

# **The Predictive Global Neuronal Workspace: A Formal Active Inference Model of Visual Consciousness**

Christopher J. Whyte<sup>1</sup> and Ryan Smith<sup>2</sup>

Author affiliations:

<sup>1</sup>School of Psychology, University of Sydney, NSW, Australia

<sup>2</sup>Laureate Institute for Brain Research, Tulsa, OK, USA

Corresponding author email: [christopherjackwhyte@gmail.com](mailto:christopherjackwhyte@gmail.com)

Key words:

active inference, consciousness, global neuronal workspace, predictive global neuronal workspace, computational neuroscience, predictive processing

**Abstract:** The global neuronal workspace (GNW) model has inspired over two decades of hypothesis driven research on the neural basis consciousness. However, recent studies have reported findings that appear inconsistent with the predictions of the model. Further, the macroanatomical focus of current GNW research has limited the specificity of predictions afforded by the model. In this paper we present a neurocomputational model – based on the active inference framework – that captures central architectural elements of the GNW and that can address these limitations. The resulting ‘predictive global workspace’ casts neuronal dynamics as approximating Bayesian inference, allowing precise, testable predictions at both the behavioural and neural levels of description. We report simulations demonstrating the model’s ability to reproduce: 1) the electrophysiological and behaviour results observed in previous studies of inattention blindness, and 2) the previously described four-way taxonomy predicted by the GNW, which describes the relationship between consciousness, attention, and sensory signal strength. We then illustrate how our model can reconcile/explain (apparently) conflicting findings, extend the GNW taxonomy to include the influence of prior expectations, and inspire novel paradigms to test associated behavioural and neural predictions.

## 1. Introduction

Global workspace theory (GWT) is one of the most widely supported neuroscientific theory of consciousness. Indeed, a survey conducted at the 2018 Association for the Scientific Study of Consciousness showed that, among experts, GWT is regarded as the most promising theory of consciousness (Michel et al., 2018).

GWT was first proposed by Baars (1988) as a cognitive architecture that identifies consciousness with the global availability of information. According to GWT, information becomes conscious when it is simultaneously made available to a wide range of localized (and individually sub-personal) processes – jointly comprising a ‘global workspace’. More recently, Dehaene and colleagues have advanced a global *neuronal* workspace (GNW) model, which identifies the global workspace with a large-scale network of excitatory pyramidal neurons with long-range axonal pathways connecting prefrontal and parietal cortices (Dehaene & Changeux, 2011; Dehaene et al., 2011).

The key working hypothesis of the GNW is that when a stimulus becomes conscious there will be a late, non-linear, all-or-nothing “ignition” of prefrontal and parietal regions (Dehaene and Changeux, 2011; Dehaene, 2014) corresponding to the large-scale influence of selected (otherwise unconscious) representations of perceptual features encoded locally within sensory cortices. In contrast, activity related to stimuli that is rendered unconscious (e.g., by masking or inattention) will fail to attain this global influence and related neuronal activity will only be observable locally within sensory cortices.

Behavioural and neurobiological predictions of the GNW can be broadly summarised in terms of a four-way taxonomy describing the relationship between consciousness, bottom-up sensory signal strength, and attention-based modulation

(Dehaene et al., 2006). Specifically, in the absence of attention the activation caused by the presence of a weak stimulus should remain within early extrastriate areas, leading to weak priming effects (i.e. only slightly above chance) and unavailability for conscious report. When stimulus strength is weak but attention is present, the signal should reach deeper levels of extrastriate cortex – leading to noticeable priming effects (i.e. above chance) yet still unavailable for report. When stimulus strength is increased, but attention is absent, this should allow for deep processing, again facilitating noticeable (e.g., semantic) priming effects but activation should remain limited to sensory areas and be unavailable for report. Finally, when a strong signal has made its way to deep levels of processing and is amplified by top-down attention, prefrontal and parietal loops will be recruited to maintain sensory information through recurrent activity – thereby making it broadly available to large-scale networks subserving domain-general (goal-directed) cognition and allowing for conscious report (among other adaptive uses).

These predictions, while by no means uncontroversial, have been largely corroborated. In a pioneering fMRI study, Dehaene and colleagues found that conscious report of rapidly presented words resulted in the wide-spread activation of prefrontal, temporal and parietal regions, whereas activity remained within sensory regions when the stimulus was rendered invisible via masking (Dehaene et al, 2001). In a meta-contrast masking paradigm combined with electroencephalography (EEG), DelCul et al (2007) found that early event-related potential (ERP) components did not display a significant difference between seen and unseen conditions, while the late P3 component showed a significant non-linear increase in amplitude between seen and unseen conditions. Importantly, objective task performance was well above chance even at the lowest visibility rating. Similarly, using an attentional blink

paradigm, Sergent et al (2005) found that early ERP components either did not vary with visibility, or had a linear relationship with visibility. In contrast, late components, such as the P3, showed non-linear increases in amplitude when T2 visibility was above 50%. Recent work using multivariate decoding in magnetoencephalography (MEG) has extended this finding, showing that, in comparison to temporally adjacent distractor stimuli, consciously reported target stimuli display a stronger and more sustained pattern of activation (Marti & Dehaene, 2017). Most convincingly, when recording from intracranial electrodes implanted in epilepsy patients, Gaillard et al (2009, p.475) found that conscious word perception had a significant effect on many frontal and parietal electrodes, whilst the electrodes showing a significant effect of unconscious perception were almost all located in the occipital and temporal lobes.

Despite the flood of research supporting the GNW, the theory also has a number of limitations. First, it is largely described at the level of gross neuroanatomy, leaving the details of cortical architecture unspecified and thereby limiting the granularity of predictions. Second, it is agnostic about the implementation of expectation, rendering the theory unable to engage with a large body of evidence highlighting the role of expectation in visual consciousness (e.g. Chang et al., 2015; Denison et al., 2011; Valuch & Kulke, 2019; van Gaal et al., 2015).

Of greater concern, as experimental paradigms have become more sophisticated, two predictions of the original theory have been falsified. For example, in a delay matching task, multivariate decoding showed that the brain represented both target presence and target orientation for an entire 800ms delay period across visibility levels. In addition, whilst target visibility correlated with decoding accuracy for target presence, unseen stimuli still exhibited a stable pattern of activation that generalised across time, suggesting that information did not have to be conscious to

enter later phases of processing (King et al., 2016). This corroborates the findings of Salti and colleagues (Salti et al., 2015) who showed that target position could be decoded from superior frontal and superior parietal cortices in both seen and unseen conditions. Together these results demonstrate that, contrary to the original formulation of the GNW, information is processed and unconsciously maintained (at least briefly) by the same structures implicated in conscious processing.

Another seemingly inconsistent result was found by Pitts et al (2014). Using a novel inattention blindness paradigm, they showed that the P3, which was initially thought to discriminate information that did or did not enter the global workspace (see Dehaene, 2014, p.180), was driven by task relevance and not conscious access. This result was recently replicated in a standard masking paradigm (Cohen, et al., 2020) showing unequivocally that the P3 is related to task relevance and is not a necessary signature of conscious access.

To the credit of GNW theorists, the model has been revised to accommodate these findings. First, a revised computational model of subjective report has been proposed, which, in addition to frontoparietal activity, requires that a stimulus representation can be separated from a noise distribution (King & Dehaene, 2014). Second, the claim that the P3 is a specific marker of conscious access (Dehaene et al., 2014) is no longer defended. However, the modified computational model of subjective report (King & Dehaene, 2014) is too idealised to make neurobiological predictions, and the abandonment of the P3 as a signature of conscious access was not accompanied by any theoretical revisions and makes no additional behavioural or neurobiological predictions. This situates the GNW in a tenuous position, in which revision primarily explains away contrary results – a recognized characteristic of degenerative research programs (Lakatos, 1970).

Here we aim to make progress in overcoming these limitations by advancing an extension of the global neuronal workspace - the *predictive* global neuronal workspace (PGNW) - that unifies essential aspects of the GNW with the more recent (Bayesian) Active Inference approach to understanding brain function. Specifically, we present a hierarchical, partially observable Markov decision process (POMDP) model of visual consciousness based upon active inference. Importantly, we leverage the neural process theory associated with active inference to make explicit links between neurobiology and the formal simulations afforded by the model.

Formalising ideas first developed by Hohwy (2013), Whyte (2019), and Friston (2018), we will argue that conscious access or “ignition” is a fundamentally inferential process that depends upon a level of processing of sufficient temporal depth to contextualise and coordinate lower levels of processing – where this longer timescale coordination is seen as necessary for the generation of subjective reports. Here, subjective reports stand in as one example of a broader set of temporally extended action plans (i.e., extended sequences of actions), the generation of which requires integration, maintenance, and manipulation of information over sufficient lengths of time – and where that information is sufficiently complex to guide the controlled generation of such goal-directed behaviours (e.g., combining conceptual contents associated with words such as “I”, “see”, “a”, “red” and “square” requires representing contents of much greater abstraction and temporal depth than is necessary for representing the perceptual property denoted by the word “red”).

To begin, we provide a brief primer on active inference and POMDPs, followed by a specification of the specific structure of the generative model we will employ, paying close attention to the importance of temporal depth. With the groundwork laid out, we show through simulations that the model can both 1) unify

seemingly contradictory previous results, and 2) reproduce the essential aspects of the four-way taxonomy predicted by the GNW, describing the relationship between conscious access, attention, and stimulus strength. Using the same generative model architecture, we then reproduce (and offer mechanistic explanations for) the electrophysiological and behavioural results of the inattentional blindness paradigm introduced by Pitts et al (2014).

Next, we turn to the role of expectation in visual consciousness and show how our model can extend the original four-way taxonomy of GNW theory to encompass paradigms that manipulate prior expectations on trial-by-trial basis – highlighting the novel predictions that emerge from this extension. We also describe one possible novel paradigm that could be used to test distinct model predictions regarding dissociable effects of expectation, attention, and stimulus strength. We end by examining the relationship between the PGNW and alternative models, and briefly address potential concerns about how phenomenal consciousness could plausibly be situated within our model. However, this paper is chiefly concerned with what Block (2005) terms “access consciousness” which is defined as the availability of information for verbal report, action, and executive processing. For brevity, we will use “conscious access” and “consciousness” interchangeably throughout the paper unless otherwise indicated.

## **2. A Primer on Active Inference**

Active inference, a corollary of the free energy principle (FEP), is a first principles approach to modelling (approximately) Bayes optimal behaviour (Friston; 2010; Friston et al, 2016, 2017a). The FEP starts from the tautology that, in order for a self-



organising system to maintain the integrity of its internal milieu, it must stay within the narrow range of states consistent with its survival. Human body temperature, for example, should ideally stay within the range of 36.5 – 37.5 degrees (Celsius). This entails that an organism's phenotype has an attracting set of bodily states. Over long timescales, these attractor states have a high probability of being observed in the sense that the organism will visit them repeatedly (Friston, 2013). Formally then, all self-organising systems must be minimising the (information theoretic) surprise of their sensory observations. However, surprise is computationally intractable. Instead, according to the FEP, organism's construct an internal (generative) model of environmental dynamics that, when accurate, acts as an upper bound on surprise (Buckley et al., 2017). The perception-action cycle is thus cast as an optimisation problem. Perception corresponds to the process of inferring the hidden state values that maximise the likelihood of observations and create a tight bound on surprise, while action is the process of inferring action sequences that either minimise uncertainty about hidden states (epistemic value) or bring about preferred observations (pragmatic value), thereby minimising surprise more directly (Friston et al, 2016, 2017). The former (perception) corresponds to the minimisation of variational free energy  $F$ , while the latter (action selection) corresponds to the minimisation of expected free energy  $G$  (see figure 1 for formal descriptions).

