



HAL
open science

Active inference, eye movements and oculomotor delays

Laurent Perrinet, Rick A Adams, Karl J Friston

► **To cite this version:**

Laurent Perrinet, Rick A Adams, Karl J Friston. Active inference, eye movements and oculomotor delays. *Biological Cybernetics (Modeling)*, 2014, 106 (8), pp.777-801. 10.1007/s00422-014-0620-8 . hal-01382350

HAL Id: hal-01382350

<https://hal.science/hal-01382350v1>

Submitted on 16 Oct 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Active inference, eye movements and oculomotor delays

Laurent U. Perrinet^{1,2}, Rick A. Adams² and Karl J. Friston²

¹Institut de Neurosciences de la Timone , CNRS / Aix-Marseille
Université - Marseille, France

²The Wellcome Trust Centre for Neuroimaging , University College
London - London, UK.

Keywords

Oculomotor delays; tracking eye movements; smooth pursuit eye movements; generalized coordinates; perception; Bayesian filtering; variational free energy; active inference

Abstract

This paper considers the problem of sensorimotor delays in the optimal control of (smooth) eye movements under uncertainty. Specifically, we consider delays in the visuo-oculomotor loop and their implications for active inference. Active inference uses a generalisation of Kalman filtering to provide Bayes optimal estimates of hidden states and action in generalized coordinates of motion. Representing hidden states in generalized coordinates provides a simple way of compensating for both sensory and oculomotor delays. The efficacy of this scheme is illustrated using neuronal simulations of pursuit initiation responses, with and without compensation. We then consider an extension of the generative model to simulate smooth pursuit eye movements — in which the visuo-oculomotor system believes both the target and its centre of gaze are attracted to a (hidden) point moving in the visual field. Finally, the generative model is equipped with a hierarchical structure, so that it can recognise and remember unseen (occluded) trajectories and emit anticipatory responses. These simulations speak to a straightforward and neurobiologically plausible solution to the generic problem of integrating information from different sources with different temporal delays and the particular difficulties encountered when a system — like the oculomotor system — tries to control its environment with delayed signals.

1 Introduction

1.1 Problem statement

This paper considers optimal motor control and the particular problems caused by the inevitable delay between the emission of motor commands and their sensory consequences. This is a generic problem that we illustrate within the context of oculomotor control where it is particularly prescient (see for instance (Nijhawan, 2008) for a review). Although we focus on oculomotor control, the more general contribution of this work is to treat motor control as a pure inference problem. This allows us to use standard (Bayesian filtering) schemes to resolve the problem of sensorimotor delays — by absorbing them into a generative or forward model. Furthermore, this principled and generic solution has some degree of biological plausibility because the resulting active (Bayesian) filtering is formally identical to predictive coding, which has become an established metaphor for neuronal message passing in the brain. We will use oculomotor control as a vehicle to illustrate the basic idea using a series of generative models of eye movements — that address increasingly complicated aspects of oculomotor control. In short, we offer a general solution to the problem of sensorimotor delays in motor control — using established models of message passing in the brain — and demonstrate the implications of this solution in the particular setting of oculomotor control.

The oculomotor system produces eye movements to deploy sensory (retinal) epithelia at very fast timescales. In particular, changes in the position of a visual object are compensated for with robust and rapid eye movements, such that the object is perceived as invariant, despite its motion (Ilg, 1997; Lisberger, Morris, and Tychsen, 1987). This near-optimal control is remarkable, given the absence of any external clock to coordinate dynamics in different parts of the visual-oculomotor system. An important constraint, in this setting, is axonal conduction, which produces delays in sensory and motor signalling within the oculomotor system. Figure 1 shows that in humans, for example, retinal signals arriving at motion processing areas report the state of affairs at least about 50 ms ago, while the action that follows is executed at least 40 ms in the future (Inui and Kakigi, 2006); for a review, see (Masson and Ilg, 2010). Different sources of delays exist — such as the biomechanical delay between neuromuscular excitation and eye movement. Due to these delays, the human smooth pursuit system responds to unpredictable stimuli with a minimum latency of around 100 ms (Wyatt and Pola, 1987). In addition, these delays may produce oscillations about a constant velocity stimulus (Robinson, 1965; Robinson, Gordon, and Gordon, 1986), whose amplitude and frequency can be altered by artificially manipulating the feedback (Goldreich, Krauzlis, and Lisberger, 1992).

Eye movements can anticipate predictable stimuli, such as the sinusoidal movement of a pendulum (Barnes and Asselman, 1991; Dodge, Travis, and Fox, 1930; Westheimer, 1954); for a review, see (Barnes, 2008). Interestingly, ocular tracking can compensate for sensorimotor delays after around one or two periods of sinusoidal motion — producing a tracking movement with little discernible delay (Barnes and Asselman, 1991). This suggests that the oculomotor system can use sensory information from the past to predict its future sensory states (including its actions), despite the fact that these sensory changes can be due to both movement of the stimulus and movement of the eyes. The

time taken to compensate for delays increases with the unpredictability of the stimulus (Michael and Jones, 1966), though the system can adapt quickly to complex waveforms, with changes of velocity (Barnes and Schmid, 2002), single cycles (Barnes, Barnes, and Chakraborti, 2000) or perturbed periodic waves – where subjects appear to estimate their frequency using an average over recent cycles (Collins and Barnes, 2009). Further studies suggest that different sources of information, such as auditory or verbal cues (Kowler, 1989) or prior knowledge about the nature of sensory inputs (Montagnini, Spering, and Masson, 2006) can evoke anticipatory eye movements.

The aim of this work was to establish a principled model of optimal visual motion processing and oculomotor control in the context of sensorimotor delays. Delays are often ignored in treatments of the visual-oculomotor system; however, they are crucial to understanding the system’s dynamics. For instance, delays may be important for understanding the pathophysiology of impaired oculomotor control: schizophrenic smooth pursuit abnormalities are due to impairments of the predictive (extra-retinal) motion signals that are required to compensate for sensorimotor delays (Nkam et al., 2010; Thaker et al., 1999). Surprisingly, delays may also explain a whole body of visual illusions (Changizi, 2001; Changizi and Widders, 2002; Changizi et al., 2008; Vaughn and Eagleman, 2013), even for visual illusions that involve a static display. Delays are also an important consideration in control theory and engineering. Finally, neuronal solutions to the delay problem speak to the representation of time in the brain, which is essential for the proper fusion of information in the central nervous system.

1.2 Existing solutions and the proposed hypothesis

A principled approach to optimal oculomotor control is provided by Bayesian filtering schemes that use probabilistic representations to estimate visual and oculomotor states. These states are *hidden*; i.e., they cannot be measured directly. A popular scheme for linear control problems is the Kalman filter (Kalman, 1960). The Kalman scheme can be extended to accommodate biomechanical constraints, such as transmission delays (e.g., fixed-lag smoothers). However, their solutions can become computationally complex when delays are large in relation to discretisation time and are not biologically plausible. We have previously considered generalized Bayesian filtering in continuous time as a metaphor for action and perception. This approach has been applied to eye movements (Friston et al., 2010b) and saccades in particular (Friston et al., 2012). However, these applications ignored sensorimotor delays and their potentially confounding effects on optimal control.

Crucially, the active inference schemes we have considered previously are formulated using representations in *generalized coordinates of motion*; that is, states (such as position) are represented along with their higher order temporal derivatives (such as speed, acceleration, jerk, etc). This means that one has an implicit representation of hidden states in the recent past and future that can be used to finesse the problems of delays. For example, it has been shown that acceleration is a necessary component of the predictive drive to eye movements (Bennett et al., 2007). In brief, generalized representations can be projected to the past and to the future using simple (linear) mixtures of generalized motion. Note that a representation of generalized motion can be explicit

or implicit by using a population coding scheme – as has been demonstrated for acceleration (Lisberger and Movshon, 1999). Representations of generalized motion may be important for modelling delays when integrating information in the brain from distal sources – such as other cortical columns in the same cortical area or other areas that are connected with fixed but different delays (Roelfsema et al., 1997). The integration of information over time becomes particularly acute in motor control, where the products of sensory processing couple back to the sampling of sensory information through action.

In the context of action, acted inference finesses the problems with delayed control signals in classical formulations of motor control by replacing command signals with descending cortico-spinal predictions. For instance, the location of receptive fields in the parietal cortex in monkeys is shown to shift transiently before an eye movement (Duhamel, Colby, and Goldberg, 1992). These predictions are fulfilled at the peripheral level, using fast closed loop mechanisms (peripheral reflex arcs). In principle, “these predictions can anticipate delays if they are part of the generative model,” (Friston, 2011): However, this anticipation has never been demonstrated formally. Here, we show how generalized Bayesian filtering – as used in active inference – can compensate for both sensory and motor delays in the visual-oculomotor loop.

It is important to mention what this work does not address. First, we focus on tracking eye movements (pursuit of a single dot stimulus for a monocular observer with fixed head position): we do not consider other types of eye movements (vergence, saccades, or the vestibulo-ocular reflex). Second, we take an approach that complements existing models, such as those of (Robinson, Gordon, and Gordon, 1986) and (Krauzlis and Lisberger, 1989). Existing models account for neurophysiological and behavioural data by refining block-diagram models of oculomotor control to describe *how* the system might work. We take a more generic approach, in which we define the imperatives for any system sampling sensory data, derive an optimal oculomotor control solution and show *why* this solution explains the data. Although the two approaches should be consistent, ours offers a principled approach to identifying the necessary solutions (such as predictive coding) to a given problem (oculomotor delays). We hope to demonstrate the approach by modelling pursuit initiation and smooth pursuit – and then consider the outstanding issue of anticipatory responses: in previous treatments (Robinson, Gordon, and Gordon, 1986), “[anticipation] has not been adequately modelled and no such attempt is offered (...) only unpredictable movements are considered.”

1.3 Outline

The main contributions of our work are described in the subsequent five sections. First, section 2 summarises the basic theory behind active inference and attempts to link generalized filtering to conventional Bayesian filters used in optimal control theory. This section then considers neurobiological implementations of generalized filtering, in terms of predictive coding in generalized coordinates of motion. This formulation allows us to consider the problem of delayed sensory input and motor output in section 3 – and how this problem can be finessed in a relatively straightforward way using generalized representations. Having established the formal framework (and putative neuronal implementation), the final three sections deal with successively harder problems

in oculomotor control. We start in Section 4 by considering *pursuit initiation* using a simple generative model of oculomotor trajectories. Using simulations, we consider the impact of motor delays, sensory delays and their interaction on responses to a single sweep of a visual target. The subsequent section turns to *smooth pursuit eye movements* – using a more sophisticated generative model of oculomotor trajectories, in which prior beliefs about eye movements enable the centre of gaze to predict target motion using a virtual or fictive target (see Section 5). In the final section, we turn to hierarchical models of target trajectories that have explicit memories of hidden dynamics, which enable anticipatory responses (see Section 6). These responses are illustrated using simulations of anticipatory pursuit movements using (rectified) hemi-sinusoidal motion. In short, these theoretical considerations lead to a partition of stimulus-bound eye movements into pursuit initiation, smooth pursuit and anticipatory pursuit, where each mode of oculomotor control calls on formal additions to the underlying generative model; however, they all use exactly the same scheme and basic principles. Where possible, we try to simulate classic empirical results in this field – at least heuristically.

In short, these theoretical considerations lead to a partition of stimulus-bound eye movements into pursuit initiation, smooth pursuit and anticipatory pursuit, where each mode of oculomotor control calls on formal additions to the underlying generative model. However, these models all use exactly the same scheme and basic principles; in particular, they all use the same solution to the oculomotor delay problem. These simulations illustrate that the active inference scheme can reproduce classical empirical results in three distinct experimental contexts.

2 From predictive coding to active inference

This section sets out the basic theory, before applying it to the special problem of oculomotor delays in the following sections. We first introduce the general framework of active inference in terms of generalized Bayesian filtering and variational free energy minimisation. In brief, active inference can be regarded as equipping standard Bayesian filtering schemes with classical reflex arcs that enable action, such as an eye movement, to fulfil predictions about hidden states of the world. Second, we will briefly describe the formalism of active inference in terms of differential equations describing the dynamics of the world and internal states of the visual-oculomotor system. The neurobiological implementation of these differential equations is considered in terms of predictive coding, which uses prediction errors on the motion of hidden states – such as the location of a visual target. In the next section, we will turn to the special problem of oculomotor delays and how this problem can be finessed using active inference in generalized coordinates of motion. This solution will be illustrated in subsequent sections using simulations of pursuit initiation responses and smooth pursuit. Finally, we shall exploit the richness of hierarchical generative models – which underlie active inference – to illustrate anticipatory eye movements that cannot be explained by simply compensating for oculomotor delays.

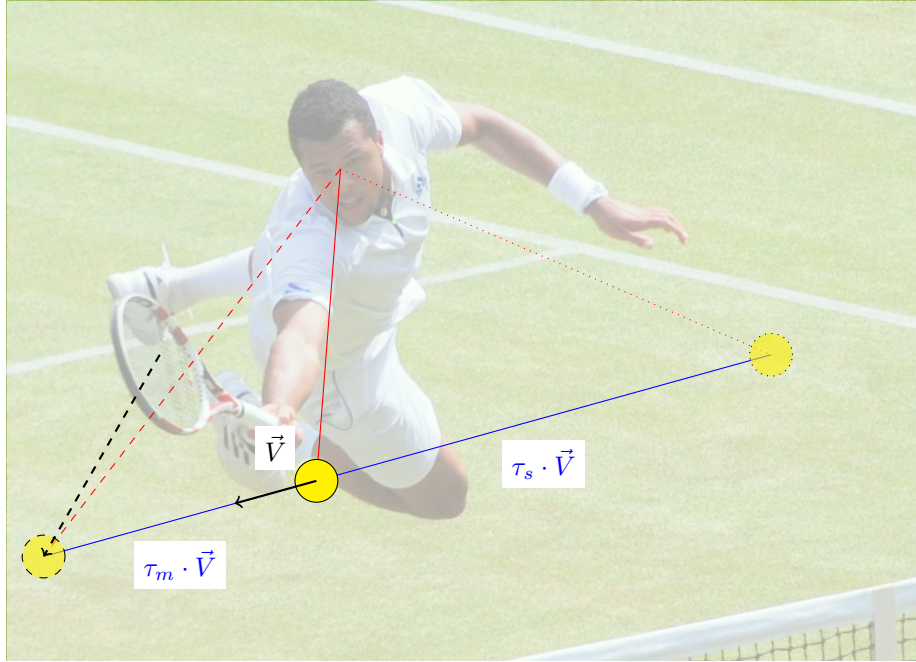


Figure 1: Problem statement: optimal motor control under axonal delays. The central nervous system has to contend with axonal delays, both at the sensory and the motor levels. For instance, in the human visuo-oculomotor system, it takes approximately $\tau_s = 50$ ms for the retinal image to reach the visual areas implicated in motion detection, and a further $\tau_m = 40$ ms to reach the oculomotor muscles. As a consequence, for a tennis player trying to intercept a ball at a speed of $20 \text{ m} \cdot \text{s}^{-1}$, the sensed physical position is 1 m behind the true position (as represented here by $\tau_s \cdot \vec{V}$), while the position at the moment of emitting the motor command will be .8 m ahead of its execution ($\tau_m \cdot \vec{V}$). Note that while the actual position of the ball when its image produced by the photoreceptors on the retina hits visual areas is approximately at 45 degrees of eccentricity (red dotted line), the player's gaze is directed to the ball at its *present* position (red line), in anticipatory fashion. Optimal control directs action (future motion of the eye) to the expected position (red dashed line) of the ball in the future — and the racket (black dashed line) to the expected position of the ball when motor commands reach the periphery (muscles).

