

Chapter: Active inference, novelty, and neglect

Thomas Parr¹, Karl J Friston¹

¹ Wellcome Centre for Human Neuroimaging, Institute of Neurology, University College London, WC1N 3BG, UK.

thomas.parr.12@ucl.ac.uk, k.friston@ucl.ac.uk,

Correspondence: Thomas Parr

The Wellcome Centre for Human Neuroimaging

Institute of Neurology

12 Queen Square, London, UK WC1N 3BG

thomas.parr.12@ucl.ac.uk

Abstract

In this chapter, we provide an overview of the principles of active inference. We illustrate how different forms of short-term memory are expressed formally (mathematically) through appealing to beliefs about the causes of our sensations, and about the actions we pursue. This is used to motivate an approach to active vision that depends upon inferences about the causes of ‘what I have seen’ and learning about ‘what I would see if I were to look there’. The former could manifest as persistent ‘delay-period’ activity – of the sort associated with working memory, while the latter is better suited to changes in synaptic efficacy – of the sort that underlies short-term learning and adaptation. We review formulations of these ideas in terms of active inference, their role in directing visual exploration, and the consequences – for active vision – of their failures. To illustrate the latter, we draw upon some of our recent work on the computational anatomy of visual neglect.

Keywords: Active inference; Novelty; Visual neglect; Saccades; Markov Decision Process; Active vision

Introduction

Active inference is a principled framework for describing Bayes optimal behaviour. It formalises the notion of perception as hypothesis testing (Gregory 1980, Friston, Adams et al. 2012), emphasising the active manipulation of sensory epithelia to perform perceptual experiments. Here, we focus

upon active vision (Wurtz, McAlonan et al. 2011, Andreopoulos and Tsotsos 2013, Ognibene and Baldassarre 2014, Mirza, Adams et al. 2016). Our formulation implies that perception of space is fundamentally tied to motor representations, as visual input at a point in space is the consequence of a saccade to that location (Zimmermann and Lappe 2016). This enactive view also brings with it perspectival aspect to the way that we sample – and internalise – our lived world: see (Rudrauf, Bennequin et al. 2017) for a treatment of the implicit phenomenology. In what follows, we provide an overview of active inference in the context of visual foraging. This takes the form of a generic description in terms of Markov Decision Processes and draws examples from our recent work to illustrate these ideas in concrete scenarios. We will illustrate visual foraging in an environment with volatile contingencies, to disclose the intimate relationship between past beliefs (i.e., working memory and attention) and future plans (i.e., working memory and intention). We then turn to scenes with multiple stimuli and show that the optimisation of beliefs about the causes of sensory impressions is an efficient way to deal with such environments. To demonstrate the importance of this sort of inference, we illustrate – via simulations – how these mechanisms can fail. This provides us with several plausible mechanisms that could account for visual neglect – a common neuropsychological syndrome that reflects a cardinal failure of active vision.

A note on terminology

In the following, we call on some technical concepts – e.g. evidence, free energy, and generative model. In this section, we briefly introduce the terms that will become important in the later parts of the chapter. Most of these terms reflect the fact that we are trying to understand behaviour in terms of a principled sampling of the world to make perceptual inferences. Formally speaking, this (active) inference can be described in terms of probability distributions or ‘beliefs’. Priors and posteriors are the probabilities before and after making an observation (respectively). Likelihoods quantify the probability of an observation given we knew how it was caused. A generative model is a probabilistic specification of beliefs about how a certain kind of sensory data is generated. It comprises priors and likelihoods and can be used to make predictions about the most probable cause of sensations; namely the posterior. Model evidence is the probability that a model could account for some given data. This is sometimes referred to as a marginal likelihood, or negative surprise. Self-evidence depends upon a ‘self-model’ that specifies the kinds of data that will be observed by an organism given the sort of creature it is. For example, given that a creature is a fish, it might expect to be surrounded by water. Observing this carries high self-evidence, while finding itself on dry-land carries little evidence for its continued existence. From this, it is intuitively sensible that self-evidence should not drop too low.

Negative free energy (sometimes referred to as an ‘evidence lower bound’) is a quantity that is always less than the evidence for a model. Minimising free energy thus maximises the lower bound on evidence and ensures that self-evidence is maintained above a certain value. This is just a mathematical way of describing behaviour that minimises surprise (i.e., maximises model evidence) and thereby resolves uncertainty about sensory exchanges with the world.

Active inference

Bayes optimal behaviour requires that action maximises self-evidence (Hohwy 2016) or, equivalently, minimises surprise (Friston, Daunizeau et al. 2010). This is not always possible (or efficient) to compute directly, but can be approximated by a free energy bound (Beal 2003), as illustrated by Jensen’s inequality:

$$F = \underbrace{-\mathbb{E}_Q \left[\ln \frac{P(\tilde{o}, \tilde{s}, \pi)}{Q(\tilde{s}, \pi)} \right]}_{\text{Jensen's inequality}} \geq -\ln \mathbb{E}_Q \left[\frac{P(\tilde{o}, \tilde{s}, \pi)}{Q(\tilde{s}, \pi)} \right] = \underbrace{-\ln P(\tilde{o})}_{\text{Surprise}}$$