Here we formulate the generative model as a partially observable Markov decision process (POMDP; see figure 1). POMDPs model discrete transitions between latent variables and the observations they generate. Such models infer states and policies based upon the mapping between different hidden state factors and distinct observation (or outcome) modalities – given by a set of  $A$  matrices (one matrix per outcome modality). Transitions between states are determined by the

transition probabilities encoded by a set of **B** matrices (*at least* one matrix per state factor; see description of policy selection below). A set of **C** matrices describes the agent's prior preferences over observations at each time point (one matrix for each outcome modality) and quantifies the degree to which agents prefer, or are averse to, particular observations. Finally, prior beliefs about initial states are determined by a set of **D** vectors (one per hidden state factor). **A**, **B**, **C** and **D** are each categorical distributions.

Such models are equipped with allowable sequences of actions that can be chosen (policies;  $\pi$ ), where each possible sequence is assigned a value (higher policy values relate to lower expected free energies **G**, defined in relation to the prior preferences encoded in **C**). In the context of this class of models, allowable policies are specified as sequences of allowable state transitions, where each allowable transition (action) at each time point is encoded by a distinct **B** matrix for a given state factor. Thus, action corresponds to the agent's direct control of state transitions. Observations and hidden states are factorised into separate outcome *modalities* and hidden state *factors* to allow for interactions between hidden states in the likelihood mapping (**A**). In hierarchical models, such as the model employed in this paper, the hidden states at the first level serve as observations at the second level (see figure 1). Crucially, hierarchical models also allow for inferences about deep temporal structure. An intuitive example of this is reading, in which the first level of a model could infer single words while the second level could infer the narrative meaning entailed by sequences of words over longer spatiotemporal scales (see Friston et al, 2017). Over the timescale of a single trial of a task (for example) belief updates are equivalent to (e.g., perceptual) *inference*, while, over longer

timescales, updating gives rise to *learning* (we refer mathematically interested readers to Da Costa et al., 2020a).

# Partially Observable Markov Decision Process


## Variational free energy

$$\begin{aligned}
 F &= \mathbf{E}_Q[\ln Q(s) - \ln P(s, o)] \\
 &= \underbrace{\mathbf{E}_Q[\ln Q(s) - \ln P(s|o)]}_{\text{Relative entropy}} - \underbrace{\ln P(o)}_{\text{Model evidence}}
 \end{aligned}$$

## Expected free energy

$$\begin{aligned}
 -G(\pi) &= \mathbf{E}_{\bar{Q}}[\ln Q(s|\pi) - \ln P(s, o, \pi)] \\
 &= \underbrace{\mathbf{E}_{\bar{Q}}[\ln Q(s|\pi) - \ln P(s|o, \pi)]}_{\text{Epistemic value}} - \underbrace{\mathbf{E}_{\bar{Q}}[\ln P(o)]}_{\text{Pragmatic value}}
 \end{aligned}$$

## Generative model (states)



## Generative model (matrices)

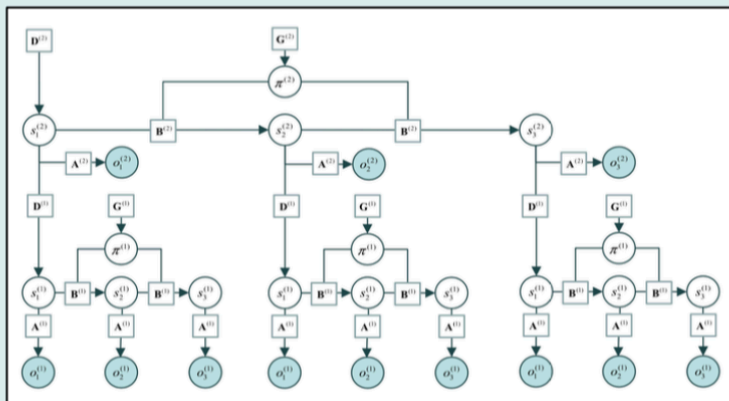


Figure 1. **First.** Decomposition of variational free energy (VFE) into relative entropy and model evidence. Because the relative entropy term is always greater than or equal to zero when the approximate posterior  $Q(s)$  approximates the true posterior  $P(s|o)$ , VFE is equal to the negative model evidence. Minimising VFE is, therefore, equivalent to maximising model evidence. For visual simplicity we have not included the policy term in the VFE, although it should be noted that states are conditioned on policies. **Second.** Decomposition of expected free energy ( $\mathbf{G}$ ) into epistemic value and pragmatic value. Here  $\bar{Q} = P(o|s)Q(s|\pi)$  means that the expectation sums over all current observations. To minimise  $\mathbf{G}$ , agents must select policies that transition them into states that maximise the difference between  $\ln Q(s|\pi)$  and  $\ln P(s|o, \pi)$ . Or, in other words, agents seek out states that reduce uncertainty about expected observations. **Third.** Bayesian network depiction of a 2-level POMDP. Arrows show the dependences between variables. This first network highlights the dependencies between first-level and second-level variables. Observations depend on hidden states at the first level. In turn, hidden states at the first level depend on policies and states at the second level. Importantly, first-level hidden states function as observations for the second level. **Fourth.** Second Bayesian network depiction of a 2-level POMDP. adapted from (Friston et al., 2017b). This representation of the network highlights the role of the matrices in determining the conditional dependencies between variables. Observations are generated by hidden states described by  $\mathbf{A}$ . The  $\mathbf{B}$  matrix determines state transitions and also functions as an empirical prior. The  $\mathbf{D}$  vector serves as the prior for initial states. When the  $\mathbf{B}$  matrix is under the control of the agent, state transitions depend upon the policy. The probability that a particular policy will be selected is dependent on the expected free energy  $\mathbf{G}$  of the policy which is, in turn, partially dependent on prior preferences specified by  $\mathbf{C}$ , although the dependence of  $\mathbf{C}$  is not depicted graphically.

In terms of neurobiological implementation, active inference has a detailed process theory that specifies how a family of possible message passing algorithms can be used to perform inference, as implemented within neurobiologically plausible structural and functional dynamics (Parr & Friston, 2018). Broadly speaking, the firing rates of certain neuronal populations - represented in figure 2 as neuronal populations within cortical columns - encode the current estimate of the posterior

probability over hidden states. The synaptic inputs to the columns carry the conditional probabilities encoded in each of the matrices described above. This means that, for example, activity levels in neuronal populations encoding posterior probabilities over states are updated by ascending signals (from observations) based on synaptic weights encoding the amount of evidence that each possible observation provides for each possible hidden state (i.e., entries within the **A** matrices).

Of particular importance for the purposes of this paper are the equations describing the expected hidden states and the time derivative of the depolarisation variable  $\dot{v}$  (see upper left of figure 2) associated with the neural process theory linked to active inference. Specifically, the posterior expectation over hidden states is a softmax (normalized exponential) function of the depolarisation variable  $v$  which represents the average membrane potential of the neuronal populations responsible for encoding the surprise of expected states. The output of the softmax function is taken as the average firing rate of the population. The use of the softmax function (which is simply a generalisation of the sigmoid function to vector inputs) to simulate average firing rate is based on the assumption made in mean-field models of neural dynamics that the average firing rate of a population can be treated as a sigmoid function of the average membrane potential (see Da Costa et al, 2020b). ERPs and local field potentials are modelled as the temporal derivative (rate of change) of this depolarisation variable (see figure 2; Friston et al., 2017; Friston et al., 2018; Parr, & Friston, 2017a).

## Process Theory

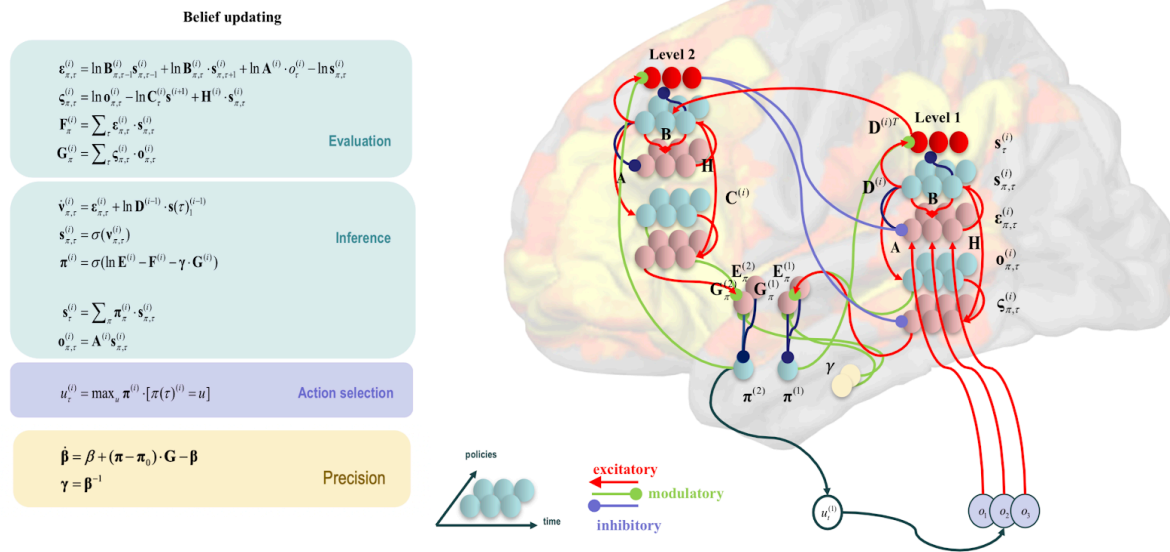


Figure 2. Neural network implementation of the message passing implied by the Bayesian network in figure 1. **Top left.** Update equations and free energy functionals. Heuristically, state prediction errors  $\epsilon$  score the evidence that observations provide for each policy (i.e., the difference between outcomes expected under each policy and those that are subsequently observed). In contrast, outcome prediction errors  $\zeta$  encode beliefs about the value of each policy (i.e., higher outcome prediction errors for a given policy roughly correspond to lower probabilities of observing preferred outcomes under that policy, as well as less informative observations expected under that policy). Directly below are expressions for variational free energy  $\mathbf{F}$  and expected free energy  $\mathbf{G}$ , expressed in terms of the above mentioned (state and outcome) prediction errors. **Middle left.** Update equations for states, policies, Bayesian model averages and the depolarisation variable. **Bottom left.** Action selection and update terms for expected precision over policies. **Right.** Schematic of message passing between cell populations. Red units encode Bayesian model averages, cyan units encode expectations over hidden states, and pink units encode state and outcome prediction errors.

### 3. A Deep Temporal Model of Visual Consciousness

To model the difference between conscious and unconscious perception, we based our simulated task on the paradigm introduced by Pitts and colleagues (2012, 2014a,

2014b). We chose this task because, with only minor changes in design, the paradigm can be used to study both inattentional blindness and phenomenal masking – allowing us to model the interaction between attention and sensory signal strength in an empirically plausible manner (it also afforded the opportunity to reproduce, and provide explanations for, empirical results obtained from this paradigm; see below).

At the beginning of each trial in our simulated task, the in silico subject - or agent - was presented with a stimulus composed of an array of bars surrounded by coloured disks. At the 2nd time point, the array of bars was replaced by a square, and at the 3rd time point the array changed back to the collection of bars. The agent was then required to self-report whether or not they had seen the square or to perform a two-alternative forced-choice task. We manipulated attention by requiring the agent to monitor the colour of the surrounding circles (either red or black) at the expense of the inner array (see figure 3).