Exchange with the environment

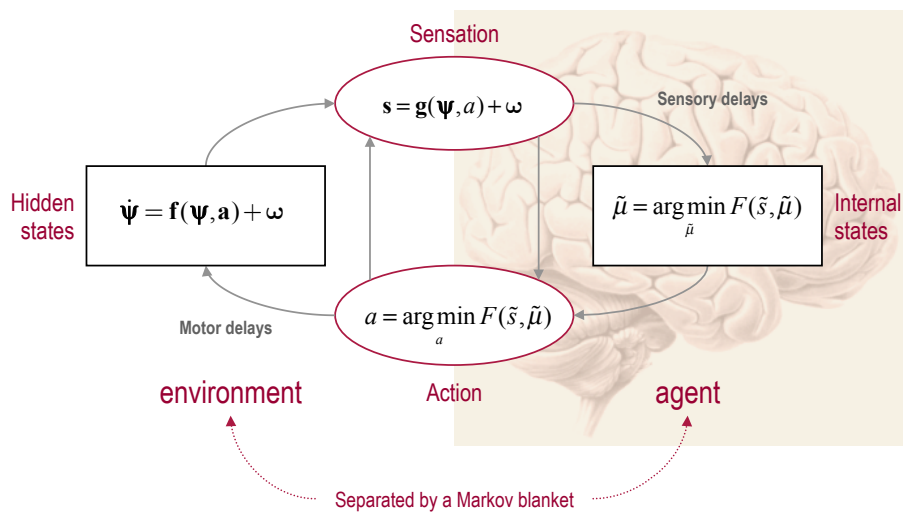


Figure 2: This schematic shows the dependencies among various quantities modelling exchanges of an agent with the environment. It shows the states of the environment and the system in terms of a probabilistic dependency graph, where connections denote directed dependencies. The quantities are described within the nodes of this graph – with exemplar forms for their dependencies on other variables (see main text). Hidden (external) and internal states of the agent are separated by action and sensory states. Both action and internal states – encoding a conditional probability density function over hidden states – minimise free energy. Note that hidden states in the real world and the form of their dynamics can be different from that assumed by the generative model; this is why hidden states are in bold. See main text for further details.

2.1 From free energy to generalized filtering

The scheme used here to model oculomotor behaviour has been used to model several other processes and paradigms in neuroscience. This active inference scheme is based upon three assumptions:

- The brain minimises the free energy of sensory inputs defined by a generative model.
- The generative model used by the brain is hierarchical, nonlinear and dynamic.
- Neuronal firing rates encode the expected state of the world, under this model.

The first assumption is the free energy principle, which leads to active inference in the context of an embodied interaction of the system with its environment; where the system can act to change its sensory inputs. The free energy here is a variational free energy that provides a computationally tractable upper bound on the negative logarithm of Bayesian model evidence (see Appendix 1). In Bayesian terms, this means that the brain maximises the evidence for its model of sensory inputs (Ballard, Hinton, and Sejnowski, 1983; Bialek, Nemenman, and Tishby, 2001; Dayan et al., 1995; Gregory, 1980; Grossberg et al., 1997; Knill and Pouget, 2004; Olshausen and Field, 1996). This is the Bayesian brain hypothesis (Yuille and Kersten, 2006). If we also allow action to maximise model evidence, we get active inference (Friston et al., 2010b). Crucially, unlike conventional optimal control schemes, there is no *ad hoc* value or loss function guiding action: Action minimises the free energy of the system’s model. This permits the application of standard Bayesian solutions and simplifies the implicit neuronal architecture; for example, there is no need for an efference copy signal (Friston, 2011). In this setting, desired movements are specified in terms of prior beliefs about state transitions or the motion of hidden states in the generative model. Action then realises prior beliefs (policies) by sampling sensory data that provides evidence for those beliefs.

The second assumption above is motivated by noting that the world is both dynamic and nonlinear – and that hierarchical structure emerges inevitably from a separation of temporal scales (Ginzburg, 1955; Haken, 1983). The third assumption is the Laplace assumption that, in terms of neural codes, leads to the *Laplace code*, which is arguably the simplest and most flexible of all neural codes (Friston, 2009). In brief, the Laplace code means that probabilistic representations are encoded explicitly by synaptic activity in terms of their mean or expectation (while the second order statistics such as dispersion or precision are encoded implicitly in terms of synaptic activity and efficacy). This limits the representation of hidden states to continuous variables, as opposed to discrete states; however, this is appropriate for most aspects of sensorimotor processing. Furthermore, it finesses the combinatoric explosion associated with discrete state space models. Restricting probabilistic representations to a Gaussian form clearly precludes multimodal representations. Having said this, the hierarchical form of the generative models allows for fairly graceful modelling of nonlinear effects (such as shadows and occlusions). For example, a Gaussian variable at one level of the model may enter the lower levels in highly non-linear

way — we will see examples of this later. See Appendix 2 for a motivation of the Laplace assumption from basic principles.

Under these assumptions, action and perception can be regarded as the solutions to coupled differential equations describing the dynamics of the real world (the first pair of equations) and the behaviour of an agent (the second pair of equations); expressed in terms of action and internal states that encode conditional expectations about hidden states of the world (Friston et al., 2010b):

$$\begin{aligned}
 \mathbf{s} &= \mathbf{g}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{a}) + \boldsymbol{\omega}_{\nu} \\
 \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x}, \boldsymbol{\nu}, \mathbf{a}) + \boldsymbol{\omega}_{\mathbf{x}} \\
 \dot{a} &= -\partial_a F(\tilde{s}, \tilde{\mu}) \\
 \dot{\tilde{\mu}} &= \mathcal{D}\tilde{\mu} - \partial_{\tilde{\mu}} F(\tilde{s}, \tilde{\mu})
 \end{aligned} \tag{1}$$

For clarity, real-world states are written in boldface, while internal states of the agent are in italics: Here $(\mathbf{s}, \mathbf{x}, \boldsymbol{\nu}, \mathbf{a})$ denote sensory input, hidden states, hidden causes and action in the real world, respectively. The variables in the second pair of equations $(\tilde{s}, \tilde{\mu}, a)$ correspond to generalized sensory input, conditional expectations and action. Generalized coordinates of motion, denoted by the $\tilde{\cdot}$ notation, correspond to a vector representing the different orders of motion of a variable: position, velocity, acceleration, and so on (Friston et al., 2010a). Using the Lagrangian notation for temporal derivatives, we get e. g. for s : $\tilde{s} = (s, s', s'', \dots)$. In the absence of delays $\tilde{s}(t) = \tilde{\mathbf{s}}(t)$ the agent receives instantaneous sensations from the real world. The differential equations above are coupled because sensory states depend upon action through hidden states and causes $(\mathbf{x}, \boldsymbol{\nu})$ while action $a(t) = \mathbf{a}(t)$ depends upon sensory states through internal states $\tilde{\mu}$.

By explicitly separating real-world states –hidden from the agent– to its internal states, one can clearly separate the generative model from the updating scheme that allows to minimise the agent’s free-energy: The first pair of coupled stochastic differential equations describes the dynamics of hidden states and causes in the world and how they generate sensory states. These equations are stochastic because sensory states and the motion of hidden states are subject to random fluctuations $(\boldsymbol{\omega}_{\mathbf{x}}, \boldsymbol{\omega}_{\nu})$.

The second pair of differential equations corresponds to action and perception respectively – they constitute a (generalized) gradient descent on variational free energy. The differential equation describing changes in conditional expectations (perception) is known as *generalized filtering* or predictive coding and has the same form as standard Bayesian (Kalman-Bucy) filters – see also (Beal, 2003; Rao and Ballard, 1999). The first term is a prediction based upon a differential operator \mathcal{D} that returns the generalized motion of the conditional expectations; namely the vector of velocity, acceleration, jerk and so on – such that $\mathcal{D}\tilde{\mu} = (\mu', \mu'', \mu''', \dots)$. However, the expected velocity is not the velocity of the expectation and comprises both prediction and update terms: The second term reflects this correction and ensures the changes in conditional expectations are Bayes-optimal predictions of hidden states of the world – in the sense that they maximise the free energy bound on Bayesian model evidence. See Figure 2 for a schematic summary of the implicit conditional dependencies implied by Equation 1.

2.2 Hierarchical form of the generative model

To perform simulations using this scheme, one simply integrates or solves Equation 1 to simulate (neuronal) dynamics that encode conditional expectations and ensuing action. Conditional expectations depend upon a generative model, which we assume has the following (hierarchical) form

$$\begin{aligned}
 s &= g^{(1)}(x^{(1)}, v^{(1)}) + \omega_\nu^{(1)} \\
 \dot{x}^{(1)} &= f^{(1)}(x^{(1)}, v^{(1)}) + \omega_x^{(1)} \\
 &\vdots \\
 \nu^{(i-1)} &= g^{(i)}(x^{(i)}, v^{(i)}) + \omega_\nu^{(i)} \\
 \dot{x}^{(i)} &= f^{(i)}(x^{(i)}, v^{(i)}) + \omega_x^{(i)} \\
 &\vdots
 \end{aligned} \tag{2}$$

Where (i) indexes the level in the hierarchical model. Note that we denote the sensory layer as $i = 0$, but this indexing is somewhat arbitrary. This equation is just a way of writing down a generative model that specifies a probability density function over sensory inputs and hidden states and causes. This probability density is needed to define the free energy of sensory input: it is specified in terms of functions $(f^{(i)}, g^{(i)})$ and Gaussian assumptions about random fluctuations $(\omega_x^{(i)}, \omega_\nu^{(i)})$ on the motion of hidden states and causes. It is these that make the model probabilistic – they play the role of sensory noise at the first level and induce uncertainty about states at higher levels. The precisions of these fluctuations are quantified by $(\Pi_x^{(i)}, \Pi_\nu^{(i)})$ which are defined as the inverse of the respective covariance matrices.

The deterministic part of the model is specified by nonlinear functions of hidden states and causes $(f^{(i)}, g^{(i)})$ that generate dynamics and sensory consequences. Hidden causes link hierarchical levels, whereas hidden states link dynamics over time. Hidden states and causes are abstract quantities that the brain uses to explain or predict sensations – like the motion of an object in the field of view. In hierarchical models of this sort, the output of one level acts as an input to the next. This input can produce complicated convolutions with deep (hierarchical) structure. We will see examples of this later in particular in the context of anticipatory movements.

2.3 Perception and predictive coding

Given the form of the generative model (Equation 2) one can write down the differential equations (Equation 1) describing neuronal dynamics in terms of prediction errors on the hidden causes and states. These errors represent the difference between conditional expectations and predicted values, under the generative model (using $A \cdot B := A^T B$ for the scalar product and omitting higher-order terms):

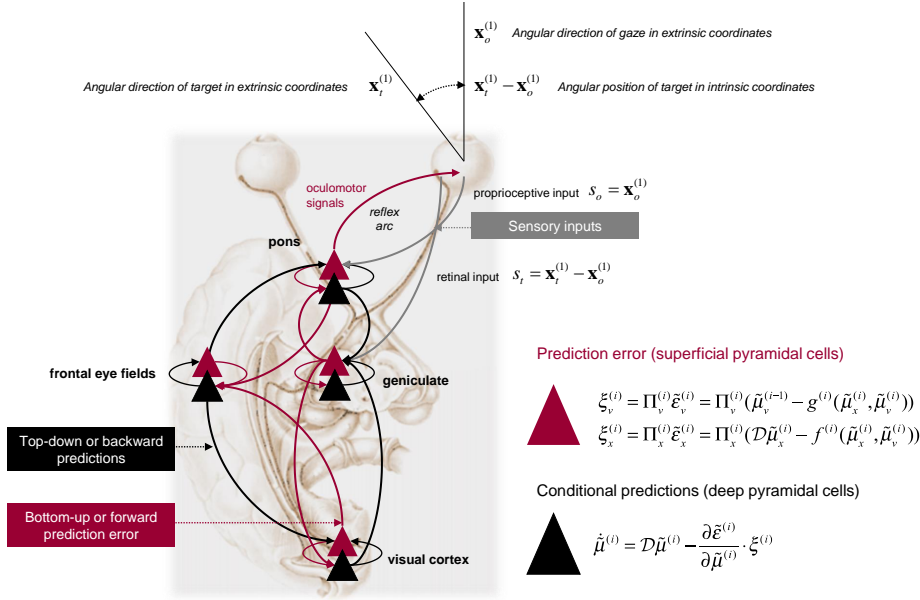


Figure 3: Schematic detailing a neuronal message passing scheme (generalized Bayesian filtering or predictive coding) that optimises conditional expectations about hidden states of the world, given sensory (visual) data and the active (oculomotor) sampling of those data. This diagram shows the speculative cells of origin of forward driving connections (in red) that convey prediction error from a lower area to a higher area and the backward connections (in black) that construct predictions (Mumford, 1992). These predictions try to explain away prediction error in lower levels. In this scheme, the sources of forward and backward connections are superficial (red) and deep (black) pyramidal cells respectively. The equations on the right represent a generalized descent on free energy under the hierarchical model described in the main text – this can be regarded as a generalisation of predictive coding or Kalman filtering: see (Friston, 2008). State-units are in black and error-units are in red. Here, we have placed different levels of some hierarchical model within the visual-oculomotor system. Visual input arrives in an intrinsic (retinal) frame of reference that depends upon the angular position of a stimulus and the direction of gaze. Exteroceptive input is then passed to the lateral geniculate nuclei (LGN) and to higher visual and prefrontal (e.g., motion sensitive, such as the frontal eye field) areas in the form of prediction errors. Crucially, proprioceptive sensations are also predicted, creating prediction errors at the level of the cranial nerve nuclei (pons). The special aspect of these proprioceptive prediction errors is that they can be resolved through classical reflex arcs – in other words, they can elicit action to change the direction of gaze and close the visual–oculomotor loop.

$$\begin{aligned}
\dot{\tilde{\mu}}_x^{(i)} &= \mathcal{D}\tilde{\mu}_x^{(i)} + \frac{\partial \tilde{g}^{(i)}}{\partial \tilde{\mu}_x^{(i)}} \cdot \Pi_\nu^{(i)} \tilde{\varepsilon}_\nu^{(i)} + \frac{\partial \tilde{f}^{(i)}}{\partial \tilde{\mu}_x^{(i)}} \cdot \Pi_x^{(i)} \tilde{\varepsilon}_x^{(i)} - \mathcal{D}\Pi_x^{(i)} \tilde{\varepsilon}_x^{(i)} \\
\dot{\tilde{\mu}}_\nu^{(i)} &= \mathcal{D}\tilde{\mu}_\nu^{(i)} + \frac{\partial \tilde{g}^{(i)}}{\partial \tilde{\mu}_\nu^{(i)}} \cdot \Pi_\nu^{(i)} \tilde{\varepsilon}_\nu^{(i)} + \frac{\partial \tilde{f}^{(i)}}{\partial \tilde{\mu}_\nu^{(i)}} \cdot \Pi_x^{(i)} \tilde{\varepsilon}_x^{(i)} - \Pi_\nu^{(i+1)} \tilde{\varepsilon}_\nu^{(i+1)} \\
\tilde{\varepsilon}_x^{(i)} &= \mathcal{D}\tilde{\mu}_x^{(i)} - \tilde{f}^{(i)}(\tilde{\mu}_x^{(i)}, \tilde{\mu}_\nu^{(i)}) \\
\tilde{\varepsilon}_\nu^{(i)} &= \tilde{\mu}_\nu^{(i-1)} - \tilde{g}^{(i)}(\tilde{\mu}_x^{(i)}, \tilde{\mu}_\nu^{(i)})
\end{aligned} \tag{3}$$

The quantities $\tilde{\varepsilon}^{(i)}$ correspond to prediction errors (on hidden states x or hidden causes ν). These are weighted by their respective precision vectors $\Pi^{(i)}$ in the update scheme. Equation 3 can be derived fairly easily by computing the free energy for the hierarchical model in Equation 2 and inserting its gradients into Equation 1. This gives a relatively simple update scheme, in which conditional expectations are driven by a mixture of prediction errors, where prediction errors are defined by the equations of the generative model.

It is difficult to overstate the generality and importance of Equation 3 – its solutions grandfather nearly every known statistical estimation scheme, under parametric assumptions about additive noise (Friston, 2008). These range from ordinary least squares to advanced variational deconvolution schemes. In this form, one can see clearly the relationship between predictive coding and Kalman-Bucy filtering – changes in conditional expectations comprise a prediction (first term) plus a weighted mixture of prediction errors (remaining terms). The weights play the role of a Kalman gain matrix and are based on the gradients of the model functions and the precision of random fluctuations.