In this equation, P is the probability distribution that represents a generative model. This probabilistic model expresses the beliefs an animal has about the way in which its sensations (\tilde{o}) are generated from hidden (unobservable) states (\tilde{s}) and the policy (π), or action sequence that it is pursuing. The tilde (\sim) notation indicates a trajectory (or sequence) through time. The distribution Q may be any arbitrary distribution but takes on an interesting interpretation when the free energy is minimised. A rearrangement of the equation above gives

$$F = D_{KL}[Q(\tilde{s}, \pi) || P(\tilde{s}, \pi | \tilde{o})] - \ln P(\tilde{o})$$

This shows that the difference between surprise and the free energy is the KL-Divergence between Q and the posterior distribution. This is minimal when the two distributions are the same, indicating that free energy minimisation entails finding Q that approximates the posterior probability. Under mean-field assumptions (Feynman 1998) we rewrite the variational free energy:

$$F = \mathbb{E}_Q[F(\pi)] + D_{KL}[Q(\pi) || P(\pi)]$$

$$F(\pi) = -\mathbb{E}_Q[\ln P(\tilde{o}, \tilde{s} | \pi) - \ln Q(\tilde{s} | \pi)]$$

This allows us to separate out perceptual inference (optimisation of $Q(\tilde{s} | \pi)$) from planning (optimisation of $Q(\pi)$). Planning requires us to define a prior probability distribution over possible action sequences. A self-consistent prior for a free energy minimising agent is that the most

Novelty and neglect

probable policies are those that minimise expected free energy ($G(\pi)$) (Friston, FitzGerald et al. 2017).

$$P(\pi) = \sigma(-G(\pi))$$

$$G(\pi) = \sum_{\tau > t} G(\pi, \tau)$$

$$G(\pi, \tau) = -\mathbb{E}_{\tilde{Q}}[\ln P(o_\tau, s_\tau | \pi) - \ln Q(s_\tau | \pi)]$$

$$\tilde{Q}(o_\tau, s_\tau | \pi) = P(o_\tau | s_\tau)Q(s_\tau | \pi)$$

Under these self-evidencing prior beliefs, the probability of each policy is defined by applying a softmax function (σ) to the negative expected free energy of policies or sequences of actions that are currently entertained (e.g., ‘where to look next’). To gain some intuition for the consequences of minimising expected free energy (i.e., expected surprise), we can express the expected free energy in the following form.

$$G(\pi, \tau) = -\mathbb{E}_{\tilde{Q}}[\underbrace{\ln P(s_\tau | o_\tau, \pi) - \ln Q(s_\tau | \pi)}_{\text{Epistemic value}} + \underbrace{\ln P(o_\tau)}_{\text{Extrinsic value}}]$$

The first pair of terms on the right shows that the expected free energy will be smaller if the (expected) difference between the posterior belief before and after the next observation is large. In other words, a policy will be more likely if its sensory consequences bring about a larger change in beliefs. This epistemic value is also known as salience, information gain, Bayesian surprise or epistemic affordance. The final term says that policies are more probable if they fulfil prior beliefs (preferences) about sensations that will be sampled under that policy. Together, these terms ensure explorative, information seeking (epistemic) and exploitative, goal seeking (pragmatic) behaviour (Friston, Rigoli et al. 2015).

It remains for us to specify the form of the generative model (P). In this chapter, we appeal to a generative model that takes the form of a Markov Decision Process, as depicted in Figure 1. This is a discrete state-space model that incorporates the variables considered above; namely, states of the world and the action sequences that change those states. In a Markov decision process, hidden states evolve through time according to a probability transition matrix (\mathbf{B}). At each time, a state depends upon only the state at the previous time, and upon the policy pursued. States give rise to observations as determined by a likelihood distribution (\mathbf{A}). The probability of initial states is given by the distribution \mathbf{D} , while preferences are defined by \mathbf{C} . In the next section, we will look at the interplay between perceptual inference (i.e., optimising beliefs about states of the world) and

planning (i.e., optimising beliefs about policies from which actions are selected). This inference calls for a form of working memory that could be mediated by persistent neuronal firing rates (Funahashi, Bruce et al. 1989, Friston, Rosch et al. 2017, Parr and Friston 2017). We then turn to beliefs about the probability distributions, and an elemental form of short term memory that could be synaptically mediated (Hempel, Hartman et al. 2000, Wang, Markram et al. 2006, Mongillo, Barak et al. 2008, Parr and Friston 2017).