### Task Structure

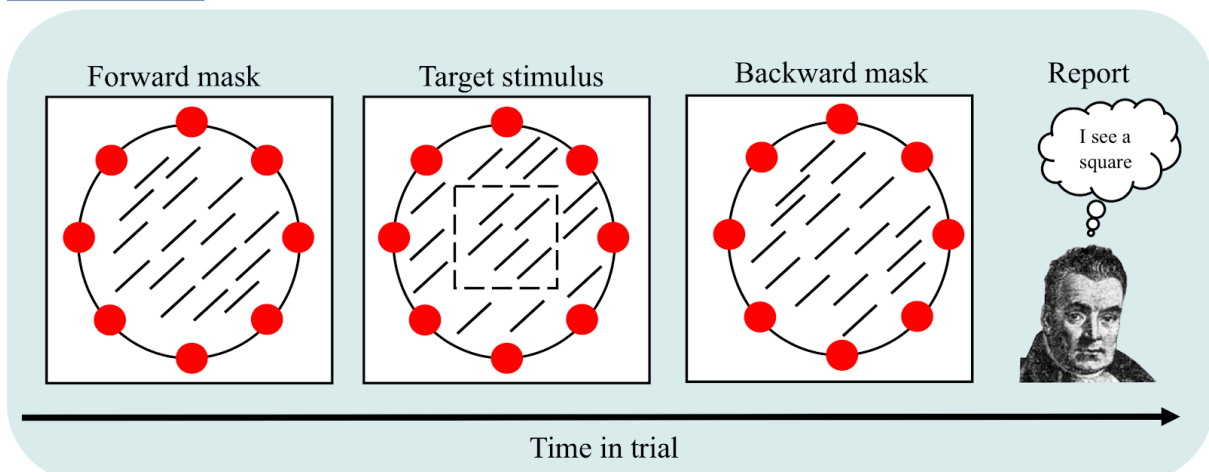


Figure 3. Illustration of the task performed by the agent. On each trial the in silico subject was presented with a stimulus composed of an array of bars surrounded by coloured discs. At the 2nd time point, the array was replaced by a square, and at the 3rd time point the array changed back to the



original bar pattern. The agent was then required to either perform a two-alternative forced-choice task or report whether they had seen the square.

To simulate the perceptual categorisation and self-report behaviour described above, we specified the task in terms of a generative model which was inverted using a variational message passing scheme (for technical details see Parr et al., 2019a). To reproduce the recurrent interactions between the frontoparietal network identified with the global workspace, and the input it receives from the visual system, we employed a two-level deep temporal model (Friston et al., 2017b; Friston et al., 2018). The first level (see figure 4), which roughly corresponds to processes within the visual system, had four hidden state factors; attention allocation, internal stimulus (bars/square), surrounding or external stimulus (red/black circles), and a set of auditory-verbal states (a number of words that could be put together in different sequences to generate verbal reports). The second, higher level, which corresponds to the frontoparietal network associated with the global workspace, included three hidden state factors: 1) sequence type (encoding beliefs about the sequence of internal and external stimuli presented on each trial), 2) time point within the trial which, in line with data from non-human primates (Kapoor et al., 2018), encodes the current phase of the task, and 3) abstract semantic representations that could be unpacked into different verbally reported sequences of words at the lower-level (dependent on the chosen policy at the higher level).

It is worth emphasising that the temporal depth of this second level is essential for simulating the self-report behaviour that defines conscious access. The language component of the model is obviously an oversimplified depiction of linguistic cognition. Yet, it remains true that, in order to coordinate the selection of a specific sequence of words (i.e., to construct a sentence describing the content of

perception), where each words is generated over a more rapid time scale, the agent *must* have a level of processing that unfolds over a slow enough time scale to abstract away from the moment by moment sensory flux and coordinate lower-level language processes to report the outcome of perceptual decision-making in a goal- (i.e., preference-) dependant manner. Thus, the higher level is necessary to integrate the moment-to-moment sensory flux with the longer-timescale, controlled processes capable of generating behaviour plans with greater temporal depth. This feature is at the core of our model and we will revisit it in more depth in the discussion.

### Generative Model

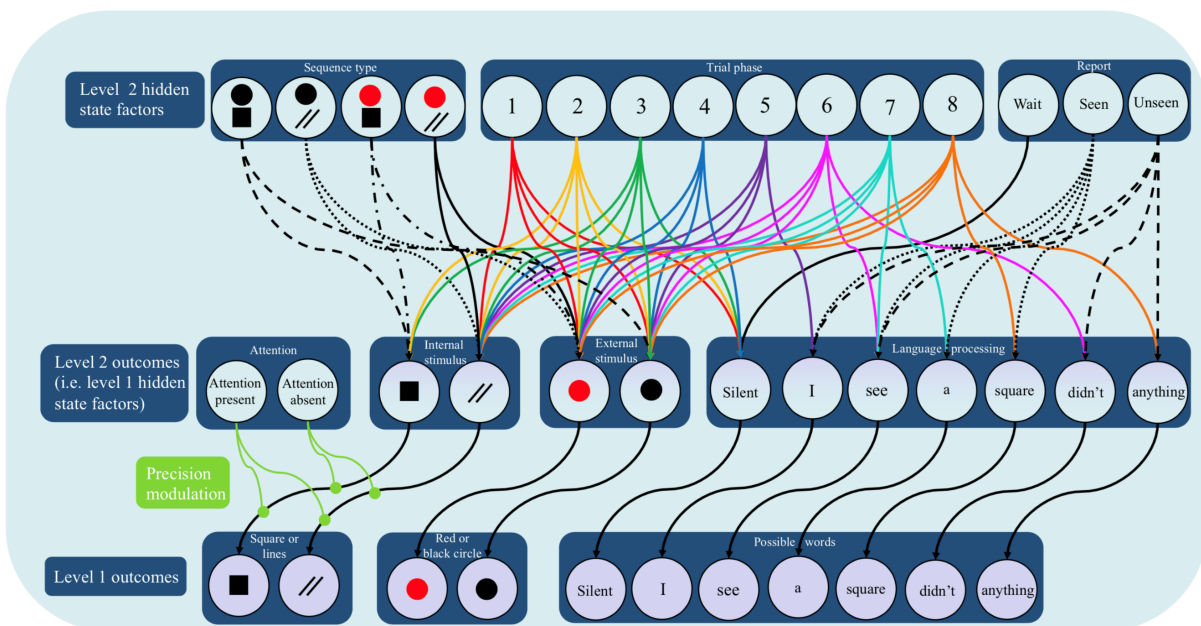


Figure 4. Bayesian network depiction of the generative model, with arrows showing the dependencies between hidden state factors and outcome modalities. At the second level, states within the sequence type and trial phase hidden state factors determine the internal stimulus and external stimulus hidden states at the first level (which function as second-level observations). This was set up such that the mapping between second-level hidden states and first-level hidden states was dependent on the time point within the trial. The report state factor determined the language processing state at the first level. In turn, first-level outcomes were dependent on first-level hidden states. Importantly, there was

only a precise mapping between the first-level internal (line/square) stimulus hidden state factor and the corresponding outcome modality when the model was in an attentive state.

At the first level of the model (see figure 5), the **D** vectors specified the initial state of the four hidden state factors; top-down attention (present, absent), internal stimulus (bars/square), external stimulus (red/black circles), and auditory-verbal states (single words: “silent”, “I”, “see”, “a”, “square”, “didn’t”, “anything”). The state transitions specified by the **B** matrices were all identity matrices, meaning that the hidden states were stable across the course of each trial. The likelihood mapping between the hidden states and outcomes, specified by the **A** matrices, is where we implemented the attention and signal strength manipulations. The external stimulus and language matrices were both fully precise (identity matrices). In contrast, we reduced the precision (denoted by  $\zeta$  for stimulus strength and  $\Sigma$  for attention) of the mapping between the internal state and the outcomes by passing what were initially identity matrices through two softmax functions controlled by precision parameters representing the effects attention and signal strength (i.e. presentation time). Higher values of these parameters made the **A** matrices more precise. We set up the interaction between the **A** matrices such that the likelihood mapping for the internal stimulus factor was more precise when the agent was in an attentive state. In contrast, stimulus strength manipulations reduced the precision of the mapping between stimuli and hidden states independent of attentional state.

## Level 1 Matrix Structure

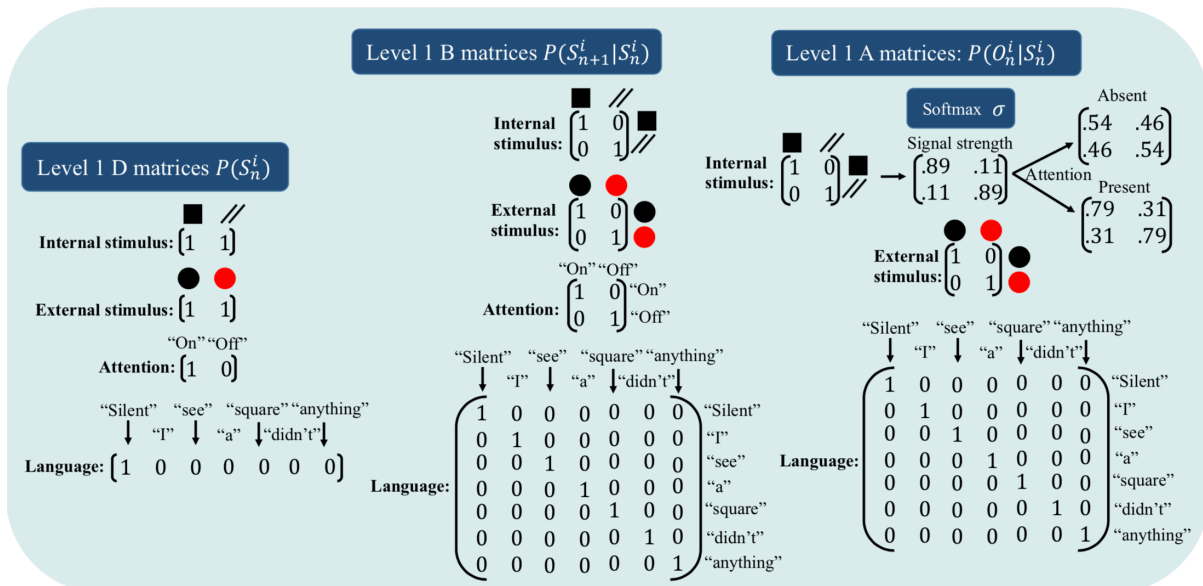


Figure 5. 1st-level generative model matrices. All matrices are passed through softmax functions such that the columns of each matrix, and the rows of each vector, always sum to one. Each column of the **D** vector represents the probability of a hidden state. Columns of the **B** matrix correspond to states at time  $t$  and rows correspond to states at  $t+1$ . Here all the **B** matrices are identity matrices meaning that the states were believed to be stable across each trial. Columns of the **A** matrix correspond to hidden states while rows correspond to observations. An identity mapping therefore implies a deterministic likelihood mapping between states and observations. To model the effects of attention and stimulus strength the **A** matrix encoding the likelihood mapping for the internal segment of the stimulus was passed through a softmax function twice with a precision multiplier representing attention ( $\Sigma$ ) which could be present or absent, and stimulus strength which could be strong or weak ( $\zeta$ ). The combined effect of the attention and signal strength multipliers determine the final precision of the matrix as depicted above. With the exception of the **A** matrix encoding the likelihood mapping for the internal segment of the stimulus, all other first-level **A** matrices were identity matrices, meaning that the mapping was deterministic. Finally, it is important to note that, for visual simplicity, the matrices displayed above have not been factorised and appear differently than how they are implemented in the code.

