocity. In other words, the rising part of the blue solid curve is not predicted by the model, whose first iteration corresponds roughly to the time of peak vertical velocity. Second, the numerical simulations use an arbitrary time step, implying the need for a *time-stretching* adjustment (see right panel).

to drive smooth pursuit eye movements, but it does not account for the motor constraints. In other words, as already mentioned, our model does not predict motor latency (that we have also dismissed in the representation

of the experimental curves), nor does it reproduce the dynamics of acceleration of pursuit initiation (the model has virtually an *infinite* acceleration) and any other non-purely visual phenomena characterising smooth pursuit

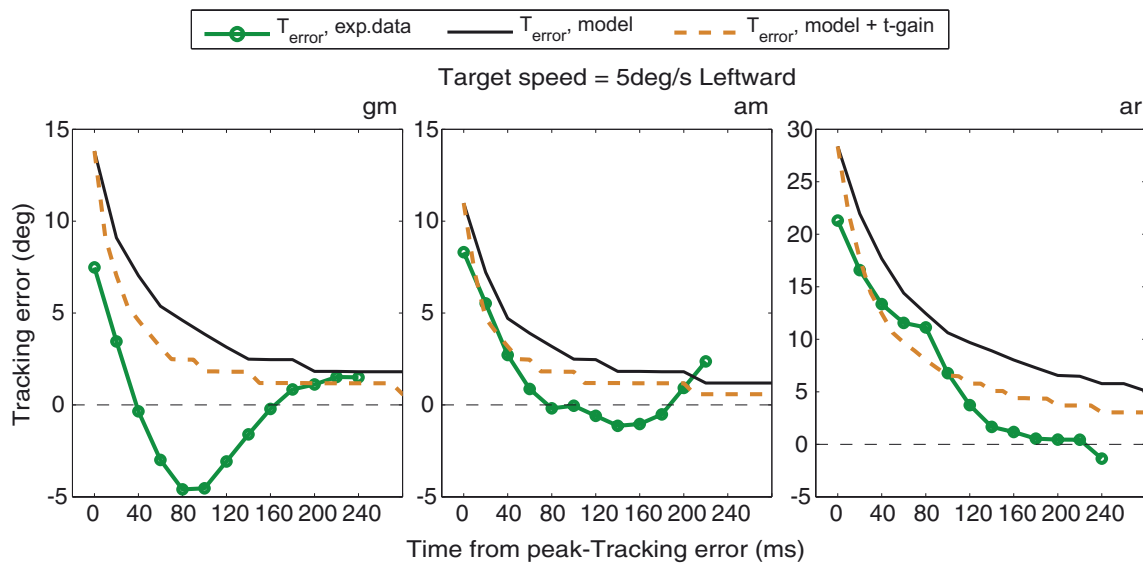


Fig. 9. Model prediction and experimental tracking error. The main adjustments applied to optimise the match between the simulated and the recorded tracking error are shown for speed = 5 deg/s and direction = Right and for all three subjects. Only the decreasing phase of the experimental tracking error is taken into account (i.e. time offset is already corrected). The green curve represents the original measured tracking error, whereas the black curve represents the original model prediction in arbitrary time steps. Finally, the orange-dashed curve accounts for an asymptotic vertical offset correction and a temporal *contraction* of the model prediction to better match the experimental data.

(see discussion in Section 5). This fact implies that at the very first iteration the model's prediction is consistent with the maximally biased velocity in the direction orthogonal to the line (it will progressively converge to the horizontal motion later); in contrast smooth pursuit eye velocity, although driven by the same visual input (the perceived biased velocity), needs a finite time to accelerate from rest to the final desired velocity (see rising part of the blue curve in Fig. 8). Second, the time constant is actually arbitrarily fixed in the recurrent model and therefore the predicted velocity curves have to be fitted to the experimental curves by means of a multiplicative temporal gain.

In Fig. 9 we present the tracking error for one particular combination of target speed and motion direction and for all three subjects. The original mean experimental tracking error can be compared with the model prediction and a transformed version of the latter, which accounts for a non-zero asymptotic vertical offset of the experimental data, as well as for the temporal rescaling needed to match the arbitrary model step to real time. The optimal adjustment for the model time step was obtained for a step corresponding to about 10 ms. This estimate was obtained by varying the number of model time steps (1, 2, 3, ...) that matched one experimental time step of 20 ms and minimizing the sum of squared errors between the experimental and the model-predicted velocity (both the horizontal and the vertical components).

5. Conclusions

Uncertainty in motion processing is reflected in the variability of the initial velocity of smooth pursuit eye movements. This type of eye movements provides also a reliable dynamic measure of the different contributions of

1D and 2D motion cues to motion integration. We have presented a simple model of motion integration dynamics, which is based on the idea of recursively updating the observer's prior about object motion by means of recent experience. The model is quantitatively constrained by oculomotor data and it replicates the well-established transient deviation of the gaze from the object's trajectory as well as its dynamical correction.