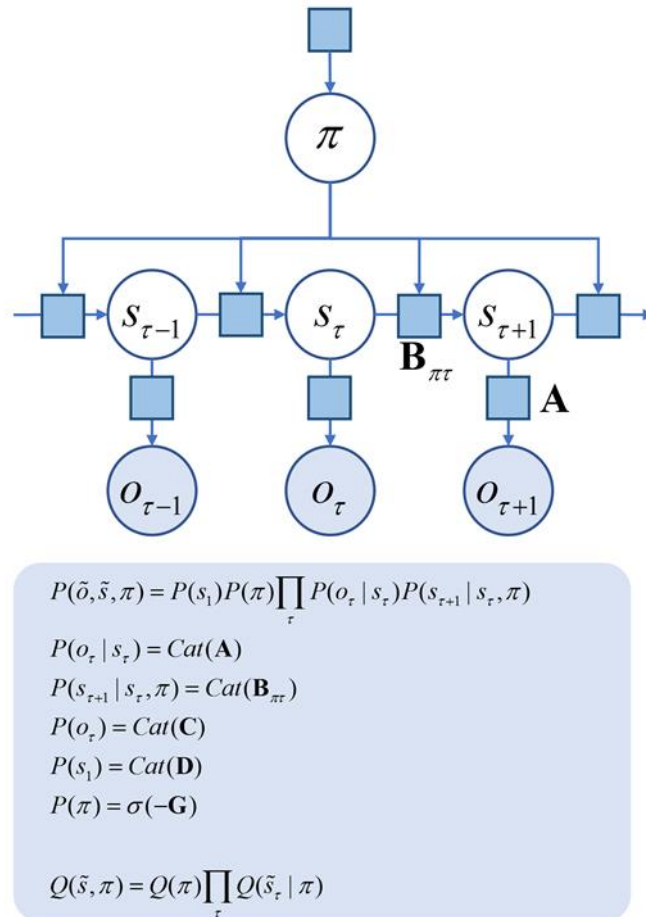


Figure 1 – Markov Decision Process This graphic shows a Markov Decision Process in factor graph form (Loeliger, Dauwels et al. 2007). Blue squares indicate probability distributions (factors of the generative model). In the lower panel, these factors are expressed in terms of probability matrices. ‘Cat’ denotes a categorical distribution. The mean-field factorisation of the approximate distribution, Q is shown. This Figure has been adapted from (Parr and Friston Accepted).

The neurobiology of active inference

Novelty and neglect

While the mathematical formalism presented above might seem a little abstract, active inference is associated with a process theory (Friston, FitzGerald et al. 2017) that interprets these quantities in terms of their biological substrates. For example, the approximate posterior beliefs (Q) can be thought of as neuronal (firing rate) representations. The conditional probability distributions (\mathbf{A}, \mathbf{B}) then represent the synaptic connections between different populations of neurons. Quantities like the expected free energy (\mathbf{G}) are computed from the above variables to evaluate competing plans of action – this evaluation could take place in cortico-basal ganglia-cortical loops. There have been many papers linking the computational anatomy and to functional anatomy in the human brain. These range from treatments of cortical hierarchies through to associative plasticity. In the next sections, we will focus on precision (the confidence in a given synapse) –that may be mediated by ascending neuromodulatory influences – and learning, which involves plastic modulation of synapses through long term potentiation or depression.

Inference and salience

Perceptual inference is the process of optimising beliefs about states (under a given policy). By minimising free energy, we find that inference or belief updating can be expressed as

$$Q(s_\tau | \pi) = \sigma \left(E_{Q(s_{\tau-1}|\pi)} [\ln P(s_\tau | s_{\tau-1}, \pi)] + E_{Q(s_{\tau+1}|\pi)} [\ln P(s_{\tau+1} | s_\tau, \pi)] + \ln P(o_\tau | s_\tau) \right)$$

The first term on the right entails a form of memory, because beliefs about a state in the past are used to inform beliefs about the present (via the transition probabilities). However, the capacity for the past to inform the present depends upon the degree of uncertainty associated with transitions. Intuitively, if the environment changes randomly from one time to the next, the past is of no use in inferring the present. Conversely, in a deterministic environment, the present can be reliably inferred using beliefs about the past. We have previously described a way to parameterise transition probabilities so that we can manipulate the precision (ω) of these transitions (Parr and Friston 2017). When $\omega = 0$, there is no relationship between past and present. As $\omega \rightarrow \infty$, transitions tend towards being deterministic. Another way to think about this is that ω scales with the negative entropy of transition probabilities, also known as volatility (Parr and Friston 2017). Neurobiologically, this suggests that this parameter may play the role of a synaptic gain – controlling the degree to which neurons representing the past can influence those representing the present. In turn, this implicates neuromodulatory transmitters, such as noradrenaline, that have been associated with signalling volatility (Dayan and Yu 2006, Marshall, Mathys et al. 2016).

This has important consequences for visual search, and the selection of saccadic targets that minimise expected free energy. If we treat eye position as a hidden state (with transitions dependent on policies or saccadic sequences), this hidden state determines which other hidden states are informed by foveal visual data (Mirza, Adams et al. 2016). Figure 2 shows four objects (squares of various shades). Performing a saccade to a location means foveal vision will provide precise information about the object at that location. If the shade of each square can change, with a certain probability, this has important consequences for our belief about the shades of squares not currently foveated. Specifically, if ω is high, the confidence in $Q(s_\tau | \pi)$ will be relatively high and the expected free energy of a saccade to that location will be high. In other words, actively sampling other squares would have a low epistemic affordance. This would make for a poor perceptual experiment, as it would do little to change beliefs. If low, the confidence in $Q(s_\tau | \pi)$ will decrease with each time step, making it an increasingly plausible (i.e., salient) saccadic target.