At the second level (see figure 6) the three hidden state factors specified by the **D** vectors were: sequence type (black disk and square, black disk and bars, red

disk and square, red disk and bars), time point within trial (1-8), and report state (wait, seen, unseen). We set the initial level of the report state to “wait”. The **B** matrix for sequence type was an identity matrix, meaning that the agent believed a priori that the sequence type would not change mid-trial. Trial phase was set up such that time point 1 transitioned to time point 2, which transitioned to time point 3 and so on until the end of the trial. For time points 1-4, all states in the “report” **B** matrix mapped to “wait”; however, at time point 5 the agent had control over the **B** matrices for the report state, meaning that it could transition to either a “seen” or “unseen” state depending on which policy best minimised expected free energy. The **A** matrices were factorised such that the mapping from hidden states to outcomes was dependent on the time point in the trial (see the time-in-trial hidden state factor in figure 4). At time point 1 both square sequence and bar sequence hidden states predicted a bar outcome (recall that second-level outcomes are also first-level hidden states). While at the 2nd and 3rd time points the square sequence predicted a square outcome and a bar hidden state predicted a bar outcome. At time point 4, both the square sequence and bar sequence once again predicted a bar outcome. To model the recurrent feedback between hierarchical levels characteristic of ignition, the square sequences mapped to the square outcomes for time point 2 and 3 – allowing the state at the second level to influence belief updating at the first level via the second-level **A** matrix. From time points 1-4 all the states in the “report” factor mapped to the “silent” first-level verbal state. However, from time point 5 on, “seen” and “unseen” states entailed a different sequence of lower-level word representations. The “seen” report state entailed the words (in order) “I” “see” “a” “square” at each successive time point, while the “unseen” report state entailed the words “I” “didn’t” “see” “anything” in that order. The report state also had a likelihood

mapping to a feedback outcome. The agent was “correct” if, at time point 8, they reported “seen” after a square sequence or “unseen” after a bar sequence, and “incorrect” if they reported “unseen” after a square sequence or “seen” after a bar sequence (this preferred feedback was used to motivate honest verbal reporting policies; see below). Finally, to account for the diffuse nature of feedback projections (Bannister, 2005; Garcia-Cabezas et al., 2019), we lowered the precision of the **A** matrix for the “sequence type” factor (precision = 0.8) providing a plausible threshold on ignition events.

### Level 2 Matrix Structure

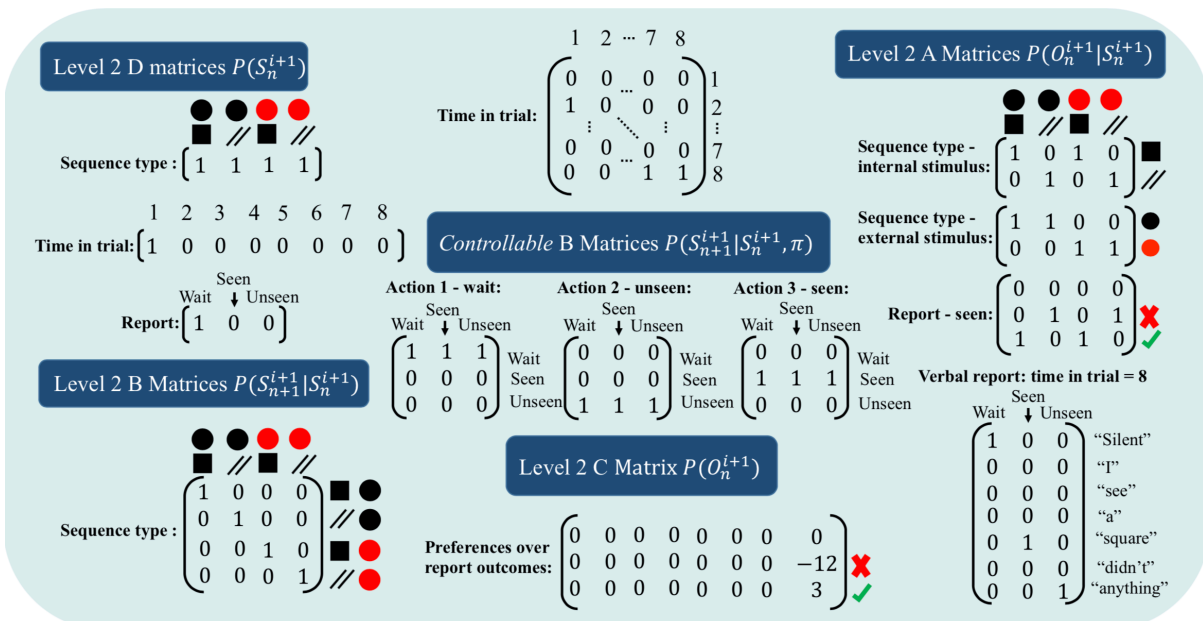


Figure 6. 2nd-level generative model matrices. **D** vectors for sequence type (unless otherwise indicated) assigned equal probability to each state. The **D** vector for the trial phase hidden state factor was initialised so that the model would always start each trial with full confidence that it was in the state corresponding to time point 1. Similarly, the **D** vector for the report state was initialised so that the model had full confidence in the “wait” state at the start of each trial. The **B** matrix for sequence type was an identity matrix meaning that the agent believed a priori that the sequence of states would be stable throughout the trial. We set up the **B** matrix for the trial phase hidden state factor such that each state successively transitioned to the next state (i.e. state 1 transitioned to state 2 and so on).

For the controllable **B** matrices there was one matrix for each possible report state (wait, unseen, and seen). From time step 1 - 4 the agent could only select the “wait” matrix but at time step 5 the agent could choose (i.e. via policy selection) between the “seen” and unseen “matrices”. The **C** matrix encoded the agents preference for each outcome and had a column for each time point and a row for each action (report state). There was one **C** matrix for each outcome modality. Here we only display the **C** matrix for the “correct/incorrect feedback” outcome modality associated with the report state, as it is the only outcome for which the agent had non-zero preferences. That is, the agent preferred to be “correct” at the end of each trial rather than “incorrect”. Finally, the **A** matrices were set up such that the sequence type hidden state factor had two corresponding outcome modalities, which mapped the sequence type hidden states to the internal stimulus and external stimulus hidden states at level 1. To provide the model with a plausible “ignition threshold,” we lowered the precision of the **A** matrix for the “sequence type” factor by passing it through a softmax function (precision = 0.8) although we do not picture this graphically. The report hidden state factor did not map to hidden states at the level below, instead it mapped to observations that informed the agent about whether they were “correct” or “incorrect” (recall that, because of the **C** matrix, the agent wanted to receive “correct” observations and was averse to “incorrect” observations). Finally, the report hidden state factor mapped to first level language processing hidden states so that after time step 5, once the model was in a “seen” or “unseen” state, the appropriate sequence of spoken word states would be initiated at the level below (i.e. “I” “see” “a” “square”). Again, for visual simplicity the matrices displayed above have not been factorised and appear differently to how they are implemented in code.

We constructed the **C** matrix so that when the agent received feedback at time point 8 they most preferred to be correct and least preferred to be incorrect when reporting whether or not they had seen the stimulus (preference values that produced sufficient motivation for accurate reporting are depicted in figure 6). To model forced-choice behaviour, we ran a separate simulation but reduced the preference for being correct versus incorrect, making the agent less conservative and more likely to guess under conditions of weaker perceptual signals. Because

active inference models are deterministic, we set policy precision (the confidence in policy selection denoted by  $\beta$ ), and motor stochasticity (randomness of action selection denoted by  $\alpha$ ) to  $\beta = 2$  and  $\alpha = 6$ , thereby allowing for a plausible level of behavioural variability reflecting the agent's relative confidence in some states over others.

We made the modelling decision to have a more liberal threshold for forced-choice behaviour than for subjective reports based on the well-replicated finding that subject's display above-chance performance even in the absence of reportability (see discussion in King & Dehaene, 2014). However, we acknowledge that this finding is largely dependent on the method of report. Specifically, there is evidence suggesting that humans have optimal introspective access to their perceptual processes in the sense that betting performance in a 2-interval forced-choice task matches that of an ideal Bayesian observer (Peters & Lau, 2015). We do not wish to take a stand on this issue here as we consider it an open empirical question. Instead, we merely note that the decision was a pragmatic one based on the method of report used in the paradigms we were aiming to simulate.

## **4. Simulating Visual Consciousness**

### **4.1 Foundational Simulations**

As a proof of principle, we simulated 200 trials, 100 of which were “square-present” trials. The model was in an attentive state ( $\Sigma = 0.5$ ) and the stimulus also had a high stimulus strength ( $\zeta = 1.5$ ), corresponding to a long presentation time that enabled evidence accumulation resulting in a relatively precise hidden state-to-outcome



mapping. We found that the model performed with 100% accuracy on the forced-choice task and when reporting the presence – and absence – of the square.

Having established the face validity of the model, we now turn to the simulation of minimal contrast paradigms.

## Foundational Simulation

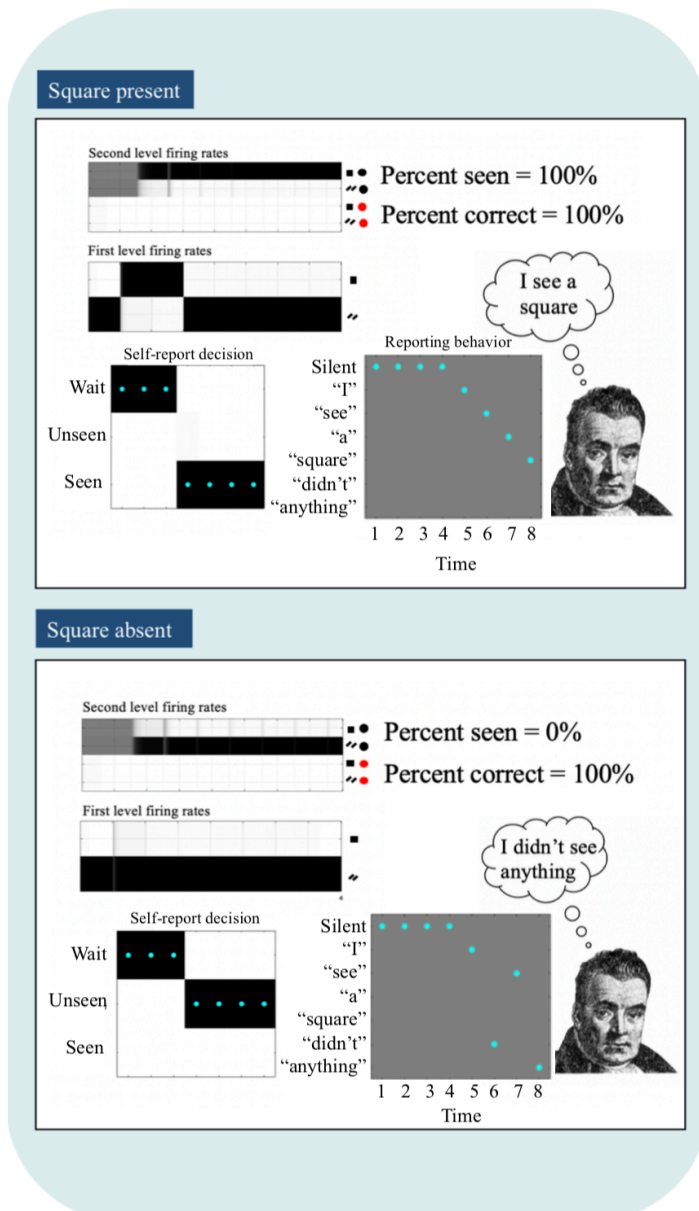


Figure 7. Simulated firing rates (darker = higher firing rate) predicted under the process theory associated with active inference (Friston et al., 2017). Each row represents the firing rate of the

neuronal populations encoding the posterior expectations over each state. Individual squares each represent the time point within the trial. Actions (cyan dots = true action chosen; colour represents the posterior confidence in the actions chosen at the end of each trial, darker = higher confidence). Self-report outcomes (cyan dots = actual observations; colour indicates the preference for each observation, darker = greater preference, here the colours are all the same which means the agent did not have a preference for language outcomes). **Top.** Simulated results for 100 “square-present” trials. Notice the outcome modality associated with the “seen” report state generates the sentence “I see a square”. **Bottom.** Simulated results for 100 “square-absent” trials. Notice that the outcomes associated with the “unseen” report state generate the sentence “I didn’t see anything”.