In neural network terms, Equation 3 says that error-units receive predictions from the same hierarchical level and the level above. Conversely, conditional expectations (encoded by the activity of state units) are driven by prediction errors from the same level and the level below. These constitute bottom-up and lateral messages that drive conditional expectations towards a better prediction to reduce the prediction error in the level below. This is the essence of recurrent message passing between hierarchical levels to suppress free energy or prediction error: see (Friston and Kiebel, 2009) for a more detailed discussion. In neurobiological implementations of this scheme, the sources of bottom-up prediction errors, in the cortex, are thought to be superficial pyramidal cells that send forward connections to higher cortical areas. Conversely, predictions are conveyed from deep pyramidal cells by backward connections, to target (polysynaptically) the superficial pyramidal cells encoding prediction error (Friston and Kiebel, 2009; Mumford, 1992). This defines an elementary circuit that may be the basis of the layered organisation of the cortex (Bastos et al., 2012). Figure 3 provides a schematic of the proposed message passing among hierarchically deployed cortical areas.

2.4 Action

In active inference, conditional expectations elicit behaviour by sending predictions down the hierarchy to be unpacked into proprioceptive predictions at the

level of (pontine) cranial nerve nuclei and spinal-cord. These engage classical reflex arcs to suppress proprioceptive prediction errors and produce the predicted motor trajectory

$$\dot{a} = -\partial_a F = -(\partial_a \tilde{\varepsilon}_\nu^{(1)}) \cdot \Pi_\nu^{(1)} \tilde{\varepsilon}_\nu^{(1)} \quad (4)$$

The reduction of action to classical reflexes follows because the only way that action can minimise free energy is to change sensory (proprioceptive) prediction errors by changing sensory signals. This highlights the tight relationship between action and perception; cf., the equilibrium point formulation of motor control (Feldman and Levin, 1995). In short, active inference can be regarded as equipping a generalized predictive coding scheme with classical reflex arcs: see (Friston, Daunizeau, and Kiebel, 2009; Friston et al., 2010b) for details. The actual movements produced clearly depend upon (changing) top-down predictions that can have a rich and complex structure. This scheme is consistent with the physiology and anatomy of the oculomotor system (for a review see (Ilg, 1997; Krauzlis, 2004); although our goal here is not to identify the role of each anatomical structure but rather to give a schematic proof-of-concept.

2.5 Summary

In summary, we have derived equations for the dynamics of perception and action using a free energy formulation of adaptive (Bayes-optimal) exchanges with the world and a generative model that is both generic and biologically plausible. A technical treatment of the material above will be found in (Friston et al., 2010a), which provides the details of the generalized filtering used to produce the simulations in the next section. Before looking at these simulations, we consider how delays can be incorporated into this scheme.

3 Active inference with sensorimotor delays

If action and sensations were not subject to delays, one could integrate (solve) Equation 1 directly; however, in the presence of sensory and motor delays (τ_s and τ_a respectively) Equation 1 becomes a (stochastic and non-linear) delay differential equation because $\tilde{s}(t) = \tilde{s}(t - \tau_s)$ and $a(t) = \mathbf{a}(t + \tau_a)$. In other words, the agent receives sensations from (sees) the past, whilst emitting motor signals that will be enacted in the future (we will only consider delays from the sensory and motor sub-systems and neglect delays between neuronal systems in this paper).

To finesse the integration of these delay differential equations one can exploit their formulation in generalized coordinates: By taking linear mixtures of generalized motion one can easily map from the present to the future, using the matrix operators:

$$T(\tau) = \exp(\tau\mathcal{D}) = \begin{bmatrix} 1 & \frac{1}{1!}\tau & \frac{1}{2!}\tau^2 & \cdots \\ 0 & 1 & \frac{1}{1!}\tau & \cdots \\ 0 & 0 & 1 & \ddots \\ 0 & 0 & 0 & \ddots \end{bmatrix} \quad (5)$$

$$\text{with } \mathcal{D} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & \ddots \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

The first differential operator simply returns the generalized motion $\mathcal{D}\tilde{x}(t) = \tilde{x}'(t)$ while the second delay operator produces generalized states in the future $T(\tau)\tilde{x}(t) = \tilde{x}(t+\tau)$ (we define delays as positive by convention). Note that shifting forward and backwards by the same amount of time produces the identity operator $T(\tau)T(-\tau) = I$ and that, more generally, $T(\tau_1)T(\tau_2) = T(\tau_1 + \tau_2)$.

These delay operators are simple to implement computationally (and neurobiologically) and allow an agent to finesse the delayed coupling above by replacing (delayed) sensory signals with future input $\tilde{s}(t) = T(\tau_s)\tilde{s}(t - \tau_s) = \tilde{s}(t)$ for subsequent action and perception. Alternatively, one can regard this compensation for sensory delays as attempting to predict the past (see below). Generalized coordinates allow the representation of the trajectory of a given variable at any time (that is, its evolution in the near past and present) and thus allow its projection into the future or past. Generalized representations are more extensive than ‘snapshots’ at a particular time and enable the agent to anticipate the future (of delayed sensory trajectories) and represent hidden states in real time – that is, representations that are synchronised with the external events. In terms of motor delays, the agent can replace its internal motor signals with action in the future $\mathbf{a}(t) = T(\tau_a)\mathbf{a}(t - \tau_a) = \mathbf{a}(t)$, such that when action signals reach the periphery they correspond to the action encoded centrally. These substitutions allow us to express action and perception in Equation 1 as¹:

$$\begin{aligned} \dot{\mathbf{a}}(t) &= -\partial_{\mathbf{a}}F(T(\tau_a)T(\tau_s)\tilde{\mathbf{s}}(t - \tau_s - \tau_a), T(\tau_a)\tilde{\boldsymbol{\mu}}(t - \tau_a)) \\ &= -\partial_{\mathbf{a}}F(T(\tau_s - \tau_s + \tau_a - \tau_a)\tilde{\mathbf{s}}(t), T(\tau_a - \tau_a)\tilde{\boldsymbol{\mu}}(t)) \end{aligned} \quad (6)$$

$$\begin{aligned} \dot{\tilde{\boldsymbol{\mu}}}(t) &= \mathcal{D}\tilde{\boldsymbol{\mu}}(t) - \partial_{\tilde{\boldsymbol{\mu}}}F(T(\tau_s)\tilde{\mathbf{s}}(t - \tau_s), \tilde{\boldsymbol{\mu}}(t)) \\ &= \mathcal{D}\tilde{\boldsymbol{\mu}}(t) - \partial_{\tilde{\boldsymbol{\mu}}}F(T(\tau_s - \tau_s)\tilde{\mathbf{s}}(t), \tilde{\boldsymbol{\mu}}(t)) \end{aligned}$$

This equation distinguishes between true delays (τ) and those assumed by the agent (τ). When the two are the same, the delay operators $T(\tau - \tau) = I : \tau = \tau$ become identity matrices and Equation 6 reduces to Equation 1. When the two differ, Equation 6 permits the simulation of a system with uncompensated delays. Notice how the dynamics of action in the first differential equation are

¹We have made a slight approximation here because $T(\tau_a)\tilde{\boldsymbol{\mu}}(t - \tau_a) = T(\tau_a - \tau_a)\tilde{\boldsymbol{\mu}}(t)$ when, and only when, the free energy gradients are zero and $\tilde{\boldsymbol{\mu}}(t) = \mathcal{D}\tilde{\boldsymbol{\mu}}(t)$. Under the assumption that the perceptual destruction of these gradients is fast, in relation to action, this can be regarded as an adiabatic approximation.

driven by a gradient descent on the free energy of sensations with composite sensory and motor delays. In other words, action in the real world depends upon sensory states generated $\tau_s + \tau_a$ in the past.

One can now solve Equation 6 to simulate active inference, with or without compensation for sensorimotor delays. We use a standard local linearisation scheme for this integration (Ozaki, 1992), where delays enter at the point at which sensory prediction error is computed and when it drives action: from Equations 3 and 4:

$$\begin{aligned}\tilde{\varepsilon}_\nu^{(1)} &= T(\tau_s)\tilde{\mathbf{s}}(t - \tau_s) - \tilde{g}^{(1)}(\tilde{\mu}_x^{(1)}, \tilde{\mu}_\nu^{(1)}) \\ &= T(\tau_s - \tau_a)\tilde{\mathbf{s}}(t) - \tilde{g}^{(1)}(\tilde{\mu}_x^{(1)}, \tilde{\mu}_\nu^{(1)}) \\ \dot{\mathbf{a}}(t) &= -(\partial_a \tilde{\varepsilon}_\nu^{(1)}) \cdot \Pi_\nu^{(1)} T(\tau_a) \tilde{\varepsilon}_\nu^{(1)}(t - \tau_a) \\ &= -(\partial_a \tilde{\varepsilon}_\nu^{(1)}) \cdot \Pi_\nu^{(1)} T(\tau_a - \tau_a) \tilde{\varepsilon}_\nu^{(1)}(t)\end{aligned}\tag{7}$$

Equation 7 means that perfect (errorless) prediction requires $T(\tau_s)\tilde{\mathbf{s}}(t - \tau_s) = \tilde{g}^{(1)}(\tilde{\mu}_x^{(1)}, \tilde{\mu}_\nu^{(1)})$. In other words, errorless prediction means that the agent is effectively predicting the future projection of the past. Note again the dependency of action, via prediction errors, on sensory states $\tau_s + \tau_a$ in the past. See Appendix 3 for further details of the integration scheme used in the simulations below.

3.1 Summary

This section has considered how the differential equations describing changes in action and internal (representational) states can be finessed to accommodate sensorimotor delays. This is relatively straightforward – in the context of generalized schemes – using delay operators that take mixtures of generalized motion to project states into the future or past. Sensory delays can be (internally) simulated and corrected by applying delays to sensory input producing sensory prediction error, while motor delays can be simulated and corrected by applying delays to sensory prediction error producing action. Neurobiologically, the application of delay operators just means changing synaptic connection strengths to take different mixtures of generalized sensations and their prediction errors. We will now use these operators to look at the effects of sensorimotor delays with and without compensation.

4 Results: pursuit initiation

This section focuses on the consequences of sensory delays, motor delays and their combination – in the context of pursuit initiation – using perhaps the simplest generative model for active inference. Our purpose is to illustrate the difficulties in oculomotor control that are incurred by delays and how these difficulties dissolve when delays are accommodated during active inference. We start with a description of the generative model and demonstrate its behaviour when delays are compensated. We then use this normal behaviour as a reference to look at failures of pursuit initiation induced by delays. In this section, responses

to a single sweep of rightward motion are used to illustrate basic responses. In the next section, we consider pursuit of sinusoidal motion (with abrupt onsets) and the implications for generative models that may be used by the brain.

4.1 Generative model of pursuit initiation

The generative model for pursuit initiation used here is very simple and is based upon the prior belief that the centre of gaze is attracted to the target location. The processes generating sensory inputs and the associated generative model can be expressed as follows:

$$\begin{aligned}
 \mathbf{s} &= \begin{bmatrix} \mathbf{s}_o \\ \mathbf{s}_t \end{bmatrix} = \begin{bmatrix} \mathbf{x}_o \\ \mathbf{x}_t - \mathbf{x}_o \end{bmatrix} + \boldsymbol{\omega}_\nu^{(1)} \\
 \dot{\mathbf{x}} &= \begin{bmatrix} \dot{\mathbf{x}}_o \\ \dot{\mathbf{x}}_t \end{bmatrix} = \begin{bmatrix} \frac{1}{t_a} a - \frac{1}{t_o} \mathbf{x}_o \\ \frac{1}{t_m} (\boldsymbol{\nu}^{(1)} - \mathbf{x}_t) \end{bmatrix} + \boldsymbol{\omega}_x^{(1)} \\
 & \hspace{20em} (8) \\
 s &= \begin{bmatrix} s_o \\ s_t \end{bmatrix} = \begin{bmatrix} x_o \\ x_t - x_o \end{bmatrix} + \omega_\nu^{(1)} \\
 \dot{x} &= \begin{bmatrix} \dot{x}_o \\ \dot{x}_t \end{bmatrix} = \begin{bmatrix} \frac{1}{t_s} (x_t - x_o) \\ \frac{1}{t_m} (\boldsymbol{\nu}^{(1)} - x_t) \end{bmatrix} + \omega_x^{(1)} \\
 \boldsymbol{\nu}^{(1)} &= \omega_x^{(2)}
 \end{aligned}$$

The first pair of equations corresponds to a noisy sensory mapping from hidden states and the equations of motion for states in the real world. These pertain to real-world variables representing the position of the target and of the eye (in boldface). The remaining equations constitute the generative model of how sensations are generated using the form of Equation 2. These define the free energy in Equation 1 – and specify behaviour under active inference. The variables constitute the first layer of the hierarchical model (see Equation 2, but for simplicity, we have written \mathbf{x} instead of $\mathbf{x}^{(1)}$ and x instead of $x^{(1)}$).

The real-world provides sensory input in two modalities: proprioceptive input from cranial nerve nuclei reports the angular displacement of the eye $\mathbf{s}_o \in \mathbb{R}^2$ and corresponds to the centre of gaze. Note that, using the approximation of relatively small displacements, we use Cartesian coordinates to follow previous treatments e.g (Friston et al., 2010a). However visual space is better described by bounded polar coordinates, and treatments of large eye movements should account for this. Exteroceptive (retinal) input reports the angular position of a target in a retinal (intrinsic) frame of reference $\mathbf{s}_t \in \mathbb{R}^2$. The indices o and t thus refer to states of the oculomotor system or of the target, respectively. Note that \mathbf{s}_t is just the difference between the centre of gaze and target location in an extrinsic frame of reference $\mathbf{x}_t - \mathbf{x}_o$. In this paper, we are modelling the online inference of target position, and we are ignoring the problem of how the causal structure of the environment is learned. We simply assume that this structure has already been learned accurately, and therefore the dynamics of the real-world and the generative model are the same. Clearly, this model of visual processing is an enormous simplification: we are assuming that place coded spatial information can be summarised in terms of displacement vectors.

However, more realistic simulations – using a set of retinotopic inputs with classical receptive fields covering visual space – produce virtually the same results. We will use more realistic models in future publications that deal with smooth pursuit and visual occlusion. Here, we use the simpler formulation to focus on delays and the different sorts of generative models that can provide top-down or extra-retinal constraints on visual motion processing.

The hidden states of this model comprise the true, real-world oculomotor displacement ($\mathbf{x}_o \in \mathbb{R}^2$) and target location ($\mathbf{x}_t \in \mathbb{R}^2$). The units of angular displacement are arbitrary, but parameters are tuned to correspond to a small displacement of 4 degrees of visual angle for one arbitrary unit (that is approximately 4 times the width of a thumb’s nail at arm’s length). The oculomotor state is driven by action with a time constant of $t_a = 64$ ms and decays (slowly) to zero through damping, with a time constant of $t_o = 512$ ms. The target location is perturbed by hidden causes $\mathbf{x}_t \in \mathbb{R}^2$ that describe the location to which the target is drawn, with a time constant of $t_m = 16$ ms. In this paper, the random fluctuations on sensory input and on the motion of hidden states are very small, with a log precision of 16. In other words, the random fluctuations have a variance of $\exp(-16)$. This completes our description of the process generating sensory information; in which hidden causes force the motion of a target location and action forces oculomotor states. Target location and oculomotor states are combined to produce sensory information about the target in an intrinsic frame of reference.

The generative model has exactly the same form as the generative process but with one important exception: there is no action and the motion of the hidden oculomotor states is driven by the displacement between the target location and the central gaze (with a time constant of $t_s = 32$ ms). In other words, the agent believes that its gaze will be attracted to the location of the target, which, itself, is being driven by some unknown exogenous force or hidden cause. The log-precisions on the random fluctuations in the generative model were four, unless stated otherwise. This means that uncertainty about sensory input, (motion of) hidden states and causes were roughly equivalent.

Having specified the generative process and model, we can now solve the active inference scheme in Equation 1 and examine its behaviour. Sensorimotor delays are implemented in the message passing from the generative process to the generative model. This generative model produces pursuit initiation because it embodies prior beliefs that the centre of gaze will follow the target location. This pursuit initiation rests on conditional expectations about the target location in extrinsic coordinates and the state of the oculomotor plant, where the location is driven by hidden causes that also have to be inferred.