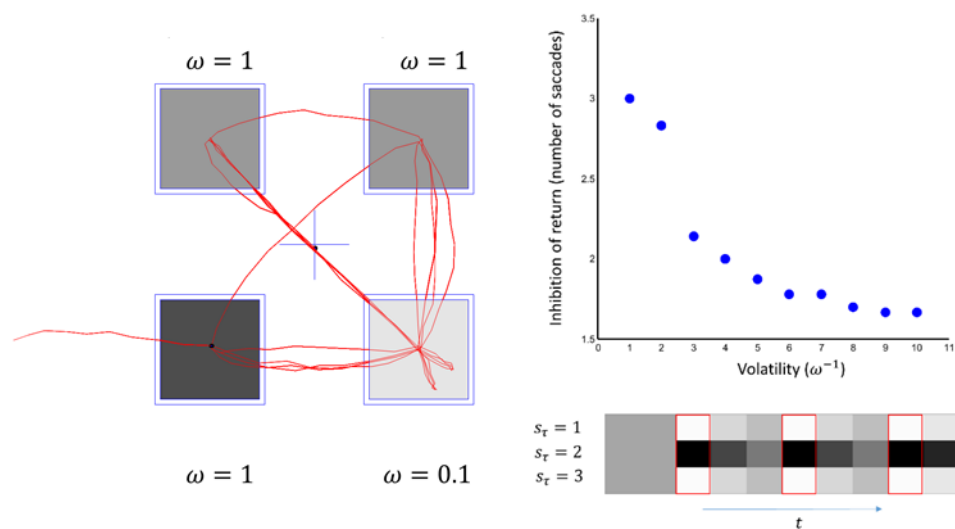


Figure 2 – Persistent activity and inhibition of return *On the left*, we show the simulated eye tracking data (over 19 saccades), when each hidden case is associated with a differing level of volatility (ω^{-1}). *The right upper plot* shows the influence of the volatility of beliefs on the simulated inhibition of return (quantified by the average number of saccades between fixations) for one of the stimulus locations. *The right lower plot* shows the simulated neuronal encoding over time for three units, each representing a different hypothesis about the identity (1, 2, or 3) of one of the hidden states. Darker shades indicate greater firing rates. The red boxes indicate the times during which the location corresponding to this hidden state is foveated. This Figure has been reproduced from (Parr and Friston 2017).

Novelty and neglect

The simulated eye-tracking trace (Parr and Friston 2017) in Figure 2 shows that those locations associated with precise (relatively deterministic) transition probabilities are less frequently the targets of saccades compared to the more volatile lower right square. This formulation of perisaccadic working memory, and its role in directing the eyes towards salient locations, reproduces the phenomenon of ‘inhibition of return’ (Posner, Rafal et al. 1985, Klein 2000), and shows how a Bayes optimal agent would modulate this at various levels of transition precisions. The lower right plot shows how expectations that a square possesses one of three alternative shades become more dispersed (as uncertainty accumulates) over time until a saccade is made to that location. This can be interpreted as a ‘raster plot’ showing the firing rates of three neurons (each row) over time. Each of these neurons encodes a probability that the square has one of three shades over time. This is exactly the profile of neuronal responses one sees in delayed saccade experiments used to elicit the neuronal correlates of working memory: see, for example (Kojima and Goldman-Rakic 1982). In short, simply inferring the best thing to do next mandate the encoding of trajectories of (hidden) states of the world that necessarily requires a neuronal encoding of beliefs about the past (and future. Crucially, this neuronal (working memory) encoding transcends the current moment, equipping the perceptual inference with a mnemonic aspect that allows for both prediction and postdiction. It is this encoding we suggest as a canonical form of working memory.

In this section, we have examined the relationship between salience (the potential to resolve uncertainty) and the sort of working memory that might be represented by persistently active neurons. In the next section, we turn to novelty, learning and their link with synaptic memory.

Learning and novelty

In this section, we draw from observations concerning visual neglect (Halligan and Marshall 1998). This is a disorder of active vision that biases saccadic exploration away from the left side of space (Husain, Mannan et al. 2001, Fruhmann Berger, Johannsen et al. 2008, Karnath and Rorden 2012). A common pen-and-paper test – used to assess this deficit – is the line cancellation task (illustrated in Figure 3) (Albert 1973, Fullerton, McSherry et al. 1986, Ferber and Karnath 2001). Patients are asked to cancel, by crossing out, all of the lines on a sheet of paper. Typically, neglect patients will cross out those on the right, but miss those on the left. Similar effects have been found in oculomotor variants of this task, but have demonstrated remarkably few re-cancellations in healthy participants (Mannan, Mort et al. 2005); even when there is no explicit (visual) cue to record that a target has been cancelled. The middle and right panels of Figure 3 make the point that maintaining a memory of previously visited locations would require many more neurons, if stored as persistent or delay

period activity, than if this information were stored through short term plastic changes (Parr and Friston 2017).




Figure 3 – Persistent activity or synaptic efficacy? *On the left*, an example of a line cancellation task is shown. The subject is presented with a sheet of paper with a set of horizontal lines, and is asked to cancel (red marks) each of these lines. *The middle panel* shows the ensemble of 192 neurons that would be required to represent the subject’s beliefs about where the lines are, and whether they have cancelled them, if the memory of previously visited locations were stored as persistent activity in a neuronal population. The currently active neurons are represented by a black outline. *The panel on the right* shows a much more efficient way to represent this information; in terms of a mapping from a representation of space to representations of each of the possible observations that could be made on visiting a particular location. Clearly it is more efficient to make use of synaptic efficacy when storing transient, high dimensional, memories. In short, synaptic efficacy represents probabilistic mappings (i.e., ‘if I were to look there, I would see that’) as opposed to beliefs about the current state of the world (i.e., ‘I am looking there’ or ‘seeing that’) encoded by synaptic activity. This figure has been adapted from (Parr and Friston 2017).