## 4.2 A Four-Way Taxonomy of the Factors Underlying Conscious Access

As a field, the neuroscience of consciousness has converged on the use of minimal contrast paradigms that, through masking, inattention, or near-threshold presentation, render nearly identical stimuli unconscious in one condition and conscious in the other. Combined with neuroimaging, this allows for the contrast of conscious and non-conscious forms of visual processing. However, early neuroimaging research reported conflicting results. Some studies found that subjective reports correlated with activation in early visual cortices (Zeki, 2003), while others found that subjective reports correlated with the activation of frontoparietal areas (for a meta-analysis see Bisenius et al., 2015). Still others argued that frontoparietal involvement is due to an attentional confound and does not reflect conscious access (Peter et al., 2005). Based upon simulations of thalamo-cortical networks (Dehaene et al., 2003), Dehaene and colleagues (Dehaene et al., 2006) created a taxonomy of factors underlying conscious access with the aim of unifying the conflicting results under the theoretical framework of the GNW.

In their exposition of the taxonomy, Dehaene and colleagues (Dehaene et al., 2006) distinguished between subliminal, preconscious and conscious forms of processing. However, for the sake of clarity, we describe the taxonomy in terms of the two factors underlying the classifications; attention, which can be present or absent, and signal (i.e. stimulus) strength, which can be strong or weak. As described in the introduction, the activation caused by a weak stimulus in the absence of attention (e.g. an unattended masked stimulus) should remain within early extrastriate regions – remaining unavailable for report and only causing weak priming effects. When attention is present, but stimulus strength is weak (an attended but masked stimulus), the signal should reach deeper levels of extrastriate cortex – remaining unavailable for report but leading to stronger priming effects. When attention is absent, but stimulus strength is high (i.e. during inattentional blindness or motion-induced blindness), there should be deep levels of processing within sensory areas – facilitating priming effects but remaining unavailable for report. Lastly, when a strong signal reaches a deep level of processing and is amplified by top-down attention, recurrent loops in frontoparietal cortices will maintain the information over longer timescale and make it available to inform the generation of verbal reports.

Crucially, different neural correlates are observed when using specific paradigms, such as the attentional blink and phenomenal masking, to contrast different parts of the taxonomy. Dehaene et al (2006) have leveraged these findings to explain a number of seemingly contradictory results.

While the theoretical backbone of the taxonomy has been revised, in that the activation of frontoparietal regions is no longer considered sufficient for conscious

access, it is still a useful starting place – as the interactions between stimulus strength and attention described by the taxonomy encompass many, if not most, of the minimal contrast paradigms reported in the literature.

As described above, we modelled the effects of both attention and signal strength on simulated behaviour by altering the precision of the state-outcome mapping between the internal stimulus hidden state and internal stimulus outcome at the first level. In terms of the task, the “low signal strength ( $\zeta = 0.1$ ) + attention absent ( $\Sigma = 0.1$ )” condition corresponds to a short presentation time, with the agent attending to the external coloured disk at the expense of the internal stimulus. The “low signal strength ( $\zeta = 0.1$ ) + attention present ( $\Sigma = 0.5$ )” condition corresponds to a short presentation time with the agent directing attention to the internal stimulus. The “high signal strength ( $\zeta = 0.65$ ) + attention absent ( $\Sigma = 0.1$ )” condition corresponds to a long presentation time with the agent attending to the external coloured disk. Finally, the “high signal strength ( $\zeta = 0.65$ ) + attention present ( $\Sigma = 0.5$ )” condition corresponds to a long presentation time with the agent attending to the internal stimulus. We settled on the specific parameters shown above by searching the parameter space to find consistent values that best reproduced the behavioural results reported in the empirical literature, while remaining within plausible limits (i.e., as when fitting model parameters to real participant behaviour in empirical studies).

#### **4.2.1 Four-Way Taxonomy: Simulated Behaviour**

We presented the model with 400 “square present” trials, 100 corresponding to each of the four parameter variants described above. When signal strength was low and attention absent, the agent displayed chance levels of forced-choice performance (51%) and did not report having seen the square on any of the trials. When signal strength increased, or when the agent attended to the stimulus, forced choice performance improved to well above chance, 70% and 72% respectively, while still not reporting ‘seen’ on more than 20% of trials (8% and 16% respectively). When stimulus strength was strong, and attention was present, the agent showed near-ceiling levels of forced-choice performance (99%; i.e., due to the small amounts of stochasticity in choice within the model) and accurately reported having seen the square on 100% of the trials. Thus, under a consistent set of parameter settings, the model accurately reproduced the self-report and forced-choice behaviour commonly reported in the literature (see figure 8). This should not be a surprise, as we fine-tuned the parameter values to capture behaviour reported in the empirical literature. However, it is worth noting that it is possible a priori that no consistent set of parameter values could be found to reproduce all of these known empirical results – which would have shown a clear insufficiency of the model. Thus, the existence of this consistent set of parameter values does support the validity of the model.

## Four-way Taxonomy

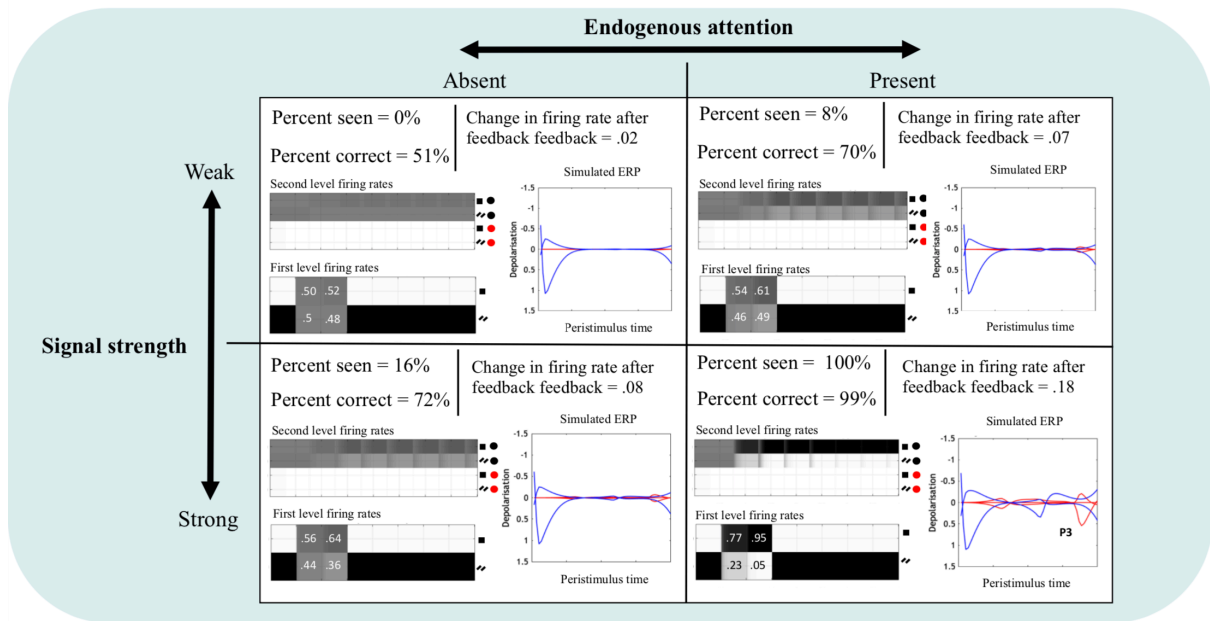


figure 8. Report frequencies, forced choice accuracy, simulated firing rates, and simulated ERPs for each quadrant of the four-way taxonomy. Numbers shown on the lower-level firing rate plots illustrate the firing rate strengths (between 0 and 1) before and after top-down feedback from the higher level (i.e., time steps 2 and 3, respectively). Blue lines on the ERP plots show the temporal derivative of the first-level firing rates, while the red lines show the temporal derivative of second-level firing rates.

### 4.2.2 Four-Way Taxonomy: Simulated Firing Rates

We next simulated the neural firing rates associated with the different quadrants of the taxonomy. The simulated first- and second-level firing rates both accurately reproduced a number of otherwise disparate empirical findings (see figure 8, left portion of each quadrant). Specifically, in all but the “weak signal + attention absent” condition, the model showed an amplification of the first-level firing rate after receiving feedback from the second level. In line with the notion of ignition, this top-down amplification was notably greater when the model reported having seen the square. Similarly, at the second level of the model, in all but “weak signal + attention

absent” condition the firing rates corresponding to the neuronal population encoding the “square present” sequence was amplified by the presentation of the square, even in the absence of report. Again, in line with the notion of ignition, the firing rate at the second level was further enhanced when the model reported having seen the square on all trials.

These results mirror the multivariate decoding results reviewed in the introduction, which found that, although increased visibility correlated with increased decoding accuracy, target orientation could be decoded across visibility levels from superior frontal and superior parietal cortices (Salti et al., 2015). What leads the agent to report the presence of the square is not simply that the second level has greater firing rates for the “square present” sequence, but that the model is sufficiently more confident about the “square present” sequence than the “random bar” (i.e. “square absent”) sequence. This feature of the model is similar to other Bayesian and hierarchical models in the literature that associate subjective report with the inference that a stimulus distribution is distinct from a noise distribution (King & Dehaene, 2014; Peters and Lau, 2015; Fleming 2019). However, the threshold for report in the present model is set by the expected free energy functional, which is composed of both an epistemic component (i.e., related to the difference between posterior expectations over states with and without an observation expected under each policy) and a pragmatic component (i.e., related to the deviation of expected and preferred observations under each policy; see Friston et al., 2017). The inclusion of the pragmatic component allows the PGNW to, in principle, explain the observed role of reward in modulating the content of visual consciousness (e.g. Marx & Einhäuser, 2015; Dong et al., 2019).

Importantly, if we were to apply a contrastive methodology and compare any of the three unconscious sections of the taxonomy with the conscious section, we would find that conscious access is associated with an enhanced firing rate at the first and second levels of the model. This mirrors the finding that conscious report correlates with the enhanced activation of sensory regions and the seemingly all-or-nothing activation of frontoparietal cortices originally cited in support of the GNW taxonomy (e.g. Dehaene et al, 2001). Based on our model, the (seemingly) contradictory finding that superior frontal and superior parietal cortices contain decodable information across visibility levels (Salti et al., 2015) can be explained as the result of using different methods to analyse the same underlying active inference architecture. that is, the use of standard univariate analysis on the one hand, which looks for voxels that are more activated by visible stimuli than invisible stimuli (e.g. Dehaene et al., 2001), and the more sensitive multivariate pattern analysis on the other, which instead looks for patterns of information present across voxels. As was noted above, using a standard univariate approach to analyse the simulated firing rates shown in figure 8 would show that firing rate is enhanced on conscious compared to unconscious trials at both the first and second level. However, this ignores the presence of stimulus relevant information in the firing rates; although firing rates are higher for conscious compared to unconscious conditions, there is still a greater than baseline firing rate in unconscious conditions at both the first- and second-level within our simulations that could easily be exploited by a classifier. Thus, our model can simulate and explain both (seemingly conflicting) results within a single neurocomputational architecture.