The generative model described in this section provides the equations required to simulate active inference using the formalism of the previous section. In short, we now consider the generative model that defines the variational free energy and (Bayes) optimal active inference.

4.2 Simulations

All simulations were performed with a time bin of 16ms and we report results in milliseconds. All results were replicated with different time bins (16ms, 8ms, 4ms, 2ms, 1ms) with minimal changes to the results. Figure 4 reports the conditional estimates of hidden states and causes during the simulation of pursuit

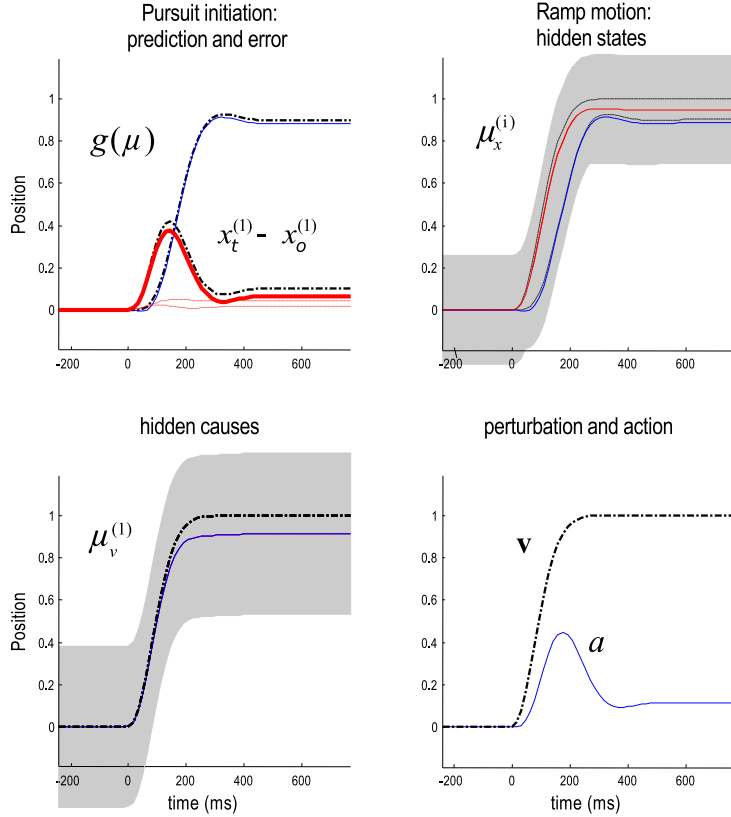


Figure 4: This figure reports the conditional estimates of hidden states and causes during the simulation of pursuit initiation, using a single rightward (positive) sweep of a visual target, while compensating for sensory motor delays. We will use the format of this figure in subsequent figures: the upper left panel shows the predicted sensory input (coloured lines) and sensory prediction errors (dotted red lines) along with the true values (broken black lines). Here, we see horizontal excursions of oculomotor angle (upper lines) and the angular position of the target in an intrinsic frame of reference (lower lines). This is effectively the distance of the target from the centre of gaze and reports the spatial lag of the target that is being followed (solid red line). One can see clearly the initial displacement of the target that is suppressed after a few hundred milliseconds. The sensory predictions are based upon the conditional expectations of hidden oculomotor (blue line) and target (red line) angular displacements shown on the upper right. The grey regions correspond to 90% Bayesian confidence intervals and the broken lines show the true values of these hidden states. One can see the motion that elicits following responses and the oculomotor excursion that follows with a short delay of about 64 ms. The hidden cause of these displacements is shown with its conditional expectation on the lower left. The true cause and action are shown on the lower right. The action (blue line) is responsible for oculomotor displacements and is driven by the proprioceptive prediction errors.

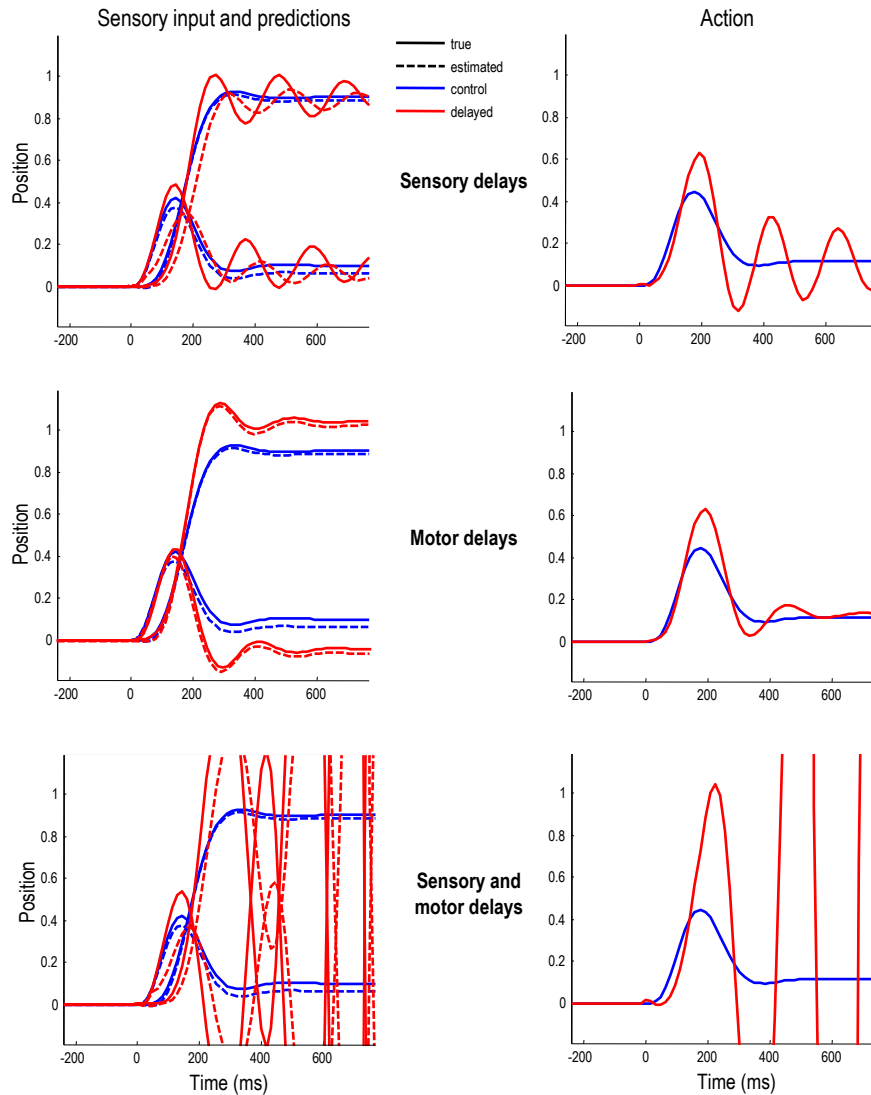


Figure 5: This figure illustrates the effects of sensorimotor delays on pursuit initiation (red lines) in relation to compensated (optimal) active inference – as shown in the previous figure (blue lines). The left panels show the true (solid lines) and estimated sensory input (dotted lines), while action is shown in the right panels. Under pure sensory delays (top row), one can see clearly the delay in sensory predictions, in relation to the true inputs. The thicker (solid and dotted) red lines correspond respectively to (true and predicted) proprioceptive input, reflecting oculomotor displacement. The middle row shows the equivalent results with pure motor delays and the lower row presents the results with combined sensorimotor delays. Of note here is the failure of optimal control with oscillatory fluctuations in oculomotor trajectories, which become unstable under combined sensorimotor delays.

initiation, using a simple rightward sweep of a visual target and compensating for sensorimotor delays: $\tau_s = \tau_s$ and $\tau_a = \tau_a$. This compensation is effectively the same as simulating responses in the absence of delays – because the delay operators reduce to the identity matrix. Target motion was induced using a hidden cause that was a ramp function of post-stimulus time. Note that ramp stimuli are often used in psychophysics, and this generative model – using velocity in place of position — produces the same results in velocity space. Indeed, most models, such as (Robinson, Gordon, and Gordon, 1986) or (Krauzlis and Lisberger, 1989), focus on modelling velocity responses. We choose to model the tracking of position for two reasons: First, it is easy to generalise position results to velocity using generalized coordinates of motion. Second, positional errors can induce slow eye movements (Kowler and Steinman, 1979; Wyatt and Pola, 1981) and we hoped to accommodate this in the model. If we assume that the units of angular displacement are 4 degrees of visual angle, then the resulting peak motion corresponds to about 20 degrees per second.

The upper left panel shows the predicted sensory input (coloured lines) and sensory prediction errors (dotted red lines) along with the true values (broken black lines). Here, we see horizontal excursions of oculomotor angle (upper lines) and the angular position of the target in an intrinsic frame of reference (lower lines). This is effectively the distance of the target from the centre of gaze and reports the spatial lag of the target that is being followed (solid red line). One can see clearly an initial retinal displacement of the target that is suppressed after approximately 20 ms. This effect confirms that the visual representation of target position is predictive and that the presentation of a smooth predictable versus an unpredictable target would induce a lag between their relative positional estimates, as is evidenced in the *flash-lag effect* (Nijhawan, 1994).

The sensory predictions are based upon the conditional expectations of hidden oculomotor (blue line) and target (red line) angular displacements shown on the upper right. The grey regions correspond to 90% Bayesian confidence intervals and the broken lines show the true values. One can see clearly the motion that elicits pursuit initiation responses, where the oculomotor excursion follows with a short delay of about 64 ms. The hidden cause of these displacements is shown with its conditional expectation on the lower left. The true cause and action are shown on the lower right. The action (blue line) is responsible for oculomotor displacements and is driven by proprioceptive prediction errors. Action does not return to zero because the sweep is maintained at an eccentric position during this simulation. This eye position slightly undershoots the target position: it is held at around 95% of the target eccentricity in the upper right panel. Note that this corresponds roughly to the steady-state gain observed in behavioural data, which was modelled explicitly by (Robinson, Gordon, and Gordon, 1986). For our purposes, these simulations can be regarded as Bayes optimal solutions to the pursuit initiation problem, in which sensorimotor delays have been accommodated (discounted) via absorption into the generative model. We can now examine the performance in the absence of compensation and see how sensory and motor delays interact to confound pursuit initiation:

The above simulations were repeated with uncompensated sensory delays ($\tau_s = 0$ ms and $\tau_s = 32$ ms), uncompensated motor delays ($\tau_a = 0$ ms and $\tau_a = 32$ ms) and combined sensorimotor delays of 64 ms ($\tau_a = \tau_s = 0$ ms and $\tau_a = \tau_s = 32$ ms). To quantify behaviour, we focus on the sensory input and un-

derlying action. The position of the target in intrinsic coordinates corresponds to spatial lag and usefully quantifies pursuit initiation performance. Figure 5 shows the results of these three simulations (red lines) in relation to the compensated (optimal) active inference shown in the previous figure (blue lines). True sensory input corresponds to solid lines and its conditional predictions to dotted lines. The left panels show the true and predicted sensory input, while action is shown in the right panels. Under pure sensory delays (top row) one can see the delay in sensory predictions, in relation to the true inputs. The thicker (solid and dotted) red lines correspond respectively to (true and predicted) proprioceptive input, reflecting oculomotor displacement. Crucially, in contrast to optimal control, there are oscillatory fluctuations in oculomotor displacement and the retinotopic location of the target that persists even after the target is stationary. These fluctuations are similar to the oscillations elicited by adding an artificial feed-back delay (Goldreich, Krauzlis, and Lisberger, 1992). Here, the fluctuations are caused by damped oscillations in action due to, and only to, sensory proprioceptive and exteroceptive delays. These become unstable (increasing in their amplitude) when the predicted value oscillates in counter-phase with the real value. Similar oscillations are observed with pure motor delays (middle row). However, here there is no temporal lag between the true and predicted sensations (solid versus dashed lines). Furthermore, there is no apparent delay in action – action appears to be emitted for longer, reaching higher amplitudes. In fact, action is delayed but the delay is obscured by the increase in the amplitude of action – that is induced by greater proprioceptive prediction errors. If we now combine both sensory and motor delays, we see a catastrophic failure of oculomotor tracking (lower row). With combined sensorimotor delays the pursuit initiation becomes unstable, with exponentially increasing oscillations as action over-compensates for delay-dependent errors.

In effect, the active inference scheme has undergone a phase transition from a stable to an unstable fixed point. We have illustrated this bifurcation by increasing sensorimotor delays under a fixed motor precision or gain in Equation 7. The results in Figure 5 used a motor gain with a log precision of 2.5. We chose this value because it produced stable responses with sensory or motor delays alone and unstable dynamics with combined delays. These results illustrate the profound and deleterious effects of sensorimotor delays on simple pursuit initiation, using biologically plausible values – namely sensorimotor delays of 64 ms and a target velocity of about 16 degrees per second. This also illustrates the necessity of compensation for these delays so that the system can achieve a more robust and stable response. One would anticipate, in the face of such failures, real subjects would engage interceptive saccades to catch the target, of the sort seen in schizophrenic patients (Levy et al., 1993). In the remainder of this paper, we will concentrate on the nature of pursuit initiation and smooth pursuit with compensated sensorimotor delays, using a reasonably high motor gain with a log precision of four.

4.3 Pursuit initiation and visual contrast

Before turning to more realistic generative models of smooth pursuit, we consider the empirical phenomena in which following responses to the onset of target movement are suppressed by reducing the visual contrast of the target (Thompson, 1982). In simulations of this sort, visual contrast is modelled

in terms of the precision of sensory information in accord with Weber’s law – see (Feldman and Friston, 2010) for details. Contrast-dependent effects are easy to demonstrate in the context of active inference. Figure 6 shows the spatial lag – the displacement in intrinsic coordinates of the target from the centre of gaze depicted by the solid red line in Figure 4 – as a function of contrast or log-precision of exteroceptive sensory input. The upper panel shows the true (solid lines) and predicted (dotted lines) spatial lag as a function of peristimulus time for different log precisions, ranging from two (low) to eight (high). The peak lags are plotted in the lower panel as a function of visual contrast or log precision. Since estimation error decreases as visual contrast increases, both curves converge, leading to a decrease to zero of the prediction error. These results show, in accord with empirical observations, how the spatial lag (position error) increases with contrast (Arnold, Ong, and Roseboom, 2009), while the true lag decreases (Spering et al., 2005). A similar difference between perception and action was recently reported (Simoncini et al., 2012). The explanation for this contrast-dependent behaviour is straightforward – because pursuit initiation is based upon proprioceptive prediction errors, it depends upon precise sensory information. Reducing the precision of visual input – through reducing contrast – increases uncertainty about visual information (sensory estimation error) and places more weight on prior beliefs and proprioceptive sensations. This reduces the perceived motion of the target and reduces the amplitude of prediction errors driving action.

4.4 Summary

In this section, we have seen that sensorimotor delays can have profound and deleterious effects on optimal oculomotor control. Here, optimal control means Bayes optimal active inference, in which pursuit initiation emerges spontaneously from prior beliefs about how a target attracts the centre of gaze. These simulations demonstrate that it is relatively easy to compensate for sensorimotor delays by exploiting representations in generalized coordinates of motion. Furthermore, the resulting scheme has some construct validity in relation to experimental manipulations of the precision or contrast of visual information. However, there are certain aspects of oculomotor tracking that suggest the pursuit initiation model above is incomplete: when presented with periodic target motion, the latency of motor gain (defined operationally in terms of the target and oculomotor velocities) characteristically reduces after the first cycle of target motion (Barnes, Barnes, and Chakraborti, 2000). This phenomenon cannot be reproduced by the pursuit initiation model above:

Figure 7 shows the responses of the pursuit initiation model to sinusoidal motion using the same format as Figure 4. Here, the hidden cause driving the target was a sine wave with a period of 512 ms that started after 256 ms. If we focus on the spatial lag (solid red line in the upper left panel), one can see that the lag is actually greater after one period of motion than at the onset of motion. This contrasts with empirical observations, which suggest that the spatial lag should be smaller after the first cycle (Barnes, Barnes, and Chakraborti, 2000). In the next section, we consider a more realistic generative model that resolves this discrepancy and takes us from simple pursuit initiation to smooth pursuit.