Under the process theory (Friston, FitzGerald et al. 2017) associated with active inference, we can think of synaptic plasticity in terms of changes in beliefs about conditional probabilities (Friston, FitzGerald et al. 2016). This is because synapses mediate the influence of beliefs represented in one population of neurons on another. For these to change, we supplement our generative model with beliefs about the parameters of the mapping from ‘where I am looking’ to ‘what I see’. When we do this for the parameters of the likelihood mapping, the expected free energy becomes:

Novelty and neglect

$$\begin{aligned}
 G(\pi, \tau) = & - \underbrace{E_{\tilde{Q}}[\ln P(s_\tau | A, o_\tau, \pi) - \ln Q(s_\tau | \pi)]}_{\text{Saliency}} \\
 & - \underbrace{E_{\tilde{Q}}[\ln P(A | o_\tau, \pi) - \ln Q(A)]}_{\text{Novelty}} \\
 & - \underbrace{E_{\tilde{Q}}[\ln P(o_\tau)]}_{\text{Extrinsic value}}
 \end{aligned}$$

Here, we have decomposed the epistemic value into the saliency term used in the previous section, and a novelty term, that relates to the expected change in beliefs about the parameters of the likelihood distribution. Figure 4, shows a simulation of an oculomotor cancellation task (Parr and Friston 2017). This incorporates beliefs about what would be seen conditioned upon the eye position. These beliefs are optimised by accumulation of Dirichlet parameters (Beal 2003, Blei, Ng et al. 2003, Friston, FitzGerald et al. 2016) – this closely resembles activity dependent plasticity (Hebb 1949, Brown, Zhao et al. 2009), as the element of the likelihood matrix (\mathbf{A}) representing a state-outcome mapping is increased whenever the two occur simultaneously. This means that, on fixating a location, uncertainty is resolved about the visual data obtained by saccades to that location and its novelty is decreased (Schwartenbeck, FitzGerald et al. 2013).

Although phenomenologically, this sort of memory is very different from the working memory of previous section, they both emerge from the minimisation expected free energy or expected surprise (i.e., entropy or uncertainty). When resolving uncertainty about states of affairs in the world we sample salient information. However, when resolving uncertainty about the contingencies the corresponding epistemic affordance becomes novelty; i.e., the opportunity to resolve uncertainty about "what would happen if I did that?" If updating beliefs about states of the world and parameters – that underwrite probabilistic contingencies – correspond to perceptual inference and learning respectively, then novelty is to learning as saliency is to inference.

Figure 4 also shows the consequences of lesioning the generative model (Parr and Friston 2017). Increasing the Dirichlet parameters for the left side of space means that (despite representing the same likelihood distribution) the capacity for the beliefs about the likelihood to change is diminished. Cutting the connection between two regions would have this effect, as no plastic changes can occur following disconnection. This means there is little novelty to resolve on the left, biasing saccades to the right. This is consistent with the idea that visual neglect is a 'disconnection syndrome' (Geschwind 1965, Catani and ffytche 2005, Bartolomeo, Thiebaut de Schotten et al. 2007, He, Snyder et al. 2007). We additionally show the effect of including a preference (\mathbf{C}) for proprioceptive outcomes on the right of space, and of directly biasing the prior beliefs over policies in favour of rightward saccades. The image on the left of Figure 4 shows common neuroanatomical

sites for lesions that induce neglect (Doricchi and Tomaiuolo 2003, Thiebaut de Schotten, Urbanski et al. 2005, Karnath and Rorden 2012) and relates these to the computational lesions shown on the right. In brief, this account suggests that the superior longitudinal fasciculus, that connects frontal to temporoparietal regions, might be the substrate of the likelihood distribution representing the visual consequences of each fixation. Plastic changes in this tract would then underwrite the learning described above, and its dysfunction would compromise the resolution of novelty following a given fixation. The dorsal pulvinar (Ungerleider and Christensen 1979), that projects to parietal cortex (Weller, Steele et al. 2002, Behrens, Johansen-Berg et al. 2003), may signal prior beliefs (preferences) about the ‘where’ data that the brain might seek out. Damage could bias these towards the left. The evaluation of competing saccadic policies is likely to implicate cortico-basal ganglia circuits (Schiller, True et al. 1980, Hikosaka and Wurtz 1985a, Schiller, Sandell et al. 1987), and disruption of these could manifest as a bias towards rightward saccades. In this – and the previous section – we have shown that disruption of any part of the expected free energy (saliency, novelty, or extrinsic value) can have drastic consequences for active vision.