### 4.2.3 Four-Way Taxonomy: Simulated Event-Related Potentials

Finally, we examined the ERPs predicted by our model under different quadrants of the taxonomy (figure 8, right portion of each quadrant). Here we found that ERPs at the first level of the model – which correspond to early components such as the P100 and N100 – are relatively unaffected by changes in visibility. In contrast, ERPs at the second level – which correspond to late components, and specifically the P3 – appear to be strongly modulated by changes in visibility. This emulates the empirical findings described in the introduction, in which early components were relatively unaffected by changes in visibility while late components displayed a non-linear increase in amplitude as visibility increased (DeICul et al., 2007; Sergent et al., 2005).

As described above, ERPs are here modelled as the time derivative of the depolarisation variable. When confidence in a state changes rapidly, the derivative of the depolarisation variable increases in magnitude. This gives us a new perspective on late ERP components in the context of visual consciousness. Since the model only reports the presence of the square when it is sufficiently more confident about the “square present” sequence than the “random bar” sequence, it makes sense that verbal reports are accompanied by late ERPs – as they reflect the rapid change in beliefs at the second level of the model when the square is presented.

However, it is important to highlight that this is not a necessary feature of conscious access. If the model is already sufficiently confident in a state – either because of a precise likelihood mapping or because *the rate of belief updating is slowed* (by a lack of precision in the likelihood; e.g., because of inattention) – the model may report the presence of a stimulus without an accompanying late ERP. This is indeed what we see empirically. Specifically, in the inattention blindness

paradigm that guided development of the particular task structure of our model, Pitts and colleagues (2012, 2014a, 2014b) found that the P3 is associated with task relevance (i.e. attentional set) and *not reportability*, which was interpreted as evidence against the standard model of the GNW. In other words, the amplitude of the P3 changed as a function of attentional set – which affects precision – rather than conscious access as such. In the next section we discuss in more detail how our model can account for this result and demonstrate how the model can shed light on the specific electrophysiological correlates of inattention blindness more generally.

### **4.3 Inattention Blindness**

The simulations reported in the previous section were focused on reproducing the findings that characterise minimal contrast paradigms associated with the GNW theory's proposed four-way taxonomy. However, minimal contrast paradigms are not without limitations. For example, as Aru and colleagues (2012) have argued, these paradigms often confound the neural correlates of consciousness with the prerequisites to, and consequences of, consciousness. With these confounds in mind, in the study referred to above Pitts and colleagues (2014a, 2014b) created a three-phase sustained inattention blindness paradigm that was designed to dissociate the electrophysiological correlates of consciousness from the correlates of task relevance.

The task used the same stimulus set discussed in the previous two sections. In phase one, participants were instructed to monitor the external disks for a change in colour. Every 600-800ms the internal section of the stimulus alternated between random bars and a square, both of which were presented for 300ms. After phase

one, participants completed a debrief and 50% of them reported not being aware of the square, replicating the findings of Mack and Rock (1998). The debrief acted as a cue alerting the participants to the presence of the stimulus. Phase two was identical to phase one except that participants had been alerted to the presence of the square. Despite not being task relevant, 100% of participants reported having seen a square in the subsequent debrief. In phase three, participants were instructed to attend to the inside stimulus while ignoring the external disks. Each phase consisted of 600 trials, 240 of which were square trials. Before the beginning of the experiment, participants performed 300 practice trials where they were asked to perform the same task as in phases one and two. Crucially, during the training session only random bar segments were shown, *establishing a strong prior expectation for random bars* (we return to this important factor below).

Contra the predictions of the GNW model, the ERP results revealed a dissociation between awareness and the P3. As expected, in phase one the P3 was absent for inattentionally blind participants. However, the P3 was also absent in phase two when all participants were conscious of the stimulus despite it not being task relevant. In contrast, when the stimulus was conscious and task relevant (in phase three), there was a large P3, showing that task relevance is the primary driver of this ERP component.

Our aim in simulating this task here is to provide a bridge between established experimental findings and the more abstract account of the relationship between the P3 and visual consciousness advanced in the previous section.

Given the training session that preceded phase one of the experiment, and the equivalent attentional demands of the distractor task in phase one and phase

two, it seems likely that expectation, in addition to attention, played a crucial role in the above results. Therefore, we modelled phase one by increasing the prior probability of the random bar segment sequence in the second level **D** vector so that the model was twice as confident a priori that it would only be presented with random bars compared to squares (we also verified through simulation that these values can be learnt upon repeated presentations of random bar sequences). We then found the value of the stimulus strength parameter ( $\zeta = 2.12$ ) that led the model to report the presence of the square ~50% of the time while in an attention-absent state ( $\Sigma = 0.1$ ; i.e., to match the empirical setup).

To model the effect of the debrief that alerted the participants to the presence of the square, in phase two we simply increased the prior probability of the “square-present” sequence so that the model was equally confident a priori that it would be presented with squares versus random bars. Phase three was the same as phase two, except that the model was in an attentive state toward the possibility of seeing the square.

### **4.3.1 Inattentional Blindness: Simulated Behaviour and Event Related Potentials**

We presented the model with 300 square trials, 100 corresponding to each of the three sets of parameter values described above. Through stimulus strength calibration, in phase one the model reported the presence of the square 49% of the time, mirroring the empirical results. Crucially, however, the electrophysiological results of all three phases, and the behavioural results of phase two and phase three, reproduced the empirical findings without any further adjustments to

parameter values. In phase one, there were only early first-level ERPs (see figure 9). In phase two, when the “square present” and “random bar” sequences had equal prior probability, the model reported the presence of the square on 99% of trials and, in line with empirical results, there were again only large first-level ERPs; no late P3-like ERPs were generated. Finally, in phase three, when the model was attending to the internal stimulus, the square was reported on 100% of trials and there was a large and late ERP at the second level resembling the P3. Considering the idealised nature of the model, the simulated ERPs displayed in figure 9 bear a striking resemblance to the empirically observed ERPs.

### Sustained Inattentional Blindness

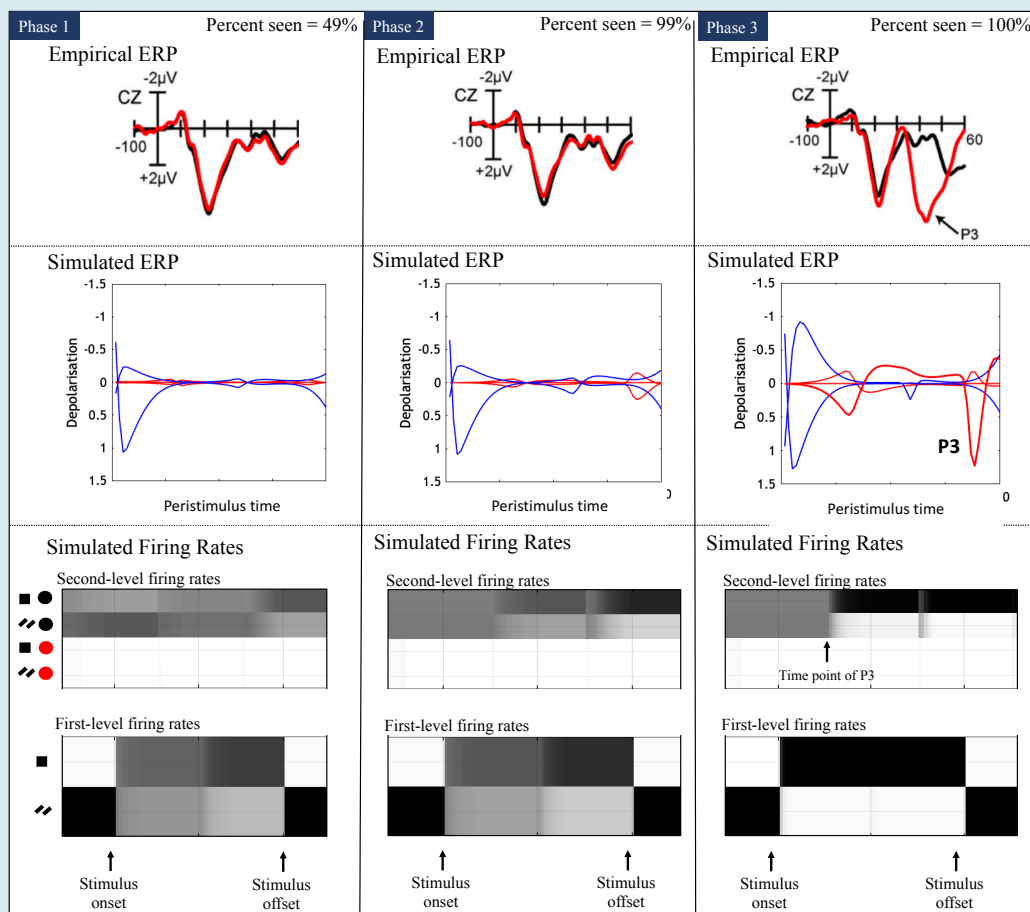


Figure 9. Report frequency, simulated ERPs, and associated firing rates predicted for each of the three phases of the sustained inattentive blindness task. The empirical ERP plots, taken from the original study conducted by Pitts et al (2014), show the amplitude for “random bar” (black) versus “square present” (red) trials at electrode CZ. Firing rate plots illustrate how more gradual updates in 2nd - level beliefs do not produce P3-like ERPs (middle panel, Phase 2) despite self-reported experience of the square, while a faster rate of change in 2nd - level beliefs does produce P3-like ERPs (right panel, Phase 3) in addition to identical self-reported experience, Note that in Phase 1, 2nd - level firing rates initially favour the lines over the square, due to prior expectations learned during the previous training phase.

This result is noteworthy. The dissociation of the P3 and visual consciousness emerges naturally out of the belief updating scheme that underwrites the PGNW. Further, it leads to a straightforward prediction; visual consciousness will be accompanied by a late ERP whenever confidence in a particular state at a high level of the hierarchy changes rapidly.

Having shown that our model reproduces the minimal contrast results cited in support of the original formulation of the GNW, explains away otherwise contradictory results in the minimal contrast literature, and accounts for the dissociation of the P3 and visual awareness, we now turn to the role of visual expectation – and illustrate how our model offers an extension of the previously introduced four-way taxonomy of factors underlying conscious report. In addition to attention and stimulus strength, our model introduces a third factor: trial-by-trial changes in prior visual expectation.

#### **4.4 Extended Taxonomy: Expectation and Visual Consciousness**

There is now a large body of evidence showing that expectation plays a fundamental role in determining the content of visual consciousness. Specifically, under conditions of continuous flash suppression predictive cues accelerate the entry of a suppressed stimulus into consciousness (van Gaal et al., 2015). In binocular rivalry paradigms, predictive context increases the dominance of stimuli congruent with that context (Denison et al., 2011; Valuch & Kulke, 2019). Cross-modal predictions also accelerate the re-entry of stimuli into consciousness after a period of motion induced blindness (Chang et al., 2015). In the absence of attention, expectations reliably induce illusory perception of absent stimuli (Aru et al., 2018). And, when viewing ambiguous figures, expectations have been shown to bias the perceived direction of rotation. (Sterzer et al., 2008). In light of this, if visual consciousness is to be fully understood, it appears essential to extend the taxonomy of factors underlying conscious access to include expectation.

Expectation is typically manipulated by either exposing subjects to repeated instances of a stimulus, or by providing an explicit cue that informs subjects about the conditional probability that a stimulus will appear. Within the active inference framework manipulating expectation via a history of exposure corresponds to learning, while the use of explicit cues corresponds to inference. Trial-by-trial inference is reflected in the firing rates of neuronal populations encoding posterior beliefs over states and policies, while learning is the result of extended periods of belief updating and occurs over a longer time-scale corresponding to the modification of synaptic weights (modelled as the accumulation of concentration parameters in a Dirichlet distribution).