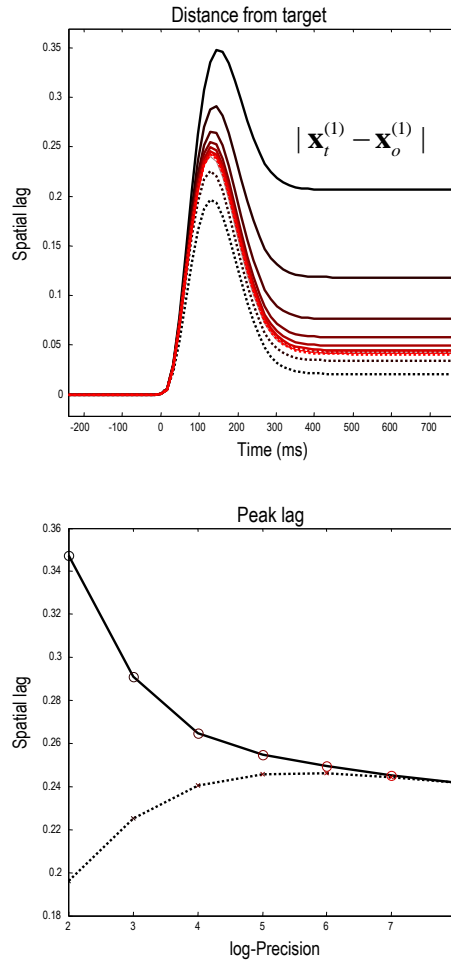


Figure 6: This figure reports the spatial lag (the displacement of the target from the centre of gaze) as a function of contrast (log precision of exteroceptive sensory input). The upper panel shows the true (solid lines) and predicted (dotted lines) spatial lag as a function of peristimulus time for different log precisions, ranging from two (black lines) to eight (red lines). The peak lags are plotted in the lower panel as a function of visual contrast or log precision. These results show how the perceived lag increases with contrast, while the true lag decreases in accord with empirical observations.

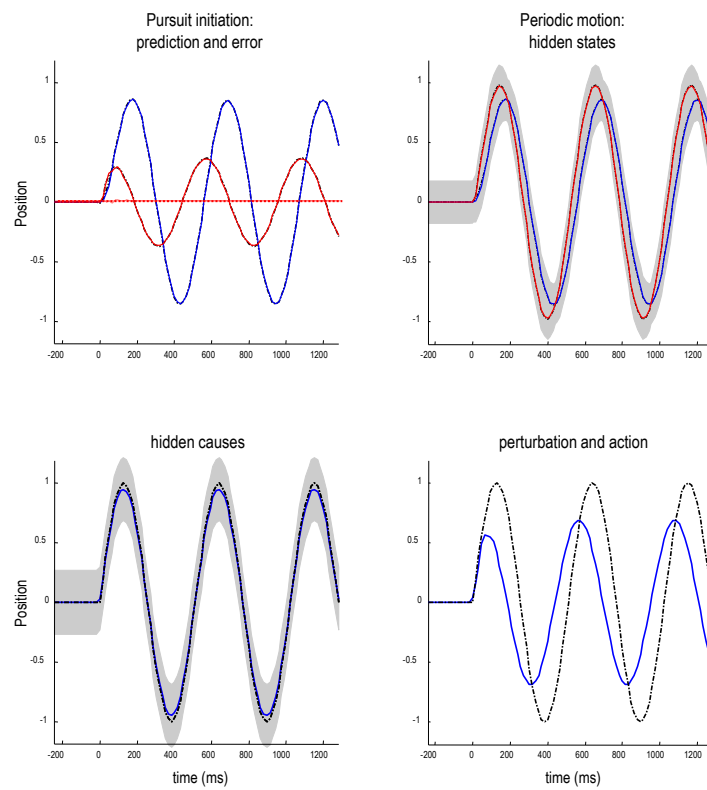


Figure 7: This figure uses the same format as Figure 4 – the only difference here is that the target motion is sinusoidal. The key thing to take from this simulation is that the peak spatial lag at the onset of the second cycle of target motion is greater than the peak lag at the onset of the first. This is contrary to empirical predictions.

5 Results: smooth pursuit

In this section, we consider a slightly more realistic generative model that replaces the prior beliefs about the target attracting the centre of gaze with the belief that both the target and centre of gaze are attracted by the same (fictive) location in visual space. This allows pursuit initiation to anticipate the trajectory of the target and pursue the target more accurately – providing the trajectories are sufficiently smooth. The idea behind this generative model is to account for the improvements in tracking performance that are not possible at the onset of motion and that are due to inference on smooth target trajectories.

5.1 Smooth pursuit model

The smooth pursuit model considered in this paper rests on a second-order generalisation of the pursuit initiation model of previous section. Previously, we have considered the motion of the oculomotor plant to be driven directly by action. This form of action can be considered as an (adiabatic) solution to a proper second-order formulation, in which action exerts a force and thereby changes the angular acceleration of oculomotor displacement. This second-order formulation can be expressed in terms of the following generative process and model

$$\begin{aligned}
 \mathbf{s} &= \begin{bmatrix} \mathbf{s}_o \\ \mathbf{s}_t \end{bmatrix} = \begin{bmatrix} \mathbf{x}_o \\ \mathbf{x}_t - \mathbf{x}_o \end{bmatrix} + \boldsymbol{\omega}_\nu^{(1)} \\
 \dot{\mathbf{x}} &= \begin{bmatrix} \dot{\mathbf{x}}_o \\ \dot{\mathbf{x}}'_o \\ \dot{\mathbf{x}}_t \end{bmatrix} = \begin{bmatrix} \mathbf{x}'_t \\ \frac{1}{t_a}a - \frac{1}{t_o}\mathbf{x}'_o \\ \frac{1}{t_m}(\boldsymbol{\nu}^{(1)} - \mathbf{x}_t) \end{bmatrix} + \boldsymbol{\omega}_x^{(1)} \\
 & \tag{9} \\
 s &= \begin{bmatrix} s_o \\ s_t \end{bmatrix} = \begin{bmatrix} x_o \\ x_t - x_o \end{bmatrix} + \omega_\nu^{(1)} \\
 \dot{x} &= \begin{bmatrix} \dot{x}_o \\ \dot{x}'_o \\ \dot{x}_t \end{bmatrix} = \begin{bmatrix} x'_t \\ \frac{1}{t_v}(\nu^{(1)} - x_o) - \frac{t_s}{t_v}x'_o \\ \frac{1}{t_m}(\nu^{(1)} - x_t) \end{bmatrix} + \omega_x^{(1)} \\
 \nu^{(1)} &= \omega_\nu^{(2)}
 \end{aligned}$$

Here, the only thing that has changed is that we have introduced new hidden states corresponding to oculomotor velocity $\mathbf{x}'_o \in \mathbb{R}^2$. Action now changes the motion of the velocity (i.e., acceleration), as opposed to the velocity directly. This difference is reflected in the generative model but with one crucial addition – the hidden oculomotor state is not driven by the displacement between the *target* and the centre of gaze but by the displacement between the *hidden cause* and the centre of gaze. In other words, the hidden oculomotor states are attracted by the hidden cause of target motion – not the target motion *per se*. The idea here is that inference about the trajectory of the hidden cause should enable an anticipatory optimisation of pursuit initiation, provided these trajectories are smooth – hence a smooth pursuit model. Note that the equation of motion in the oculomotor model $\dot{x}_o = \frac{1}{t_s}(x_t - x_o)$ (see Equation 8) is the (adiabatic) solution to the equation used to model smooth pursuit: $\frac{1}{t_v}(\nu^{(1)} - x_o) - \frac{t_s}{t_v}x'_o = 0$

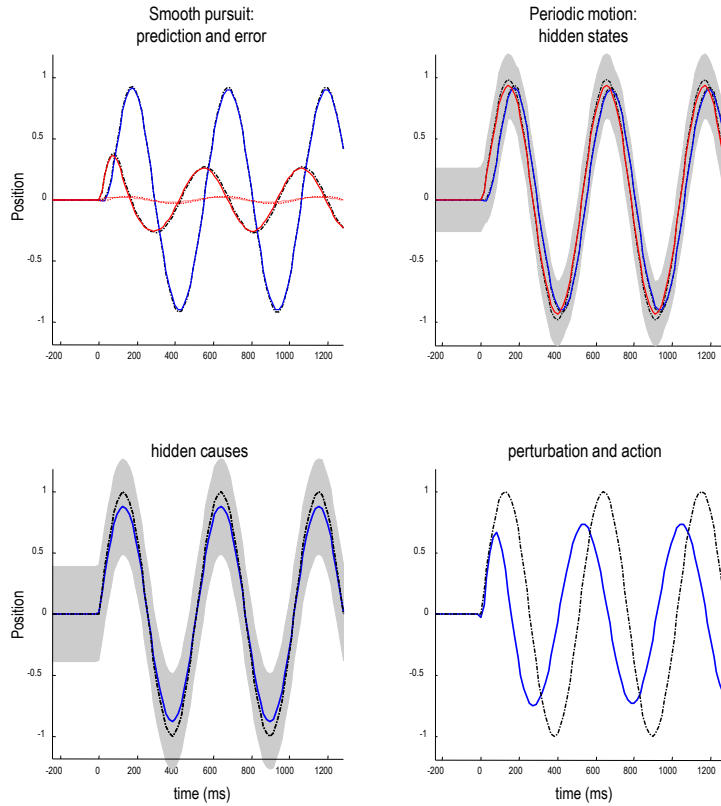


Figure 8: This figure uses the same format as the previous figure – the only difference here is that we have replaced the pursuit initiation model with a smooth pursuit model. In the smooth pursuit model the centre of gaze is attracted by a hidden cause of target motion, as opposed to the target *per se*. Note that, in comparison to the previous figure, the peak lag at the onset of the second cycle of target motion is now smaller than at the onset to the first.

when $\nu^{(1)} = x_t$ (see Equation 9). As a result (and as confirmed by simulations) this model behaved similarly for the sweep stimulus used in Figures 4, 5, 6.

5.2 Simulations

We repeated the simulation reported in Figure 7 using the smooth pursuit generative model. The results of this simulation are shown in Figure 8 using the same format as Figure 7. The key difference – in terms of performance – is that the peak spatial lag after one cycle of motion is now less than the peak lag at the onset of motion. The response to the sinusoid trajectory contrasts with simple pursuit initiation and is more consistent with empirical observations. The true and expected hidden states show that the oculomotor trajectory now follows the target trajectory more accurately, particularly at the peaks of right and leftward displacement. Interestingly, the amplitude of action has not changed very much

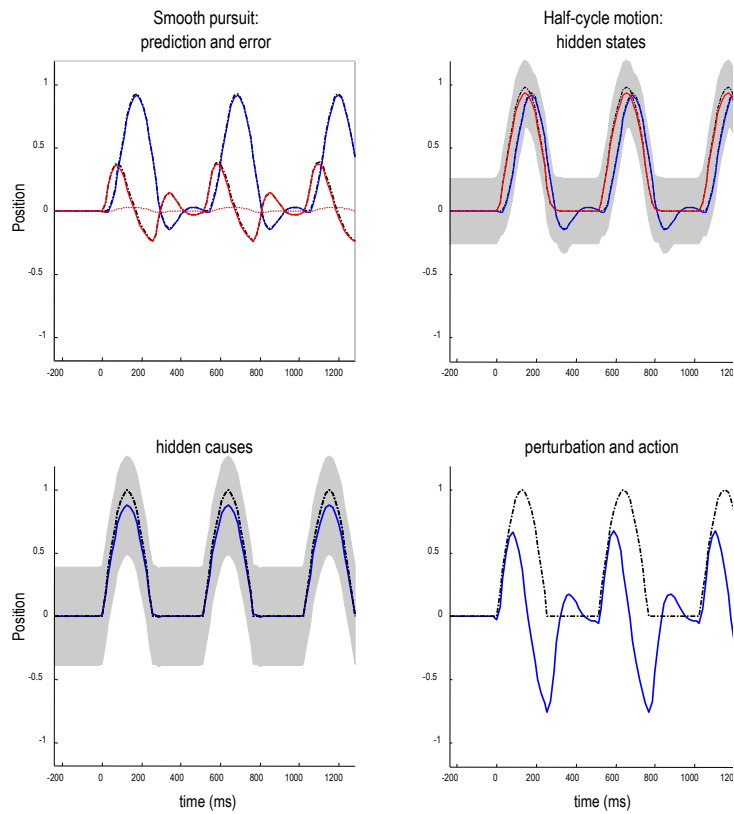


Figure 9: This figure uses the same format as the previous figure – the only difference is that the target motion has been rectified so that it is (approximately) hemi-sinusoidal. The thing to note here is that the improved accuracy of the pursuit previously apparent at the onset of the second cycle of motion has now disappeared – because active inference does not have access to the immediately preceding trajectory. This failure of an anticipatory improvement in tracking is contrary to empirical predictions.

(compare Figures 7 and 8, upper right panels). However, action is initiated with a slightly shorter latency, which is sufficient to account for the improved pursuit when informed by the prior beliefs about the smooth trajectory of the target.

5.3 Summary

In summary, by simply replacing the target with the hidden cause of target motion – as the attractor of oculomotor trajectories – we can account for empirical observations of improved pursuit during periodic target motion. In the context of active inference, this smooth trajectory can only be recognised – and used to inform action – after the onset of periodic motion. However, this smooth pursuit model still fails to account for anticipatory effects that are not directly available in sensory trajectories. Empirical observations suggest that any systematic or regular structure in target motion can facilitate the accuracy of smooth pursuit, even if this information is not represented explicitly in target motion. A nice example of this rests on the use of rectified periodic motion, in which only rightward target excursions are presented. Experimentally, subjects can anticipate the periodic but abrupt onset of motion, provided they recognise the underlying periodic behaviour of the target. We can emulate this hemi-periodic motion by thresholding the hidden cause to suppress leftward deflections. Figure 9 shows the results of simulating smooth pursuit using the same format as Figure 8. The only difference here is that we replaced the sinusoidal hidden cause $\nu(t) = \sin(2\pi f \cdot t)$ with $\nu(t) = \exp(4(\sin(2\pi f \cdot t) - 1))$. This essentially suppresses motion before rightward motion. This suppression completely removes the benefit of smooth pursuit after a cycle of motion – compare Figures 8 and 9. Here, the peak spatial lag at the onset of the second cycle of motion is exactly the same as the lag at the onset of motion; in other words, there is no apparent benefit of modelling the hidden causes of motion in terms of pursuit accuracy. This failure to model the anticipatory eye movements seen experimentally leads us to consider a full hierarchical model for anticipatory pursuit.

6 Results: anticipatory pursuit

This section presents a full hierarchical model of anticipatory smooth pursuit eye movements that tries to account for anticipatory oculomotor responses that are driven by extra-retinal beliefs about the periodic behaviour of targets. This entails adding a hierarchical level to the model that enables the agent to recognise and remember the latent structure in target trajectories and suitably optimise its pursuit movements – which are illustrated here in terms of an improvement in the accuracy of target following after the onset of rectified target motion.