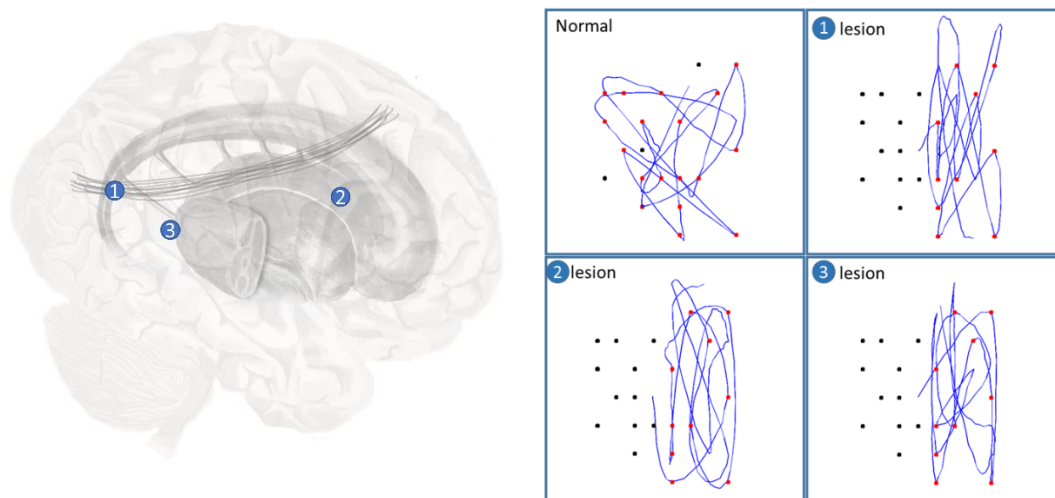


Figure 4 – The computational anatomy of visual neglect *On the left*, three lesions implicated in visual neglect are highlighted here. 1 – Disconnection of the second branch of the right superior longitudinal fasciculus (a white matter tract that connects dorsal frontal with ventral parietal regions (Makris, Kennedy et al. 2004, Bartolomeo, Thiebaut de Schotten et al. 2012)); 2 – Unilateral lesion to the right putamen (Karnath, Himmelbach et al. 2002); 3 – Unilateral lesion to the right pulvinar (a thalamic nucleus). This has been reproduced from (Parr, Rees et al. 2018). *On the right*, each panel shows the simulated eye tracking data (blue) during 20 saccades. In all cases, the target array was the same. The upper left panel shows the performance of the model with no simulated lesions. The upper right panel

Novelty and neglect

shows the results when the Dirichlet parameters associated with the likelihood (i.e., hidden state to outcome) mapping were increased for the left hemifield, corresponding to a functional disconnection of the second branch of the right superior longitudinal fasciculus. The lower left panel shows performance when there is a biasing of policy selection, simulating a lesion of the putamen. The lower right panel represents a lesion of the prior beliefs about proprioceptive outcomes, which relates to a deficit in the inputs to the dorsal parietal cortex, likely from the pulvinar. Notably, the lesioned scan paths look very similar to one another. By eye, it is almost impossible to infer which path was generated by which lesion. Despite this, it is possible to disambiguate between the cortical disconnection (1) and the two subcortical lesions (2 and 3) using Bayesian model comparison (i.e. fitting each lesioned model to synthetic data, and comparing the evidence for each model) – for details, please see (Parr and Friston 2017). The scan paths here have been reproduced from (Parr and Friston 2017)

Conclusion

In this chapter, we have reviewed the basic principles of active inference, and have illustrated the importance of holding beliefs about hidden states and parameters for planning and action. In brief, uncertainty about hidden states renders those actions that resolve uncertainty salient. Uncertainty about parameters presents the opportunity for resolving uncertainty by selecting actions that entail novelty. Interestingly, each of these beliefs equips actively inferring creatures with a form of memory. Optimisation of beliefs about states relies in part upon the propagation of beliefs about the past to the present, suggesting a form of working memory mediated by persistent activity in neuronal populations representing relevant beliefs. However, it is more efficient, for higher dimensional memories, to represent context sensitive states of the sampled world as relationships between variables; e.g. the mapping from ‘where I look’ to ‘what I see’. This implicates changes in the connections between neuronal populations, and short-term plasticity. We have illustrated the consequences of the requisite belief updating – for active vision – and the repertoire of computational lesions that could lead to visual neglect.

Acknowledgements

TP is supported by the Rosetrees Trust (Award Number 173346). KJF is a Wellcome Principal Research Fellow (Ref: 088130/Z/09/Z).

Disclosure statement

The authors have no disclosures or conflict of interest.