Here we focus exclusively on inference. Paradigms that involve learning (e.g. Mayer et al., 2015), although arguably more ecologically valid (Gordon et al., 2017; 2019), do not allow stimulus visibility and attention to be dissociated or controlled for in as tight a manner as the minimal contrast paradigms around which the taxonomy is built – making it difficult to meaningfully compare interactions between factors.

To integrate the role of expectation – and the violation of expectation – into the taxonomy, we altered the prior probability of the “square present” sequence in the second level **D** vector for each of the four parameter settings used in the four-way taxonomy, such that the model was twice as confident a priori in either the “square present” sequence (consistent prior expectation condition) or the “random bar” sequence (inconsistent prior expectation condition). We treated each trial as independent, so the manipulation of prior expectations most plausibly corresponds to the use of explicit cues.

We retained the same generative model architecture used in the previous two sections, allowing us to independently manipulate expectation, attention, and stimulus strength. We are aware, however, that independently manipulating these factors in an experimental setting is far from trivial. In the interest of making our model empirically useful, we end this section by proposing a novel Posner cueing paradigm aimed at empirically validating the results of our simulations.

Finally, it must be highlighted that the behavioural results of the following simulations should be interpreted as directional hypotheses, as opposed to precise predictions about the percentage of seen vs unseen trials. In contrast to firing rates and ERPs, the model’s report behaviour (policy selection) depends on a number of



parameters (policy precision and motor stochasticity) that need to be fitted to each participant individually and will substantially shift the effect of prior expectation.

#### **4.4.1 Extend Taxonomy: Simulated Behaviour**

Here we presented the model with 800 “square present” trials, 100 corresponding to “consistent prior” and “inconsistent prior” conditions, for each of the four parameter settings used in the four-way taxonomy. Across all taxonomy conditions, we found that consistent prior expectations increased both the accuracy of forced-choice behaviour and the percentage of trials where the model reported having seen the stimulus (see figure 10). Similarly, inconsistent prior expectations decreased both the accuracy and the number of seen trials. The enhancing effect of consistent priors was particularly strong for the “weak signal + attention present” and “strong signal + attention absent” conditions – increasing the number of trials reported as seen in the four-way taxonomy from 8% to 79% and 16% to 87% respectively (compare figure 8 to figure 10). This result makes intuitive sense; expectations will have the greatest enhancing effect on visibility when a stimulus is only marginally below the threshold for “ignition”.

While these results are consistent with a number of studies showing an enhancing effect of consistent prior expectations for both forced-choice performance and stimulus detection (Aru et al., 2016; Stein, & Peelen, 2015) there is some evidence showing that when expectations are induced by explicit cues, they boost subjective visibility but do not alter accuracy (Andersen et al., 2019). Similarly, Lamy et al (2017) found that prior experience of a target increased visibility but did not alter response priming. Here forced-choice behaviour, like subjective report, depends on

policy selection at the second level of the model. However, if forced choice behaviour depends upon a distinct neural substrate that operates over a shorter timescale than subjective report, it would be better modelled by policy selection at the first level of the model (c.f. Maniscalco & Lau, 2016). If this were the case increasing the prior probability of a state at the second level (as we have done here) would have a marginal effect on forced-choice behaviour. This represents an important possible extension of our model architecture that will be addressed in future work.

### Extended Taxonomy: Behavior and Firing Rates

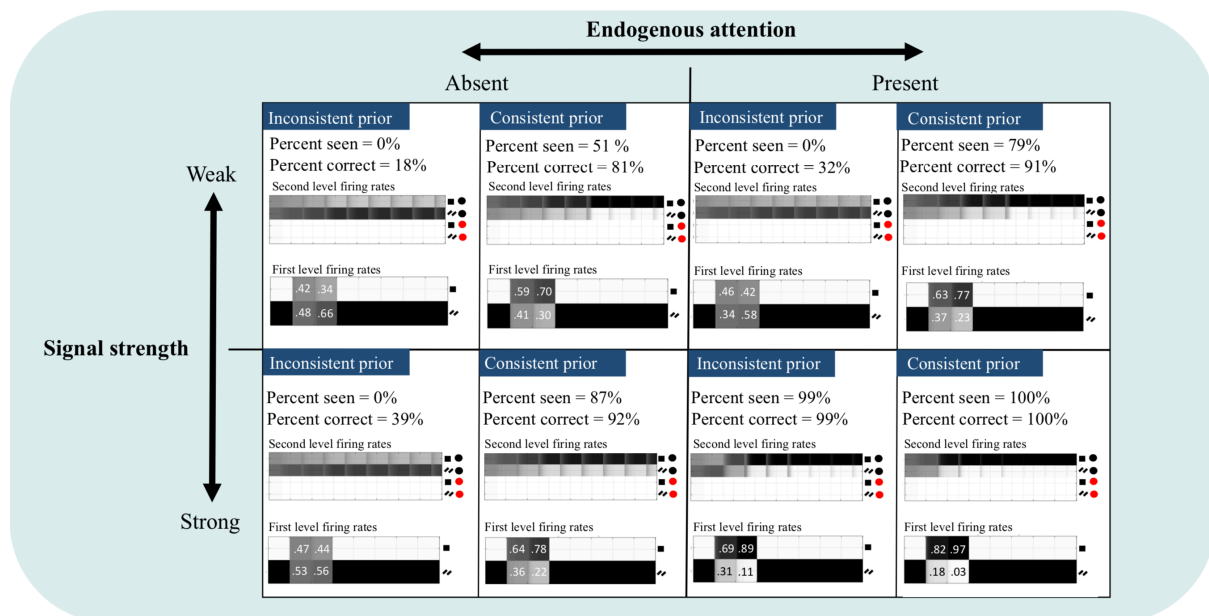


Figure 10. Report frequency, forced choice behaviour, and simulated firing rates predicted for each consistent-inconsistent prior combination of the quadrants shown in the four-way taxonomy. Relative to the results of the four-way taxonomy, consistent expectations increased forced choice accuracy, the percentage of trials reported as “seen”, and boosted the enhancing effect of feedback from the higher level. Inconsistent priors had the opposite effect, reducing accuracy, the percentage of “seen” trials and first level firing rates.

#### 4.4.2 Extending the Taxonomy: Simulated Firing Rates

Next, we simulated the firing rates associated with consistent and inconsistent priors. In all sections of the extended-taxonomy, relative to the four-way taxonomy, consistent priors enhanced the amplifying effect of feedback. Correspondingly, inconsistent priors dampened the effect of feedback. Interestingly, in the “strong signal + attention present” condition, although inconsistent priors dampened the amplifying effect of feedback (i.e., in comparison to the four-way taxonomy and the consistent prior condition), the net effect of feedback still raised the firing rate of first level neuronal populations.

This makes intuitive sense; feedback is driven by posterior confidence at the second level, and, in all but the “strong signal + attention present” condition, the sensory likelihood mapping is relatively imprecise. As such, posterior confidence at the second level is dominated by the effect of prior expectation. Consistent expectations increase second-level posterior confidence to a sufficiently large degree to shift otherwise unconscious trials over the threshold for “ignition,” thus magnifying the feedback to the first level (while inconsistent priors reduce posterior confidence and dampen top-down feedback). In contrast, when the first-level likelihood mapping is more precise, as is the case in the “strong signal + attention present” condition, inconsistent priors carry less influence and second-level posterior confidence in the “square-present” sequence is still high enough to enhance top-down feedback (i.e., despite it being dampened relative to the four-way taxonomy and consistent prior condition).

This result produces three novel predictions. First, when consistent expectations raise posterior confidence past the threshold for ignition, feedback from

frontoparietal to sensory regions will be enhanced. Second, when a stimulus is well above threshold, inconsistent expectations will reduce the amplification of feedback relative to consistent expectations. Third, based on the extrinsic connectivity predicted by the neural process theory (see figure 2), feedback from frontoparietal to sensory regions will be associated with a specific pattern of effective connectivity. Specifically, consistent prior expectations are predicted to disinhibit granular layers in the relevant neuronal populations in sensory cortices (via feedback connections originating in superficial pyramidal cells in frontoparietal regions), while inconsistent priors are predicted to inhibit granular layers in the same lower-level neuronal populations. This last prediction, although highly specific, can be readily tested via dynamic causal modelling (e.g. Parr et al., 2019b).

#### **4.4.3 Extend Taxonomy: Simulated Event Related Potentials**

Lastly, we simulated the ERPs predicted by our model for all combinations of consistent and inconsistent prior expectation conditions (see figure 11). In the “weak signal + attention absent” condition, relative to the four-way taxonomy, both inconsistent priors and consistent priors increased the amplitude of second level ERPs. When the signal was ambiguous in the “weak signal + attention present” and “strong signal +attention absent” conditions, consistent priors increased the amplitude of 2nd-level ERPs relative to both inconsistent priors and the four-way taxonomy. However, the increases in amplitude are small and may not be detectable at the level of the scalp. In contrast, for the “strong signal + attention present” condition consistent priors greatly increased the amplitude of the 2nd-level ERPs

relative to the four-way taxonomy, while inconsistent expectations decreased the amplitude of 2nd-level ERPs.

Again, this makes intuitive sense, upon presentation of the critical stimulus consistent priors increase the rate of second-level posterior confidence updates increasing the amplitude of 2nd-level ERPs, and boosting otherwise unconscious trials past the threshold for “ignition”.

Importantly, these results lead to another novel prediction. Contingent upon sensory input being sufficiently precise, consistent prior expectations induced via an explicit cue should increase the amplitude of the P3 (relative to expectation neutral conditions), whilst inconsistent prior expectations should decrease the amplitude of the P3.

Finally, it must be emphasised that this prediction is specific to paradigms that involve inference as opposed to learning. In the context of learning, the repeated presentation of a stimulus would be expected to cause the second level of the model to become increasingly confident in one sequence over another – to the point that the presentation of a critical stimulus would result in smaller belief updates, and consequently, result in smaller ERPs. Indeed, this is exactly what we see empirically. Melloni et al (2011) found that expectations induced via a history of prior exposure both increased the proportion of trials reported as seen and decreased the amplitude of the P3.

We describe this, not to provide a concrete hypothesis, but to distinguish the experimental settings appropriate to the predictions of the current model from similar but ultimately distinct settings.