6.1 Anticipatory pursuit

The generative process used in these simulations is exactly the same as in the above (smooth pursuit) scheme (see Equation 9); however, the generative model of this process is equipped with an extra level in place of the model for the hidden cause of target motion in the generative model:

$$\begin{aligned}
s &= \begin{bmatrix} s_o \\ s_t \end{bmatrix} = \begin{bmatrix} x_o \\ x_t - x_o \end{bmatrix} + \omega_\nu^{(1)} \\
\dot{x} &= \begin{bmatrix} \dot{x}_o \\ \dot{x}'_o \\ \dot{x}_t \end{bmatrix} = \begin{bmatrix} x'_t \\ \frac{1}{t_v}(\nu^{(1)} - x_o) - \frac{t_s}{t_v}x'_o \\ \frac{1}{t_m}(\nu^{(1)} - x_t) \end{bmatrix} + \omega_x^{(1)}
\end{aligned} \tag{10}$$

$$\begin{aligned}
\nu^{(1)} &= \begin{bmatrix} \sigma(x_1^{(2)}) \\ 0 \end{bmatrix} + \omega_\nu^{(2)} \\
\dot{x}^{(2)} &= \begin{bmatrix} \dot{x}_1^{(2)} \\ \dot{x}_2^{(2)} \end{bmatrix} = \nu^{(2)} \begin{bmatrix} x_2^{(2)} \\ -x_1^{(2)} \end{bmatrix} + \omega_x^{(2)} \\
\nu^{(2)} &= \eta + \omega_\nu^{(3)}
\end{aligned}$$

The first level of the generative model is exactly the same as above. However, the hidden causes are now informed by the dynamics of hidden states at the second level. These hidden states model underlying periodic dynamics using a simple periodic attractor that produces sinusoidal fluctuations of any amplitude or phase and a frequency that is determined by a second level hidden cause with a prior expectation of a frequency of η (in Hz). It is somewhat similar to a control system model that attempted to achieve zero-latency target tracking by fitting the trajectory to a (known) periodic signal (Bahill and McDonald, 1983). Our formulation ensures a Bayes optimal estimate of periodic motion in terms posterior beliefs about its frequency. In these simulations, we used a fixed Gaussian prior centred on the correct frequency with a period of 512 ms. This prior reproduces a typical experimental setting in which the oscillatory nature of the trajectory is known, but its amplitude and phase (onset) are unknown. Indeed, it has been shown that anticipatory responses are cofounded when randomising the inter-cycle interval (Becker and Fuchs, 1985). In principle, we could have considered many other forms of generative model, such as models with prior beliefs about continuous acceleration (Bennett et al., 2010).

As above, all the random fluctuations were assumed to have a log precision of four. Crucially, the mapping between the second level (latent) hidden states and the motion of first level hidden states encoding trajectories in visual (extrinsic) space is nonlinear. This means that latent periodic motion can be distorted in any arbitrary way. Here, we use a soft thresholding function $\sigma(x) = \exp(4(x - 1))$ to suppress negative (rightward) excursions of the target to model hemi-sinusoidal motion. This is the same function we used to generate the motion in Figure 9. Note that if the precision of the noise at the second level falls to zero and there is no (precise) information at this level, the generative model assumes that the random fluctuations have an infinite variance. As a consequence, the prediction at the level below in the hierarchical model simplifies to $\nu^{(1)} = \omega_\nu^{(2)}$, and we recover Equation 9 describing the smooth pursuit model. As a consequence this parameter tunes the relative strength of anticipatory modulation.

Figure 10 shows the results of simulating active inference under this anticipatory model, using the same format as Figure 9. However, there is now an extra level of hidden states encoding latent periodic motion. It can be seen

that expectations about hidden states attain nonzero amplitudes shortly after motion onset and are periodic thereafter. These provide predictions about the onset of rightward motion after the first (latent) cycle, enabling a more accurate oculomotor response. This is evidenced by the reduction in the spatial lag at the onset of the second cycle of motion, relative to the first (solid red lines on the upper left). This improvement in accuracy should be compared to the previous figure and reflects Bayes optimal anticipatory responses of the sort observed empirically (Barnes, Barnes, and Chakraborti, 2000). Further evidence of anticipatory inference can be seen by examining the conditional expectations about hidden causes at the second level. Note the substantial reduction in prediction error on the hidden cause (dotted red lines), when comparing the onset of the second cycle to the onset of the first. This reflects the fact that the conditional expectations about the hidden cause show a much-reduced latency at the onset of the second cycle due to top-down conditional predictions provided by the second level hidden states. This recurrent and hierarchically informed inference provides the basis for anticipatory oculomotor control and may be a useful metaphor for the hierarchical anatomy of the visual-oculomotor system.

6.2 Summary

In conclusion, to account for anticipatory pursuit movements that are not immediately available in target motion, one needs to equip generative models with a hierarchal structure that can accommodate latent dynamics – that may or may not be expressed at the sensory level. It is important to note that this model is a gross simplification of the complicated hierarchies that may exist in the brain. For instance, while some anticipation may be induced in smooth pursuit eye movements, some aspects, such as the aperture problem, may not be anticipated (Montagnini, Spering, and Masson, 2006). In this model, the second level hidden causes are simply driven by prediction errors and assume a constant frequency. As a consequence, prior beliefs about frequency are modelled as stationary. In the real brain one might imagine that models of increasing hierarchical depth might allow for nonstationary frequencies and other dynamics – that would better fit behavioural data. We have chosen to illustrate the basic ideas using a minimalistic example of anticipation in eye movements. Hierarchical extensions of this sort emphasise the distinction between visual motion processing and attending oculomotor control based purely upon retinal and proprioceptive input – they emphasise extra-retinal processing that is informed by prior experience and beliefs about the latent causes of visual input. We will exploit this anticipatory smooth pursuit model in future work, where visual occluders are used to disclose beliefs about latent motion.

7 Discussion

In this paper, we have considered optimal motor control in the context of pursuit initiation and anticipatory smooth pursuit. In particular, we have taken a Bayesian perspective on optimality and have simulated various aspects of eye movement control using predictive coding and active inference. This provides a solution to the problem of sensorimotor delays that reproduces the results of earlier solutions — but using a neuronally plausible (predictive coding) scheme

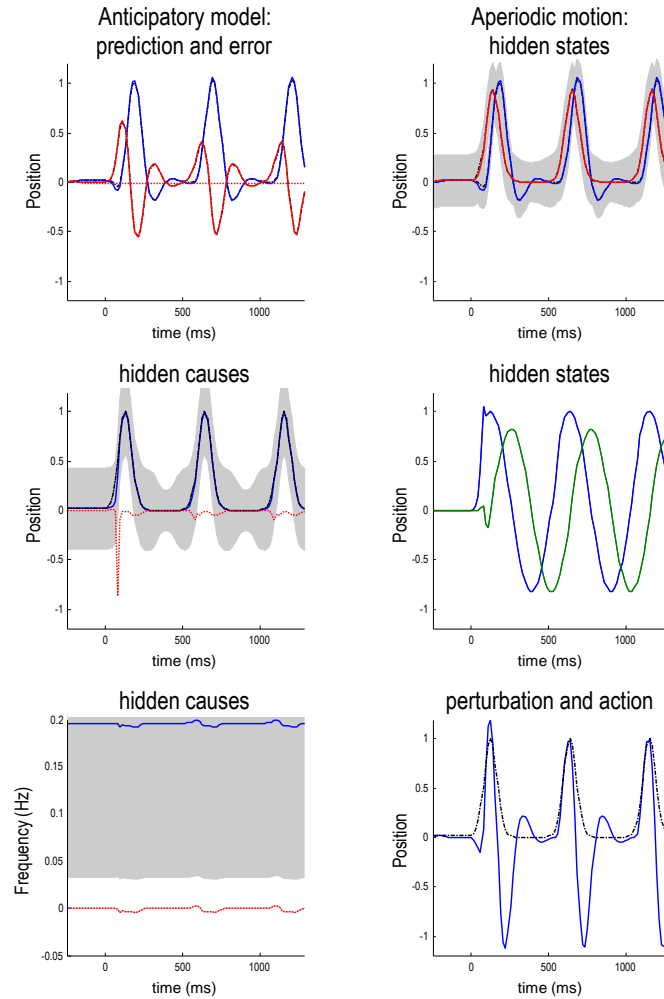


Figure 10: This figure uses the same format as the previous figure – the only difference is that the generative model has been equipped with a second hierarchical level that contains hidden states, modelling latent periodic behaviour of the (hidden) causes of target motion. With this addition, the improvement in pursuit accuracy apparent at the onset of the second cycle of motion is reinstated. This is because the model has an internal representation of latent causes of target motion that can be called upon even when these causes are not expressed explicitly in the target trajectory.

that has been applied to a whole range of perceptual, psychophysical, decision theoretic and motor control problems beyond oculomotor control. Active inference depends upon a generative model of stimulus trajectories and their active sampling through movement. This requires a careful consideration of the generative models that might be embodied by the visual-oculomotor system – and the sorts of behaviours one would expect to see under these models. The treatment in this paper distinguishes between three levels of predictive coding with respect to oculomotor control: the first is at the lowest level of sensorimotor message passing between the sensorium and internal states representing the causes of sensory signals. Here, we examined the potentially catastrophic effects of sensorimotor delays and how they can easily render oculomotor tracking inherently unstable. This problem can be finessed – in a relatively straightforward way – by exploiting representations in generalized coordinates of motion. These can be used to offset both sensory and motor delays, using simple and neurobiologically plausible mixtures of generalized motion. We then motivated a model of smooth pursuit eye movements by noting that a simple model of target following cannot account for the improvement in visual tracking after the onset of smooth and continuous target trajectories. In this paper, smooth pursuit was modelled in terms of hidden causes that attracted both the target and centre of gaze simultaneously – enabling the trajectory of the target to inform estimates of the hidden cause that, in turn, provide predictions about oculomotor consequences. While this extension accounted for experimentally observed tracking improvements – under continuous trajectories – it does not account for anticipatory movements that have to accumulate information over time. This anticipatory behaviour could only be explained with a deeper hierarchical model that has an explicit representation of latent (periodic) structure causing target motion. When the generative model was equipped with a deeper structure, it was then able to produce anticipatory movements of the sort seen experimentally. Clearly, the simulations in this paper are just heuristic and do not represent a proper simulation of neurobiological processing. However, they can be taken as proof of principle that the basic computational architecture – in terms of generalized representations and hierarchical models – can explain some important and empirical facts about eye movements. In what follows, we consider the models in this paper in relation to other models and how modelling of this sort may have important implications for understanding the visual-oculomotor system.

7.1 Comparison with other models

The model that we have presented here speaks to and complements several existing models of the oculomotor system. First, it shares some properties with computer vision algorithms used for image stabilisation. Such models often use motion detection coupled with salient feature detection for the registration of successive frames (Lucas and Kanade, 1981). A major difference is that these models are often applied to very specific problems or configurations for which they give an efficient, yet *ad hoc* solution. A more generic approach is to use – as our model does – a probabilistic method, for instance particle filtering (Isard and Blake, 1998). Our model provides a constructive extension – as we integrate the dynamics of both sensation and action. In principle, this could improve the on-line response of feature tracking algorithms.

Second, using our modelling approach, we reproduce similar behaviours

shown by other neuromimetic models of the oculomotor system. For example, the pursuit of a dot with known uncertainty can be modelled as the response of a Kalman filter (Kalman, 1960). Both generalized Bayesian (active inference) and Kalman filtering predict the current state of the system using prior knowledge (about previous target locations) and refine these predictions using sensory data (prediction errors). This analogy with block-diagrams from control theory was first highlighted by (Robinson, Gordon, and Gordon, 1986) and (Krauzlis and Lisberger, 1989) – and has since been used widely (Grossberg et al., 1997). For a recent treatment involving the neuromorphic modelling of cortical areas, see (Shibata et al., 2005). However, it should be noted that the link with Kalman filtering is rarely explicit (but see (Xivry et al., 2013)); most models have been derived heuristically, rather than as optimal solutions under a generative model. One class of such neuromimetic models uses neural networks that mimic the behaviour of the Kalman filter (Haykin, 2001). This model was used to fit and predict the response of smooth pursuit eye movements under different experimental parameters (Montagnini et al., 2007) or while interrupting information flow (Bogadhi et al., 2011). Developing this methodology – and by analogy with modular control theory architectures – these building blocks can be assembled to accommodate increasingly complex behavioural tasks. This can take the form of a multi-layered model for transparency processing (Raudies, Mingolla, and Neumann, 2011) or of an interconnected graph connecting the form and motion pathways (Beck, Ognibeni, and Neumann, 2008). Such models have been used to understand adaptation to blanking periods and to tune the balance between sensory and proprioceptive inputs (Madelain and Krauzlis, 2003). Our model is different in a key aspect: The Kalman filter is indeed the (Bayes) optimal solution under a linear generative model but a cascade of such solutions is not the optimal solution to (non-linear) hierarchical models (Balaji and Friston, 2011). The active inference approach considers the (embodied) system as a whole and furnishes an optimal solution in the form of generalized Bayesian filtering. In particular, given the delays at the sensory and motor levels, it provides an optimal solution that accommodates (or compensates for) these delays. As shown in the results, the ensuing behavior reproduces experimental results from pursuit initiation (Masson, Montagnini, and Ilg, 2010) to anticipatory responses (Avila et al., 2006; Barnes, Barnes, and Chakraborti, 2000). The approach thus provides an inclusive framework, compared with heuristics used in neuromimetic models that focus on specific aspects of oculomotor control (see below).

The model presented here shares many features with other probabilistic models. First, representations are encoded as probability density is. This allows processing and control to be defined in terms of probabilistic inference; for instance, by specifying a prior belief that favours slow speeds (Weiss, Simoncelli, and Adelson, 2002). This approach has been successful in explaining a wide variety of physiological and psychophysical results. For example, it allows one to model spatial (Perrinet and Masson, 2007) or temporal (Montagnini et al., 2007) integration of information, using conditional independence assumptions. Furthermore, recent developments have addressed the estimation of the shape and parameters of priors for slow speeds (Stocker and Simoncelli, 2006) and for the integration of ambiguous versus non-ambiguous information (Bogadhi, Montagnini, and Masson, 2011). The active inference scheme used here relies on generative models that entail exactly the same sorts of priors. It has also been shown that free energy minimisation extends the type of probabilistic

models described above to encompass retinal stabilisation and oculomotor reflexes (Friston et al., 2010b). A crucial difference here is that we have explicitly considered the problem of dynamics and delays. Our goal was to understand how the system could provide an optimal solution, when it knows (or can infer) the delay between sensing input (in the past) and processing information that informs action (in the future). This endeavour allowed us to build a model – using simple priors over the dynamics of the hidden causes – that reproduces the sorts of anticipatory behaviour seen empirically.

7.2 Limitations

Clearly, there are many aspects of oculomotor control we have ignored in this theoretical work. Foremost, we have used a limited set of stimuli to validate the model. Pursuit initiation was only simulated using a simple sweep of a dot, while smooth pursuit was studied using a sinusoidal trajectory. However, these types of stimuli are commonly used in the literature, as they best characterise the type of behaviour (following, pursuit) that we have tried to characterise: see (Barnes, 2008) for a review. We have not attempted to reproduce the oscillations at steady state as in (Robinson, Gordon, and Gordon, 1986) or (Goldreich, Krauzlis, and Lisberger, 1992), although this may help to optimise the parameters of our model in relation to empirical data. The hemi-sinusoidal stimulus is also a typical stimulus for studying anticipatory responses (Avila et al., 2006; Barnes, Barnes, and Chakraborti, 2000). Further validations of this model would call on a wider range of stimuli and consider and accumulated wealth of neurophysiological and behavioural data (Tlapale et al., 2010).

In this paper, we have focused on inference under a series of generative models of oculomotor control. We have not considered how these models are acquired or learned. In brief, the acquisition of generative models and their subsequent optimisation in terms of their parameters (i.e., synaptic connection strengths) is an important, if distinct, issue. In the context of active inference, model acquisition and perceptual learning can be cast in terms of model selection and parameter optimisation through the minimisation of free energy. Under certain simplifying assumptions, this learning reduces to associative plasticity. A discussion of these and related issues can be found in Friston (2008).

The generative model used in this paper has no explicit representation of space but only the uncertain, vectorial position of a target. We have previously studied the role of prediction in solving problems that are associated with the detection of motion using a dynamical and probabilistic model of spatial integration (Perrinet and Masson, 2012). Both that model and the current model entertain a similar problem: that of the integration of local information into a global percept, in both the temporal (this manuscript) and spatial (Perrinet and Masson, 2012) domains. We have considered integrating sensory information in the spatial domain: terms of the prediction of sensory causes and their sampling by saccades (Friston, Thornton, and Clark, 2012), and of the effects on smooth pursuit of reducing the precision. This manipulation can account for several abnormalities of smooth pursuit eye movements typical of schizophrenia (Adams, Perrinet, and Friston, 2012). In this paper, we have limited ourselves to integrating information over time. It would be nice, in the future, to consider temporal and spatial integration simultaneously.