References

- Albert, M. L. (1973). "A simple test of visual neglect." *Neurology* **23**(6): 658.
- Andreopoulos, A. and J. Tsotsos (2013). "A computational learning theory of active object recognition under uncertainty." *International journal of computer vision* **101**(1): 95-142.
- Bartolomeo, P., M. Thiebaut de Schotten and A. B. Chica (2012). "Brain networks of visuospatial attention and their disruption in visual neglect." *Frontiers in Human Neuroscience* **6**: 110.
- Bartolomeo, P., M. Thiebaut de Schotten and F. Doricchi (2007). "Left Unilateral Neglect as a Disconnection Syndrome." *Cerebral Cortex* **17**(11): 2479-2490.
- Beal, M. J. (2003). Variational algorithms for approximate Bayesian inference, University of London United Kingdom.
- Behrens, T. E. J., H. Johansen-Berg, M. W. Woolrich, S. M. Smith, C. A. M. Wheeler-Kingshott, P. A. Boulby, G. J. Barker, E. L. Sillery, K. Sheehan, O. Ciccarelli, A. J. Thompson, J. M. Brady and P. M. Matthews (2003). "Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging." *Nat Neurosci* **6**(7): 750-757.
- Blei, D. M., A. Y. Ng and M. I. Jordan (2003). "Latent dirichlet allocation." *J. Mach. Learn. Res.* **3**: 993-1022.
- Brown, T. H., Y. Zhao and V. Leung (2009). Hebbian Plasticity A2 - Squire, Larry R. *Encyclopedia of Neuroscience*. Oxford, Academic Press: 1049-1056.
- Catani, M. and D. H. ffytche (2005). "The rises and falls of disconnection syndromes." *Brain* **128**(10): 2224-2239.
- Dayan, P. and A. J. Yu (2006). "Phasic norepinephrine: A neural interrupt signal for unexpected events." *Network: Computation in Neural Systems* **17**(4): 335-350.
- Doricchi, F. and F. Tomaiuolo (2003). "The anatomy of neglect without hemianopia: a key role for parietal–frontal disconnection?" *Neuroreport* **14**(17): 2239-2243.
- Ferber, S. and H.-O. Karnath (2001). "How to assess spatial neglect-line bisection or cancellation tasks?" *Journal of clinical and experimental neuropsychology* **23**(5): 599-607.
- Feynman, R. P. (1998). *Statistical Mechanics: A Set Of Lectures*, Avalon Publishing.
- Friston, K., R. A. Adams, L. Perrinet and M. Breakspear (2012). "Perceptions as Hypotheses: Saccades as Experiments." *Frontiers in Psychology* **3**: 151.
- Friston, K., T. FitzGerald, F. Rigoli, P. Schwartenbeck, J. O'Doherty and G. Pezzulo (2016). "Active inference and learning." *Neuroscience & Biobehavioral Reviews* **68**: 862-879.
- Friston, K., T. FitzGerald, F. Rigoli, P. Schwartenbeck and G. Pezzulo (2017). "Active Inference: A Process Theory." *Neural Comput* **29**(1): 1-49.
- Friston, K., F. Rigoli, D. Ognibene, C. Mathys, T. Fitzgerald and G. Pezzulo (2015). "Active inference and epistemic value." *Cognitive Neuroscience* **6**(4): 187-214.
- Friston, K. J., J. Daunizeau, J. Kilner and S. J. Kiebel (2010). "Action and behavior: a free-energy formulation." *Biological Cybernetics* **102**(3): 227-260.
- Friston, K. J., R. Rosch, T. Parr, C. Price and H. Bowman (2017). "Deep temporal models and active inference." *Neuroscience & Biobehavioral Reviews* **77**: 388-402.
- Fruhmann Berger, M., L. Johannsen and H.-O. Karnath (2008). "Time course of eye and head deviation in spatial neglect." *Neuropsychology* **22**(6): 697-702.
- Fullerton, K. J., D. McSherry and R. W. Stout (1986). "Albert's Test: a Neglected Test of Perceptual Neglect." *The Lancet* **327**(8478): 430-432.
- Funahashi, S., C. J. Bruce and P. S. Goldman-Rakic (1989). "Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex." *Journal of Neurophysiology* **61**(2): 331.