Our modeling approach represents one of the first extensions of the Bayesian framework, already widely applied in visual perception studies (Kersten et al., 2004; Weiss et al., 2002; Weiss and Fleet, 2002), to the oculomotor domain. Eye movement measurements have a consistent advantage with respect to other experimental measures like perceptual judgements, because they can be considered as a continuous dynamic account of the inference process output.

Our attempt to estimate the uncertainty of the internal inferential representation of motion information on the basis of smooth pursuit variability provided plausible quantitative results for the variance of prior and likelihood functions. Importantly, these estimates are in agreement with recent findings obtained with perceptual judgments and used to constrain a similar Bayesian model (Stocker and Simoncelli, 2006).

However, at the present stage, there is an apparent coarse mismatch between the model prediction and the observed ocular tracking traces. One simple reason can explain this mismatch: our model does not really account for the visuo-oculomotor transformations, nor for the oculomotor feedback information sent to the early visual stages during the closed-loop phase, i.e. during the late period of the tracking error correction. In Section 4 we have mentioned the need to adjust the model-predicted eye velocity to take account of the necessary temporal gain

factor and the temporal offset. The temporal offset is justified by the fact that ocular muscles can apply only a finite force on the eye, and therefore eye movements, different from our infinite-acceleration model, undergo necessarily an initial phase of velocity increase with finite acceleration. Only at the end of this phase it is reasonable to assume that eye-velocity variations (and tracking error) are related to velocity estimates. There exists at least another apparent difference between model predictions and oculomotor recordings, which can be explained in terms of the visuo-oculomotor loop characteristic dynamics. Eye-velocity curves tend to oscillate, with decreasing amplitude, around the final desired velocity (i.e. the target velocity, in the steady state). Oculomotor oscillations are due to delayed transmission in the visuo-oculomotor loop in conjunction with an online adaptation of pursuit gain (Goldreich et al., 1992): the retinal velocity error is translated into an oculomotor command aiming to correct it, but the latter is executed after a delay, thereby not accounting for the ongoing dynamics and introducing a new retinal velocity error with opposite sign, which has to be corrected and so on. In contrast, our model does clearly not suffer from this problem, given that a virtual zero-delay between visual input and model output is assumed, and the model curves reach the asymptotic velocity following a simple non-oscillating monotonic decay law. This fact, again, is a source of mismatch between our model predictions and the data.

We are currently planning to improve our model by simulating the missing stages of the visuo-oculomotor transformations. Classical models of the oculomotor plant based on Control-theory (Robinson et al., 1986; Krauzlis and Lisberger, 1994) provide a good estimate of the observed pursuit dynamics for a simple visual motion input. We will apply this type of oculomotor transformation to our dynamic velocity estimates, thereby obtaining true theoretical predictions for the oculomotor output. By doing so, we will be able to compare our experimental oculomotor data with congruent model predictions.

The model time step represents the internal delay for the recurrent loop, or, in other words, the time needed to update the internal inferred motion information. At this stage it is hard to judge about the significance of the best-fitting time constant (~ 10 ms), especially because we do not have a clear hypothesis about the neural implementation of our dynamic recurrent model of motion processing. However, cortical interaction between early visual areas (either within the primary visual cortex, V1, or between MT and V1) can occur over a typical time scale of 10–20 ms, which makes our estimate not completely unpalatable. Note that the dynamic role of the V1-MT loop in solving the aperture problem has already been proposed by other authors (Bayerl and Neumann, 2004). We believe, in general, that the application of our dynamic model to simulate specific observed experimental effects may provide interesting indications about the underlying cortical dynamical processes.

To conclude, a number of recent experimental findings (Wallace et al., 2005; Born et al., 2006) have provided important hints about the impact of several properties of moving stimuli on motion integration and the time course of the tracking error. In perspective, this experimental database constitutes a rich benchmark to test the validity and the robustness of our model and will probably allow to improve it.

Acknowledgements

We are deeply thankful to the patient volunteers who participated in the oculomotor experiments and in particular the naïve subject AR. Anna Montagnini was supported by a Marie Curie European Individual Fellowship.