A final limitation of our model is the simplified modelling of the physical

properties of the oculomotor system – due to the biophysics of the eyes and photoreceptors, sensory input contains motion streaks that can influence the detection of motion (Barlow and Olshausen, 2004). Furthermore, we have ignored delays in neuronal message passing among and within different levels of the hierarchy: for a review of quantitative data from monkeys see (Salin and Bullier, 1995). Finally, we have not considered in any depth the finer details of how predictive coding or Bayesian filtering might be implemented neuronally. It should be noted that predictive coding in the cortex was attended by some early controversies; for example, paradoxical increases in visual evoked responses were observed when prediction error should be minimal. For example, a match between sensory signals and descending predictions can lead to the enhancement of neuronal firing (Roelfsema, Lamme, and Spekreijse, 1998). The neuronal implementation assumed in our work (see 2) finesses many of these issues. In this (hypothetical) scheme, predictions and prediction errors are encoded by the neuronal activity of deep and superficial pyramidal cells respectively (Mumford 1992; Bastos et al. 2012). In this scheme, the enhancement of evoked responses is generally thought to reflect attentional gain, which corresponds to the optimization of the expected precision (inverse variance) of prediction errors, via synaptic gain control (Feldman and Friston, 2010). Put simply, attention increases the gain of salient or precise prediction errors that the predictions are trying to suppress. Indeed, the orthogonal effects of expectations and attention in predictive coding have been established empirically using fMRI (Kok et al., 2011). See Bastos et al. (2012) for a review of the anatomical and electrophysiological evidence that is consistent with the scheme used here.

7.3 Perspectives

Notwithstanding the limitations above, this approach may provide some interesting perspectives on neural computations in the oculomotor system. First, the model presented here can be compared to existing models of the oculomotor system. In particular, any commonalities of function suggest that extant neuromimetic models may be plausibly implemented using a generic predictive coding architecture. Second, the Bayes optimal control solution rests on a computational (anatomical) architecture that can be informed by electrophysiological or psychophysical studies. For example, we have considered only delays at the motor and sensory level. However, delays in axonal conduction between hierarchical levels – within the visual-oculomotor system – may have implications for intrinsic and extrinsic connectivity: in visual search, predictions generated in higher areas (say supplementary and frontal eye fields) may exploit a shorter path, by stimulating the actuator to sample more information (by making an eye movement) rather than accumulating evidence by explaining away prediction errors in lower (striate and extrastriate) cortical levels (Masson, Montagnini, and Ilg, 2010). By studying the structure of connections implied by theoretical considerations (see Figure 3), our modelling approach could provide a formal framework to test these sorts of hypotheses. A complementary approach would be to apply dynamic causal modelling (Friston, Harrison, and Penny, 2003) to electrophysiological data, using predictive coding architectures, such that transmission delays (and their compensation or modeling) among levels of the visual-oculomotor system could be evaluated empirically. A recent example of using dynamic causal modelling to test hypotheses based upon pre-

dictive coding architectures can be found in Brown and Friston (2013). This example focuses on attentional gain control in visual hierarchies.

Second, this work may provide a new perspective for experiments, in particular for the generation of stimuli. We have previously considered such a line of research by designing naturalistic, texture-like pseudo-random visual stimuli to characterise spatial integration during visual motion detection (Leon et al., 2012). We were able to show that the oculomotor system exhibits an increased following gain, when stimuli have a broad spatial frequency bandwidth. Interestingly, the velocities of these stimuli were harder to discriminate relative to narrow bandwidth stimuli – in a two alternative forced-choice psychophysical task (Simoncini et al., 2012). In this work, the authors used competitive dynamics based on divisive normalisation. Moreover, textured stimuli were based on a simple forward model of motion detection (Leon et al., 2012). This may call for the use of more complex generative models to generate such textures. In addition, the use of gaze contingent eye-tracking systems allows real-time manipulation of the configuration (position, velocity, delays) of the stimulus, with respect to eye position and motion. By targeting different sources of uncertainty, at the different levels of the hierarchical model, one might be able to get a better characterisation of the oculomotor system.

The confounding influence of delays inherent in neuronal processing is a strong biophysical constraint on neuronal dynamics. Representations in generalized coordinates of motion provide a potential resolution that may have enjoyed positive evolutionary pressure. However, it remains unclear how neural information, represented in a distributed manner across the nervous system, is integrated with exteroceptive, operational time. The “binding” of different information, without a central clock, seems essential, but the correlate of such a temporal representation of sensory information (independent of delays) has never been observed explicitly in the nervous system. Elucidating the neural representation of temporal information would greatly enhance our understanding of both neural computations themselves and our interpretation of measured electromagnetic (EEG and MEG) signals that are tightly coupled to those computations.

Acknowledgments

LuP was supported by EC IP project FP6-015879, “FACETS” and FP7-269921, “BrainScaleS” and by the Wellcome Trust Centre for Neuroimaging. RAA and KJF are supported by the Wellcome Trust Centre for Neuroimaging.

8 Appendix

8.1 Appendix 1: Variational free energy

Here, we derive various formations of free-energy and show they relate to each other. We start with the quantity we want to bound and implicitly minimise — namely, surprise or the negative log-evidence associated with sensory states $\tilde{s}(t)$ that have been caused by some unknown quantities $\Psi(t)$. These hidden causes correspond to the (generalized) motion (that is, position, velocity, acceleration, ...) of a target that the oculomotor system is tracking.

$$-\ln p(\tilde{s}) = -\ln \int p(\tilde{s}, \Psi) d\Psi \quad (11)$$

We now simply add a non-negative cross-entropy or divergence between some arbitrary (conditional) density $q(\Psi) = q(\Psi|\tilde{\mu})$ and the posterior density $p(\Psi|\tilde{s})$ to create a free-energy bound on surprise

$$\begin{aligned} F &= -\ln p(\tilde{s}) + \int q(\Psi) \ln \frac{q(\Psi)}{p(\Psi|\tilde{s})} d\Psi \\ &= -\ln p(\tilde{s}) + D(q(\Psi)||p(\Psi|\tilde{s})) \end{aligned} \quad (12)$$

The cross entropy term is non-negative by Gibb's inequality. Because surprise depends only on sensory states, we can bring it inside the integral and use $p(\tilde{s}, \Psi) = p(\Psi|\tilde{s})p(\tilde{s})$ to show free-energy is a Gibb's energy $G = -\ln p(\tilde{s}, \Psi)$ expected under the conditional density minus the entropy of the conditional density

$$\begin{aligned} F &= \int q(\Psi) \ln \frac{q(\Psi)}{p(\Psi|\tilde{s})p(\tilde{s})} d\Psi \\ &= \int q(\Psi) \ln \frac{q(\Psi)}{p(\Psi, \tilde{s})} d\Psi \\ &= -\int q(\Psi) \ln p(\Psi, \tilde{s}) d\Psi + \int q(\Psi) \ln q(\Psi) d\Psi \end{aligned} \quad (13)$$

This is a useful formulation because it can be evaluated in a relatively straightforward way given a probabilistic generative model $p(\tilde{s}, \Psi)$. A final rearrangement, using $p(\tilde{s}, \Psi) = p(\tilde{s}|\Psi)p(\Psi)$, shows free-energy is also complexity minus accuracy, where complexity is the divergence between the recognition density $q(\Psi)$ and the prior density $p(\Psi)$

$$\begin{aligned} F &= \int q(\Psi) \ln \frac{q(\Psi)}{p(\Psi|\tilde{s})p(\tilde{s})} d\Psi \\ &= -\int q(\Psi) \ln p(\tilde{s}|\Psi) d\Psi + D(q(\Psi)||p(\Psi)) \end{aligned} \quad (14)$$

8.2 Appendix 2: The maximum entropy principle and the Laplace assumption

If we admit an encoding of the conditional density up to second order moments, then the maximum entropy principle (Jaynes, 1957), implicit in the definition of free energy above, requires $q(\Psi|\tilde{\mu}) = \mathcal{N}(\tilde{\mu}, \Sigma)$ to be Gaussian. This is because a Gaussian density has the maximum entropy of all forms that can be specified with two moments. Assuming a Gaussian form is known as the Laplace assumption and enables us to express the entropy of the conditional density in terms of its first moment or expectation. This follows because we can minimise free energy with respect to the conditional covariance as follows:

$$\begin{aligned}
F &= G(\tilde{s}, \tilde{\mu}) + \frac{1}{2} \text{tr}(\Sigma \partial_{\tilde{\mu}} G) - \frac{1}{2} \ln |\Sigma| \\
G &= -\ln p(\tilde{s}, \Psi) \\
\partial_{\Sigma} F &= \frac{1}{2} \partial_{\tilde{\mu}} G - \frac{1}{2} \Pi
\end{aligned} \tag{15}$$

so that $\partial_{\Sigma} F = 0$ implies

$$\begin{aligned}
\Pi &= \partial_{\tilde{\mu}} G \\
F &= G(\tilde{s}, \tilde{\mu}) + \frac{1}{2} \ln |\partial_{\tilde{\mu}} G|
\end{aligned} \tag{16}$$

Here, the conditional precision $\Pi(\tilde{s}, \tilde{\mu})$ is the inverse of the conditional covariance $\Sigma(\tilde{s}, \tilde{\mu})$. In short, free energy is a function of generalized conditional expectations and sensory states.

8.3 Appendix 3: Integrating or solving active inference schemes using generalized descents.

Given a generative model or its associated Gibbs energy function, one can now simulate active inference by solving the following set of ordinary differential equations for a system that includes generalized real-world states and internal states of the agent mediating (delayed) action and perception:

$$\dot{\mathbf{u}} = \begin{bmatrix} \dot{\tilde{s}} \\ \dot{\tilde{x}} \\ \dot{\tilde{\nu}} \\ \dot{\tilde{\omega}}_{\nu} \\ \dot{\tilde{\omega}}_x \\ \dot{\tilde{\mu}} \\ \dot{\tilde{\eta}} \\ \dot{\tilde{\mathbf{a}}} \end{bmatrix} = \begin{bmatrix} \mathcal{D}\tilde{g}(\tilde{x}, \tilde{\nu}, \tilde{\mathbf{a}}) + \mathcal{D}\tilde{\omega}_{\nu} \\ \tilde{f}(\tilde{x}, \tilde{\nu}, \tilde{\mathbf{a}}) + \tilde{\omega}_x \\ \mathcal{D}\tilde{\nu} \\ \mathcal{D}\tilde{\omega}_{\nu} \\ \mathcal{D}\tilde{\omega}_x \\ \mathcal{D}\tilde{\mu} - \partial_{\tilde{\mu}} F(T(\tau_s - \tau_s)\tilde{s}, \tilde{\mu}) \\ \mathcal{D}\tilde{\eta} \\ -\partial_a F(T(\tau_s - \tau_s + \tau_a - \tau_a)\tilde{s}, T(\tau_a - \tau_a)\tilde{\mu}) \end{bmatrix} \tag{17}$$

generalized action $\tilde{\mathbf{a}}(t)$ is approximated using discrete values of $\mathbf{a}(t)$ from the past. Note that we have included a prior expectation $\tilde{\eta}(t)$ of hidden causes to complete the agent's generative model of its world. Integrating or solving Equation 17 corresponds to simulating active inference. The updates of the collective states over time steps of Δt use a local linearisation scheme (Ozaki, 1992):

$$\begin{aligned}
\Delta u &= (\exp(\Delta t \cdot \partial_u \dot{u}) - I)(\partial_u \dot{u})^{-1} \\
\partial_u \dot{u} &= \begin{bmatrix} 0 & \mathcal{D}\partial_{\tilde{x}}\tilde{\mathbf{g}} & \mathcal{D}\partial_{\tilde{\nu}}\tilde{\mathbf{g}} & \mathcal{D} & \dots & & \mathcal{D}\partial_a\tilde{\mathbf{g}} \\ & \partial_{\tilde{x}}\tilde{\mathbf{f}} & \partial_{\tilde{\nu}}\tilde{\mathbf{f}} & I & & & \partial_a\tilde{\mathbf{f}} \\ \vdots & & \mathcal{D} & & \vdots & \vdots & \\ & & & \mathcal{D} & & & \\ -\partial_{\tilde{\mu}s}F & \dots & \dots & \mathcal{D} & \dots & \dots & \\ -\partial_{a\tilde{s}}F & & & & -\mathcal{D}\partial_{\tilde{\mu}\tilde{\mu}}F & -\partial_{\tilde{\mu}\tilde{\eta}}F & -\partial_{\tilde{\mu}a}F \\ & & & & -\partial_{\tilde{\eta}\tilde{\mu}}F & \mathcal{D} & \\ & & & & -\partial_{a\tilde{\mu}}F & -\partial_{a\tilde{\eta}}F & -\partial_{aa}F \end{bmatrix} \quad (18)
\end{aligned}$$

Details about how to compute the gradients and curvatures pertaining to the conditional expectations can be found in (Friston et al., 2010a). These are generally cast in terms of prediction errors using straightforward linear algebra. Because action can only affect free-energy through the sensory states, its dynamics are prescribed by the following gradients and curvatures (ignoring higher-order terms):

$$\begin{aligned}
\partial_a F &= (\partial_a \tilde{\varepsilon}_{\nu}^{(1)}) \cdot \Pi_a^{(1)} T(\tau_a - \boldsymbol{\tau}_a) \tilde{\varepsilon}_{\nu}^{(1)} \\
\partial_{aa} F &= (\partial_a \tilde{\varepsilon}_{\nu}^{(1)}) \cdot \Pi_a^{(1)} T(\tau_a - \boldsymbol{\tau}_a) (\partial_a \tilde{\varepsilon}_{\nu}^{(1)})
\end{aligned} \quad (19)$$

$$\begin{aligned}
\partial_a \tilde{\varepsilon}_{\nu}^{(1)} &= T(\tau_s - \boldsymbol{\tau}_s) \partial_a \tilde{\mathbf{s}}(t) \\
\partial_a \tilde{\mathbf{s}} &= \partial_a \tilde{\mathbf{g}} + \partial_{\tilde{x}} \tilde{\mathbf{g}} \left(\sum_i \mathcal{D}^{-i} (\partial_{\tilde{x}} \tilde{\mathbf{f}})^{i-1} \right) \partial_a \tilde{\mathbf{f}}
\end{aligned}$$

The partial derivative of the sensory states with respect to action and is specified by the generative process. In biologically plausible instances of this scheme, this derivative would have to be computed on the basis of a mapping from action to sensory consequences. It is generally assumed that agents are equipped with $\partial_a \tilde{\mathbf{s}}$ epigenetically, because it has a simple form. For example, contracting a muscle fibre elicits a proprioceptive stretch signal in a one-to-one fashion. The precision matrix $\Pi_a^{(1)}$ in Equation 19 is specified such that only proprioceptive prediction errors with these simple forms have nonzero precision. This can be regarded as the motor gain in response to proprioceptive prediction errors. Equation 18 may look complicated but can be evaluated automatically using numerical derivatives for any given generative model. All the simulations in this paper used just one routine — `toolbox/DEM/spm_ADEM.m` — and summarised in the script `toolbox/DEM/ADEM_oculomotor_delays.m`. Both are available as part of the SPM software (<http://www.fil.ion.ucl.ac.uk/spm>).

References

Adams, Rick A., Laurent U. Perrinet, and Karl Friston (2012). “Smooth Pursuit and Visual Occlusion: Active Inference and Oculomotor Control in Schizophrenia”. In: *PLoS ONE* 7.10. Ed. by Xiang Y. Zhang, e47502+.