- Geschwind, N. (1965). "DISCONNEXION SYNDROMES IN ANIMALS AND MAN1." Brain **88**(2): 237-237.
- Gregory, R. L. (1980). "Perceptions as Hypotheses." Philosophical Transactions of the Royal Society of London. B, Biological Sciences **290**(1038): 181.
- Halligan, P. W. and J. C. Marshall (1998). "Neglect of Awareness." Consciousness and Cognition **7**(3): 356-380.
- He, B. J., A. Z. Snyder, J. L. Vincent, A. Epstein, G. L. Shulman and M. Corbetta (2007). "Breakdown of Functional Connectivity in Frontoparietal Networks Underlies Behavioral Deficits in Spatial Neglect." Neuron **53**(6): 905-918.
- Hebb, D. O. (1949). "The first stage of perception: Growth of the assembly." The Organization of Behavior: 60-78.
- Hempel, C. M., K. H. Hartman, X. J. Wang, G. G. Turrigiano and S. B. Nelson (2000). "Multiple Forms of Short-Term Plasticity at Excitatory Synapses in Rat Medial Prefrontal Cortex." Journal of Neurophysiology **83**(5): 3031.
- Hikosaka, O. and R. H. Wurtz (1985a). "Modification of saccadic eye movements by GABA-related substances. I. Effect of muscimol and bicuculline in monkey superior colliculus." Journal of Neurophysiology **53**(1): 266.
- Hohwy, J. (2016). "The Self-Evidencing Brain." Noûs **50**(2): 259-285.
- Husain, M., S. Mannan, T. Hodgson, E. Wojciulik, J. Driver and C. Kennard (2001). "Impaired spatial working memory across saccades contributes to abnormal search in parietal neglect." Brain **124**(5): 941-952.
- Karnath, H.-O. and C. Rorden (2012). "The anatomy of spatial neglect." Neuropsychologia **50**(6): 1010-1017.
- Karnath, H. O., M. Himmelbach and C. Rorden (2002). "The subcortical anatomy of human spatial neglect: putamen, caudate nucleus and pulvinar." Brain **125**(2): 350-360.
- Klein, R. M. (2000). "Inhibition of return." Trends in Cognitive Sciences **4**(4): 138-147.
- Kojima, S. and P. S. Goldman-Rakic (1982). "Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response." Brain Res. **248**(1): 43-49.
- Loeliger, H. A., J. Dauwels, J. Hu, S. Korl, L. Ping and F. R. Kschischang (2007). "The Factor Graph Approach to Model-Based Signal Processing." Proceedings of the IEEE **95**(6): 1295-1322.
- Makris, N., D. N. Kennedy, S. McInerney, A. G. Sorensen, R. Wang, J. V. S. Caviness and D. N. Pandya (2004). "Segmentation of Subcomponents within the Superior Longitudinal Fascicle in Humans: A Quantitative, In Vivo, DT-MRI Study." Cerebral Cortex **15**(6): 854-869.
- Mannan, S. K., D. J. Mort, T. L. Hodgson, J. Driver, C. Kennard and M. Husain (2005). "Revisiting Previously Searched Locations in Visual Neglect: Role of Right Parietal and Frontal Lesions in Misjudging Old Locations as New." Journal of Cognitive Neuroscience **17**(2): 340-354.
- Marshall, L., C. Mathys, D. Ruge, A. O. de Berker, P. Dayan, K. E. Stephan and S. Bestmann (2016). "Pharmacological Fingerprints of Contextual Uncertainty." PLOS Biology **14**(11): e1002575.
- Mirza, M. B., R. A. Adams, C. D. Mathys and K. J. Friston (2016). "Scene Construction, Visual Foraging, and Active Inference." Front Comput Neurosci **10**: 56.
- Mirza, M. B., R. A. Adams, C. D. Mathys and K. J. Friston (2016). "Scene Construction, Visual Foraging, and Active Inference." Frontiers in Computational Neuroscience **10**(56).
- Mongillo, G., O. Barak and M. Tsodyks (2008). "Synaptic Theory of Working Memory." Science **319**(5869): 1543-1546.
- Ognibene, D. and G. Baldassarre (2014). Ecological Active Vision: Four Bio-Inspired Principles to Integrate Bottom-Up and Adaptive Top-Down Attention Tested With a Simple Camera-Arm Robot. IEEE Transactions on Autonomous Mental Development, IEEE.
- Parr, T. and K. J. Friston (2017). "The active construction of the visual world." Neuropsychologia **104**: 92-101.
- Parr, T. and K. J. Friston (2017). "The Computational Anatomy of Visual Neglect." Cerebral Cortex: 1-14.

- Parr, T. and K. J. Friston (2017). "Uncertainty, epistemics and active inference." Journal of The Royal Society Interface **14**(136).
- Parr, T. and K. J. Friston (2017). "Working memory, attention, and salience in active inference." Scientific Reports **7**(1): 14678.
- Parr, T. and K. J. Friston (Accepted). "From decisions to movements - and back again." Neural Computation.
- Parr, T., G. Rees and K. J. Friston (2018). "Computational Neuropsychology and Bayesian Inference." Frontiers in Human Neuroscience **12**(61).
- Posner, M. I., R. D. Rafal, L. S. Choate and J. Vaughan (1985). "Inhibition of return: Neural basis and function." Cognitive Neuropsychology **2**(3): 211-228.
- Rudrauf, D., D. Bennequin, I. Granic, G. Landini, K. Friston and K. Williford (2017). "A mathematical model of embodied consciousness." J Theor Biol **428**: 106-131.
- Schiller, P. H., J. H. Sandell and J. H. Maunsell (1987). "The effect of frontal eye field and superior colliculus lesions on saccadic latencies in the rhesus monkey." Journal of Neurophysiology **57**(4): 1033.
- Schiller, P. H., S. D. True and J. L. Conway (1980). "Deficits in eye movements following frontal eye-field and superior colliculus ablations." Journal of Neurophysiology **44**(6): 1175.
- Schwartenbeck, P., T. FitzGerald, R. J. Dolan and K. Friston (2013). "Exploration, novelty, surprise, and free energy minimization." Frontiers in Psychology **4**: 710.
- Thiebaut de Schotten, M., M. Urbanski and H. Duffau (2005). "Direct evidence for a parietal-frontal pathway subserving spatial awareness in humans." Science **309**.
- Ungerleider, L. G. and C. A. Christensen (1979). "Pulvinar lesions in monkeys produce abnormal scanning of a complex visual array." Neuropsychologia **17**(5): 493-501.
- Wang, Y., H. Markram, P. H. Goodman, T. K. Berger, J. Ma and P. S. Goldman-Rakic (2006). "Heterogeneity in the pyramidal network of the medial prefrontal cortex." Nat Neurosci **9**(4): 534-542.
- Weller, R. E., G. E. Steele and J. H. Kaas (2002). "Pulvinar and other subcortical connections of dorsolateral visual cortex in monkeys." The Journal of Comparative Neurology **450**(3): 215-240.
- Wurtz, R. H., K. McAlonan, J. Cavanaugh and R. A. Berman (2011). "Thalamic pathways for active vision." Trends Cogn Sci **5**(4): 177-184.
- Zimmermann, E. and M. Lappe (2016). "Visual Space Constructed by Saccade Motor Maps." Frontiers in Human Neuroscience **10**(225).