References

- Albright, T., 1984. Direction and orientation selectivity of neurons in visual area MT of the macaque. *Journal of Neurophysiology* 52, 1106–1130.
- Bayerl, P., Neumann, H., 2004. Disambiguating visual motion through contextual feedback modulation. *Neural Computation* 16, 2041–2066.
- Born, R., Pack, C., Ponce, C., Yi, S., 2006. Temporal evolution of 2-dimensional direction signals used to guide eye movements. *Journal of Neurophysiology* 95, 284–300.
- Castet, E., Lorenceau, J., Shiffrar, M., Bonnet, C., 1993. Perceived speed of moving lines depends on orientation, length, speed and luminance. *Vision Research* 33, 1921–1936.
- Collewijn, H., van der Mark, F., Jansen, T., 1975. Precise recordings of human eye movements. *Vision Research* 15, 447–450.
- Gegenfurtner, K., Xing, D., Scott, B., Hawken, M., 2003. A comparison of pursuit eye movement and perceptual performance in speed discrimination. *Journal of Vision* 3, 865–876.
- Goldreich, D., Krauzlis, R., Lisberger, S., 1992. Effect of changing feedback delay on spontaneous oscillations in smooth pursuit eye movements of monkeys. *Journal of Neurophysiology* 67, 625–638.
- Hürlimann, F., Kiper, D., Carandini, M., 1992. Testing the Bayesian model of perceived speed. *Vision Research* 42, 2257–2553.
- Kalman, R., 1958. Design of self-optimizing control system. *Transactions of ASME* 80, 468–478.
- Kersten, D., Mamassian, P., Yuille, A., 2004. Object perception as Bayesian inference. *Annual Review Psychology* 55, 10.1–10.32.
- Krauzlis, R., Lisberger, S., 1994. A model of visually-guided smooth pursuit eye movements based on behavioral observations. *Journal of Computational Neuroscience* 1, 265–283.
- Lisberger, S., Morris, E., Tychsen, L., 1987. Visual motion processing and sensory-motor integration for smooth pursuit eye movements. *Annual Reviews in Neuroscience* 10, 97–129.
- Lorenceau, J., Shiffrar, M., 1992. The influence of terminators on motion integration across space. *Vision Research* 32, 263–273.
- Lorenceau, J., Shiffrar, M., Wells, N., Castet, E., 1993. Different motion sensitive units are involved in recovering the direction of moving lines. *Vision Research* 33, 1207–1217.
- Mamassian, P., Landy, M., 2001. Interaction of visual prior constraints. *Vision Research* 41, 2653–2668.
- Masson, G., Rybarczyk, Y., Castet, E., Mestre, D., 2000. Temporal dynamics of motion integration for the initiation of tracking eye movements at ultra-short latencies. *Visual Neuroscience* 17, 753–767.
- Masson, G., Stone, L., 2002. From following edges to pursuing object. *Journal of Neurophysiology* 88, 2869–2873.

- Montagnini, A., Spering, M., Masson, G., 2006. Combining 1D visual motion and 2D predictive signals to control smooth pursuit eye movements. *Journal of Neurophysiology* 96, 3545–3550.
- Osborne, L., Hohl, S., Bialek, W., Lisberger, S.G., 2007. Time course of precision in smooth-pursuit eye movements of monkeys. *The Journal of Neuroscience* 27, 2987–2998.
- Osborne, L., Lisberger, S., Bialek, W., 2005. A sensory source for motor variation. *Nature* 437, 412–416.
- Pack, C., Born, R., 2001. Temporal dynamics of a neural solution to the aperture problem in visual area MT of macaque brain. *Nature* 409, 1040–1042.
- Perrinet, L., Barthélemy, F., Castet, E., Masson, G., 2005. Dynamics of motion representation in short-latency ocular following: a two-pathways bayesian model. In: *Perception*, vol. 34, p. 38.
- Perrinet, L., Kremkow, J., Barthélemy, F., Masson, G., Chavane, F., 2006. Input–output transformation in the visuo-oculomotor loop: modeling the ocular following response to center-surround stimulation in a probabilistic framework. In: *FENS Meeting Proceedings* A216.1610.
- Priebe, N., Lisberger, S., 2004. Estimating target speed from the population response in visual area mt. *The Journal of Neuroscience* 24, 1907–1916.
- Rao, R., 2004. Bayesian computation in recurrent neural circuits. *Neural Computation* 16, 1–38.
- Robinson, D., Gordon, J., Gordon, S., 1986. A model of the smooth pursuit eye movement system. *Biological Cybernetics* 55, 43–57.
- Stocker, A., Simoncelli, E., 2006. Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience* 9, 578–585.
- Thompson, P., 1982. Perceived rate of movement depends on contrast. *Vision Research* 22, 377–380.
- Wallace, J., Stone, L., Masson, G., 2005. Object motion computation for the initiation of smooth pursuit eye movements in humans. *Journal of Neurophysiology* 93, 2279–2293.
- Weiss, Y., Adelson, E., 1998. Slow and smooth: a Bayesian theory for the combination of local motion signals in human vision. MIT Technical report, A.I. Memo 1624.
- Weiss, Y., Fleet, D., 2002. Velocity likelihoods in biological and machine vision. In: Rao, R., Olshausen, B., Levicki, M. (Eds.), *Probabilistic Models of the Brain*. MIT Press, pp. 77–96.
- Weiss, Y., Simoncelli, E., Adelson, E., 2002. Motion illusions as optimal percepts. *Nature Neuroscience* 5, 598–604.