- Arnold, Derek H., Yolanda Ong, and Warrick Roseboom (2009). “Simple differential latencies modulate, but do not cause the flash-lag effect”. In: *Journal of Vision* 9.5.
- Avila, Matthew T. et al. (2006). “Role of anticipation in schizophrenia-related pursuit initiation deficits.” In: *Journal of neurophysiology* 95.2, pp. 593–601.
- Bahill, A. T. and J. D. McDonald (1983). “Model emulates human smooth pursuit system producing zero-latency target tracking.” In: *Biological cybernetics* 48.3, pp. 213–222.
- Balaji, Bhashyam and Karl Friston (2011). “Bayesian state estimation using generalized coordinates”. In: Orlando, Florida, USA, 80501Y–80501Y–12.
- Ballard, Dana H., Geoffrey E. Hinton, and Terrence J. Sejnowski (1983). “Parallel visual computation”. In: *Nature* 306.5938, pp. 21–26.
- Barlow, Horace B. and Bruno A. Olshausen (2004). “Convergent evidence for the visual analysis of optic flow through anisotropic attenuation of high spatial frequencies.” In: *Journal of vision* 4.6, pp. 415–426.
- Barnes, G. R. and P. T. Asselman (1991). “The mechanism of prediction in human smooth pursuit eye movements.” In: *The Journal of physiology* 439, pp. 439–461.
- Barnes, Graham R. (2008). “Cognitive processes involved in smooth pursuit eye movements.” In: *Brain and cognition* 68.3, pp. 309–326.
- Barnes, Graham R., D. M. Barnes, and S. R. Chakraborti (2000). “Ocular pursuit responses to repeated, single-cycle sinusoids reveal behavior compatible with predictive pursuit”. In: *Journal of Neurophysiology* 84.5, pp. 2340–2355.
- Barnes, Graham R. and A. M. Schmid (2002). “Sequence learning in human ocular smooth pursuit.” In: *Experimental Brain Research* 144.3, pp. 322–35.
- Bastos, Andre M. et al. (2012). “Canonical microcircuits for predictive coding”. In: *Neuron* 76.4, pp. 695–711.
- Beal, Matthew J. (2003). *Variational algorithms for approximate Bayesian inference*.
- Beck, Cornelia, Thilo Ognibeni, and Heiko Neumann (2008). “Object segmentation from motion discontinuities and temporal occlusions: A biologically inspired model”. In: *PLoS ONE* 3.11. Ed. by Ernest Greene, e3807+.
- Becker, W. and A. F. Fuchs (1985). “Prediction in the oculomotor system: smooth pursuit during transient disappearance of a visual target.” In: *Experimental brain research* 57.3, pp. 562–575.
- Bennett, S. J. et al. (2007). “Target acceleration can be extracted and represented within the predictive drive to ocular pursuit”. In: *Journal of Neurophysiology* 98.3, pp. 1405–1414.
- Bennett, Simon J. et al. (2010). “Oculomotor prediction of accelerative target motion during occlusion: long-term and short-term effects”. In: *Experimental Brain Research* 204.4, pp. 493–504.
- Bialek, William, Ilya Nemenman, and Naftali Tishby (2001). “Predictability, complexity, and learning”. In: *Neural Computation* 13.11, pp. 2409–2463.
- Bogadhi, Amarendra, Anna Montagnini, and Guillaume Masson (2011). “Interaction between retinal and extra retinal signals in dynamic motion integration for smooth pursuit”. In: *Journal of Vision* 11.11, p. 533.
- Bogadhi, Amarendra et al. (2011). “Pursuing motion illusions: a realistic oculomotor framework for Bayesian inference”. In: *Vision research* 51.8, pp. 867–880.

- Changizi, Mark A. (2001). “‘Perceiving the present’ as a framework for ecological explanations of the misperception of projected angle and angular size”. In: *Perception* 30.2, pp. 195–208.
- Changizi, Mark A. and David M. Widders (2002). “Latency correction explains the classical geometrical illusions”. In: *Perception* 31.10, pp. 1241–1262.
- Changizi, Mark A. et al. (2008). “Perceiving the Present and a Systematization of Illusions”. In: *Cognitive Science: A Multidisciplinary Journal* 32.3, pp. 459–503.
- Collins, C. J. S. and Graham R. Barnes (2009). “Predicting the unpredictable: weighted averaging of past stimulus timing facilitates ocular pursuit of randomly timed stimuli.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 29.42, pp. 13302–14.
- Dayan, P. et al. (1995). “The Helmholtz machine.” In: *Neural computation* 7.5, pp. 889–904.
- Dodge, R., R. C. Travis, and J. C. Fox (1930). “Optic nystagmus: III. Characteristics of the slow phase”. In: *Archives of Neurology* 24, pp. 21–34.
- Duhamel, J. R., C. L. Colby, and M. E. Goldberg (1992). “The updating of the representation of visual space in parietal cortex by intended eye movements”. In: *Science* 255.5040, pp. 90–92.
- Feldman, Anatol G. and Mindy F. Levin (1995). “The origin and use of positional frames of reference in motor control”. In: *Behavioral and Brain Sciences* 18.04, pp. 723–744.
- Feldman, Harriet and Karl J. Friston (2010). “Attention, uncertainty, and Free-Energy”. In: *Frontiers in Human Neuroscience* 4.
- Friston, K. J., L. Harrison, and W. Penny (2003). “Dynamic causal modelling.” In: *NeuroImage* 19.4, pp. 1273–1302.
- Friston, Karl (2008). “Hierarchical models in the brain”. In: *PLoS Computational Biology* 4.11. Ed. by Olaf Sporns, e1000211+.
- Friston, Karl (2009). “The free-energy principle: a rough guide to the brain?” In: *Trends in cognitive sciences* 13.7, pp. 293–301.
- Friston, Karl (2011). “What Is optimal about motor control?” In: *Neuron* 72.3, pp. 488–498.
- Friston, Karl and Stefan Kiebel (2009). “Cortical circuits for perceptual inference”. In: *Neural Networks* 22.8, pp. 1093–1104.
- Friston, Karl, Christopher Thornton, and Andy Clark (2012). “Free-Energy minimization and the Dark-Room problem”. In: *Frontiers in Psychology* 3.
- Friston, Karl et al. (2010a). “Generalised Filtering”. In: *Mathematical Problems in Engineering* 2010, pp. 1–35.
- Friston, Karl et al. (2012). “Perceptions as hypotheses: Saccades as experiments”. In: *Frontiers in Psychology* 3.
- Friston, Karl J., Jean Daunizeau, and Stefan J. Kiebel (2009). “Reinforcement learning or active inference?” In: *PLoS ONE* 4.7, e6421+.
- Friston, Karl J. et al. (2010b). “Action and behavior: a free-energy formulation”. In: *Biological Cybernetics* 102.3, pp. 227–260.
- Ginzburg, V. (1955). “On the theory of superconductivity”. In: *Il Nuovo Cimento (1955-1965)* 2.6, pp. 1234–1250.
- Goldreich, D., R. J. Krauzlis, and S. G. Lisberger (1992). “Effect of changing feedback delay on spontaneous oscillations in smooth pursuit eye movements of monkeys.” In: *Journal of neurophysiology* 67.3, pp. 625–638.

- Gregory, R. L. (1980). "Perceptions as hypotheses". In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 290.1038, pp. 181–197.
- Grossberg, S. et al. (1997). "A neural model of multimodal adaptive saccadic eye movement control by superior colliculus." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 17.24, pp. 9706–9725.
- Haken, Hermann (1983). *Synergetik*. Vol. 2. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Haykin, Simon (2001). *Kalman filtering and neural networks*. Ed. by Simon Haykin. New York, USA: John Wiley & Sons, Inc.
- Ilg, Uwe J. (1997). "Slow eye movements". In: *Progress in Neurobiology* 53.3, pp. 293–329.
- Inui, Koji and Ryusuke Kakigi (2006). "Temporal analysis of the flow from V1 to the extrastriate cortex in humans". In: *Journal of Neurophysiology* 96.2, pp. 775–784.
- Isard, Michael and Andrew Blake (1998). "Condensation: conditional density propagation for visual tracking". In: *International Journal of Computer Vision* 29.1, pp. 5–28.
- Jaynes, E. (1957). "Information Theory and Statistical Mechanics. II". In: *Physical Review* 108.2, pp. 171–190.
- Kalman, R. E. (1960). "A new approach to linear filtering and prediction problems". In: *Journal of Basic Engineering* 82.1, pp. 35–45.
- Knill, D. and A. Pouget (2004). "The Bayesian brain: the role of uncertainty in neural coding and computation." In: *Trends in neurosciences* 27.12, pp. 712–719.
- Kok, P. et al. (2011). "Attention Reverses the Effect of Prediction in Silencing Sensory Signals". In: *Cerebral Cortex* 22.9, pp. 2197–2206.
- Kowler, Eileen (1989). "Cognitive expectations, not habits, control anticipatory smooth oculomotor pursuit". In: *Vision Research* 29.9, pp. 1049–1057.
- Kowler, Eileen and Robert M. Steinman (1979). "The effect of expectations on slow oculomotor control - I. Periodic target steps". In: *Vision Research* 19.6, pp. 619–632.
- Krauzlis, R. J. and S. G. Lisberger (1989). "A control systems model of Smooth Pursuit Eye Movements with realistic emergent properties". In: *Neural Computation* 1.1, pp. 116–122.
- Krauzlis, Richard J. (2004). "Recasting the smooth pursuit eye movement system". In: *Journal of Neurophysiology* 91.2, pp. 591–603.
- Leon, P. S. et al. (2012). "Motion clouds: model-based stimulus synthesis of natural-like random textures for the study of motion perception". In: *Journal of Neurophysiology* 107.11, pp. 3217–3226.
- Levy, D. L. et al. (1993). "Eye tracking dysfunction and schizophrenia: a critical perspective." In: *Schizophrenia bulletin* 19.3, pp. 461–536.
- Lisberger, S. G., E. J. Morris, and L. Tychsen (1987). "Visual motion processing and sensory-motor integration for smooth pursuit eye movements." In: *Annual Review of Neuroscience* 10.1, pp. 97–129.
- Lisberger, S. G. and J. A. Movshon (1999). "Visual motion analysis for pursuit eye movements in area MT of macaque monkeys." In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 19.6, pp. 2224–2246.
- Lucas, B. D. and T. Kanade (1981). "An Iterative Image Registration Technique with an Application to Stereo Vision". In: *Proceedings of the Seventh Inter-*

- national Joint Conference on Artificial Intelligence, IJCAI 1981*. Vancouver, Canada, pp. 674–679.
- Madelain, Laurent and Richard J. Krauzlis (2003). “Effects of learning on smooth pursuit during transient disappearance of a visual target.” In: *Journal of neurophysiology* 90.2, pp. 972–982.
- Masson, Guillaume S. and Uwe J. Ilg, eds. (2010). *Dynamics of visual motion processing: neuronal, behavioral and computational approaches*. First. Berlin-Heidelberg: Springer.
- Masson, Guillaume S., Anna Montagnini, and Uwe J. Ilg (2010). “When the Brain Meets the Eye: Tracking Object Motion”. In: *Dynamics of Visual Motion Processing*. Ed. by Uwe J. Ilg and Guillaume S. Masson. Boston, MA: Springer US. Chap. 8, pp. 161–188.
- Michael, J. A. and G. M. Jones (1966). “Dependence of visual tracking capability upon stimulus predictability.” In: *Vision research* 6.12, pp. 707–716.
- Montagnini, Anna, Miriam Spering, and Guillaume S. Masson (2006). “Predicting 2D target velocity cannot help 2D motion integration for smooth pursuit initiation.” In: *Journal of neurophysiology* 96.6, pp. 3545–3550.
- Montagnini, Anna et al. (2007). “Bayesian modeling of dynamic motion integration”. In: *Journal of Physiology (Paris)* 101.1-3, pp. 64–77.
- Mumford, D. (1992). “On the computational architecture of the neocortex. II. The role of cortico-cortical loops.” In: *Biological cybernetics* 66.3, pp. 241–251.
- Nijhawan, Romi (1994). “Motion extrapolation in catching”. In: *Nature* 370.6487.
- Nijhawan, Romi (2008). “Visual prediction: Psychophysics and neurophysiology of compensation for time delays”. In: *Behavioral and Brain Sciences* 31.02, pp. 179–198.
- Nkam, Irene et al. (2010). “Impaired smooth pursuit in schizophrenia results from prediction impairment only.” In: *Biological psychiatry* 67.10, pp. 992–997.
- Olshausen, B. A. and D. J. Field (1996). “Emergence of simple-cell receptive field properties by learning a sparse code for natural images.” In: *Nature* 381.6583, pp. 607–609.
- Ozaki, T. (1992). “A bridge between nonlinear time series models and nonlinear stochastic dynamical systems: a local linearization approach”. In: *Statistica Sinica* 2.1, pp. 113–135.
- Perrinet, Laurent U. and Guillaume S. Masson (2007). “Modeling spatial integration in the ocular following response using a probabilistic framework”. In: *Journal of Physiology (Paris)* 101.1–3, pp. 46–55.
- Perrinet, Laurent U. and Guillaume S. Masson (2012). “Motion-Based Prediction Is Sufficient to Solve the Aperture Problem”. In: *Neural Computation* 24.10, pp. 2726–2750.
- Rao, R. P. and D. H. Ballard (1999). “Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects”. In: *Nat Neurosci* 2.1, pp. 79–87.
- Raudies, Florian, Ennio Mingolla, and Heiko Neumann (2011). “A Model of Motion Transparency Processing with Local Center-Surround Interactions and Feedback”. In: *Neural Computation* 23.11, pp. 2868–2914.
- Robinson, D. A. (1965). “The mechanics of human smooth pursuit eye movement.” In: *The Journal of Physiology* 180, pp. 569–591.

- Robinson, D. A., J. L. Gordon, and S. E. Gordon (1986). “A model of the smooth pursuit eye movement system”. In: *Biological Cybernetics* 55.1, pp. 43–57.
- Roelfsema, P. R., V. A. Lamme, and H. Spekreijse (1998). “Object-based attention in the primary visual cortex of the macaque monkey.” In: *Nature* 395.6700, pp. 376–381.
- Roelfsema, Pieter R. et al. (1997). “Visuomotor integration is associated with zero time-lag synchronization among cortical areas.” In: *Nature* 385.6612, pp. 157–161.
- Salin, P. A. and J. Bullier (1995). “Corticocortical connections in the visual system: structure and function.” In: *Physiological reviews* 75.1, pp. 107–154.
- Shibata, Tomohiro et al. (2005). “A model of smooth pursuit in primates based on learning the target dynamics”. In: *Neural Networks* 18.3, pp. 213–224.
- Simoncini, Claudio et al. (2012). “More is not always better: adaptive gain control explains dissociation between perception and action”. In: *Nat Neurosci* 15.11, pp. 1596–1603.
- Spering, Miriam et al. (2005). “Effects of contrast on smooth pursuit eye movements”. In: *Journal of Vision* 5.5.
- Stocker, Alan A. and Eero P. Simoncelli (2006). “Noise characteristics and prior expectations in human visual speed perception”. In: *Nature Neuroscience* 9.4, pp. 578–585.
- Thaker, G. K. et al. (1999). “Smooth pursuit eye movements to extra-retinal motion signals: deficits in patients with schizophrenia.” In: *Psychiatry research* 88.3, pp. 209–219.
- Thompson, P. (1982). “Perceived rate of movement depends on contrast.” In: *Vision research* 22.3, pp. 377–380.
- Tlapale, milien et al. (2010). *Towards a bio-inspired evaluation methodology for motion estimation models*. Tech. rep. RR-7317. INRIA, p. 18.
- Vaughn, Don A. and David M. Eagleman (2013). “Spatial warping by oriented line detectors can counteract neural delays”. In: *Frontiers in Psychology* 4.
- Weiss, Yair, Eero P. Simoncelli, and Edward H. Adelson (2002). “Motion illusions as optimal percepts”. In: *Nature Neuroscience* 5.6, pp. 598–604.
- Westheimer, G. (1954). “Eye movement responses to a horizontally moving visual stimulus”. In: *Archives of Ophthalmology* 52, pp. 932–43.
- Wyatt, H. J. and J. Pola (1981). “Slow eye movements to eccentric targets.” In: *Investigative ophthalmology & visual science* 21.3, pp. 477–483.
- Wyatt, H. J. and J. Pola (1987). “Smooth eye movements with step-ramp stimuli: the influence of attention and stimulus extent.” In: *Vision research* 27.9, pp. 1565–1580.
- Xivry, Jean-Jacques O. de et al. (2013). “Kalman Filtering Naturally Accounts for Visually Guided and Predictive Smooth Pursuit Dynamics”. In: *The Journal of Neuroscience* 33.44, pp. 17301–17313.
- Yuille, Alan and Daniel Kersten (2006). “Vision as Bayesian inference: analysis by synthesis?” In: *Trends in Cognitive Sciences* 10.7, pp. 301–308.