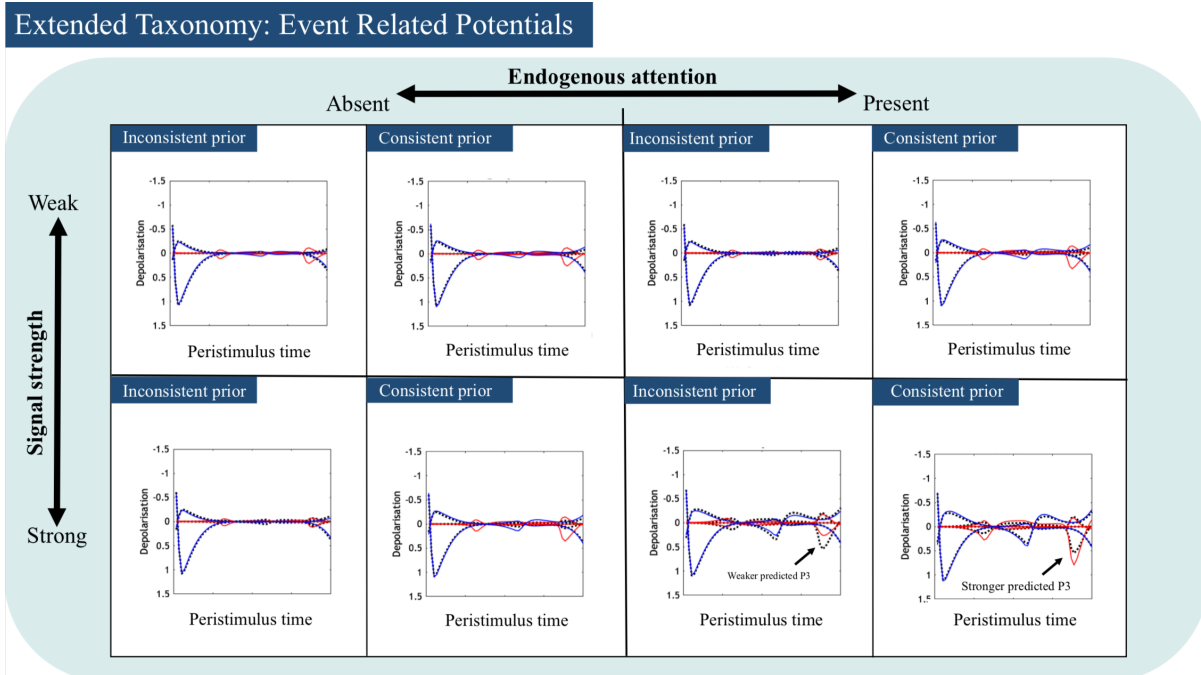


Figure 11. ERPs predicted for each consistent-inconsistent prior combination of the quadrants of the four-way taxonomy. Dotted black lines show the ERPs predicted by the original four-way taxonomy without expectation manipulations. Solid lines show the ERPs predicted for the specific consistent and inconsistent prior manipulations.

#### 4.4.4 Extend Taxonomy: A Novel Paradigm for Dissecting the Influences of Signal Strength, Attention, and Expectation on Conscious Access.

The structure of our generative model is generic enough to generalise across paradigms that involve inference; however, as noted above, in an empirical setting independently manipulating expectation, signal strength, and attention poses a number of methodological challenges, with expectation often being confounded with attention (e.g. Rahnev et al., 2011). In the interest of making the predictions of our model as straightforward to test as possible, we here outline a possible extension of the Posner cueing paradigm introduced by Kok and colleagues (2012) that would allow for the independent manipulation of expectation, signal strength, and attention.

The key feature of the design is the orthogonal manipulation of expectation and attention (see figure 12). Expectation can be manipulated in a block-wise manner with a predictive cue appearing at the beginning of every block consisting of a word (“left”, “right” or “neutral”) indicating the likelihood with which the stimulus will appear in a particular hemifield on each trial. Attention, in contrast, will be manipulated in a trial-wise manner, with a cue appearing at the start of every trial indicating the hemifield to which the subject should covertly direct their attention. However, the attention cue will contain no information about the likelihood of the stimulus’ location. Finally, stimulus strength will be manipulated by altering the time between the stimulus and the backward mask. Each block would begin with a predictive cue, and each trial would begin with the presentation of an attention cue, followed by a briefly presented stimulus (a grating in the above figure) that is congruent with the prediction on 75% of trials (paired with either a backwards mask or a blank). Subsequent to the presentation of the mask (or blank), subjects would be given a forced-choice task and asked to provide a subjective report.

Crucially, since expectation, attention, and stimulus strength are manipulated independently of each other, this paradigm could allow all 12 quadrants of the extended taxonomy to be studied within one paradigm.

## Extended Taxonomy: Novel Posner Paradigm

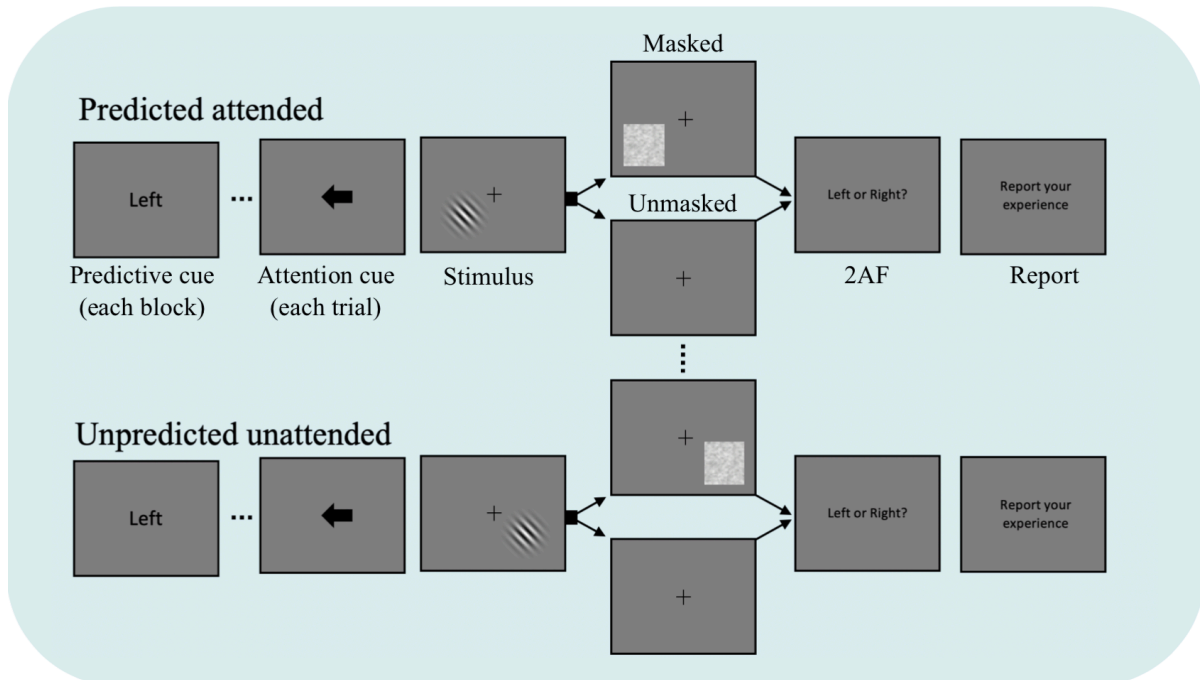


Figure 12. A potential extension of Posner cueing paradigm introduced by Kok and colleagues (2012) that could allow for the independent manipulation of expectation, stimulus strength, and attention. By manipulating expectation in a block-wise manner, and attention and stimulus strength in a trial-wise manner, the paradigm would allow all twelve combinations of expectation (consistent, neutral, inconsistent) by attention (present, absent) by stimulus strength (strong, weak) to be studied within one paradigm. Shown above are predicted/attended and unpredicted/unattended combinations.

## 5. Discussion

### 5.1 The role of deep temporal structure

The defining (measurable) feature of conscious access is subjective report (Baars, 1988; Dehaene, 2014; Fleming, 2019), which requires the coordination of processing between perceptual, motor, and auditory-verbal systems, all of which evolve over a rapid temporal scale. The core notion underlying our model is that conscious access is a fundamentally inferential process that can *only* occur at a level of processing that is sufficiently temporally deep to integrate information from lower levels of the



hierarchy and contextualise processing at these lower levels. To produce subjective reports, a system must infer the state of a lower level perceptual system, integrate this information into a representation that is not tethered to the moment by moment sensory flux, and use this representation to modulate more controlled, slowly evolving trajectories of action over time. Put another way, temporally deeper levels are necessary to encode patterns of covariance in lower-level sensory and motor representations over time under different goal states. Without a sensory representation updating predicted patterns of covariance at this higher level (to a sufficient degree), the use of that piece of sensory information by the more complex cognitive processes carried out at the higher level would be greatly limited (i.e., only promoting implicit biases through small changes in higher-level posterior distributions).

The first major insight afforded by our model is that many previous electrophysiological results can be reproduced based only on assuming a simple 2-level model with deep temporal structure. Self-reported conscious vs. unconscious percepts can then be accounted for by specifying the integrative functions of the higher level of the model that are plausible based on the known neural correlates of consciousness, and how the different hidden state factors (which could perhaps correspond to distributed processing hubs within association cortices implicated in domain-general cognition i.e. van den Heuvel et al., 2012) operate on those contents. The second major insight comes from how our model illustrates the way prior expectation can fit seamlessly within this structure – affording a number of novel, testable predictions.

## 5.2 Relationship to Other Models

The PGNW is a formal extension of the GNW and, as such, the models share many similarities. However, as should be clear by this point, the PGNW diverges substantially from the original GNW model. Ultimately, the point of departure for the PGNW is its implementation in an active inference architecture, which, as we have shown, has a number of important consequences. Foremost, by leveraging the process theory that accompanies active inference, we are able to reproduce/explain previous findings and make predictions about the neurobiological implementation of the inferential machinery that we argue underlies conscious access. Crucially, this also allows the PGNW to make specific predictions about the role of visual expectation. In contrast, the inferential machinery cited in support of the standard GNW model (i.e. King & Dehaene 2014) remains at the algorithmic level of description (in the sense of Marr, 1980). What the PGNW retains is the fundamental insight that conscious access makes information widely available to domain-general cognitive processes (i.e., represented by the different state factors at the second level of our model). It is for this reason that we retain the workspace label. However, unlike the initial conceptual account of the PGNW introduced by Whyte (2019) which, as Marvan and Tomáš (2020) point out, relies on explanatory machinery external to the active inference framework to explain conscious access, here we identify the global availability of information with temporally deep processing, and conscious access with the posterior confidence threshold required for report (broadly construed as goal-directed verbal report, button presses, saccades or any other method of goal-directed reporting of subjective content). As such the PGNW explains conscious access exclusively in terms of the explanatory tools of active inference.

The model that is most similar to the PGNW is Fleming's (2019) higher-order state-space (HOSS). Both the PGNW and HOSS are implemented in hierarchical generative models, and as such they entail similar predictions. However, there are two key points of separation. First, HOSS casts conscious access as a metacognitive inference about the presence or absence of a stimulus in the content of a generative model of perception. This inference relies on an abstract state representing presence or absence that is independent of the specific content of a stimulus. The function of this metacognitive state is to differentiate stimulus representations from noise distributions at lower levels of the model (c.f. Lau, 2007, 2019). A stimulus representation becomes available for report according to HOSS when the model infers that a distribution is different enough from a noise distribution to be classified as present.

There is a sense in which HOSS is simply a higher-order version of the PGNW. In fact, it would be relatively simple to introduce a new hidden state factor corresponding to the presence or absence of a stimulus and factorise a generative model such that the presence of the stimulus would be independent of the content of the stimulus. Computationally, however, this state would have no function. Precision estimation is an inbuilt feature of active inference architectures (see Parr & Friston 2017 for technical details) that modulates the updating process in response to the estimated reliability of a bottom-up signal, without needing to posit an additional abstract state representing presence and absence.

This brings us to a second key difference. According to the HOSS model, absence of a stimulus is explicitly represented in addition to all the possible states of a stimulus making the state space asymmetric. Fleming (2019) leverages this

asymmetry to explain the ignition response that sweeps across frontoparietal cortices during conscious access. Because there are many more ways that a stimulus can be present than absent, the presence of a stimulus causes larger belief updates than when a stimulus is absent. Although this is ultimately an empirical question, we regard the explicit representation of the absence of a stimulus in the perceptual state space to be a less realistic idealisation. Participants in minimal contrast paradigms are fully aware of the background screen and the task requirements, and frontoparietal cortices are, of course, active the whole time. The wide scale activation of these regions when a stimulus is seen is only apparent because we isolate the processes underlying the reportability of the stimulus, while holding all of these other variables constant. When participants do not report a stimulus, they are not perceiving “absence”; at a minimum the content of perception will include the background screen. It may be a mistake, therefore, to explicitly represent the absence of a stimulus in a generative model of perception. Instead, as we have done in the current model, the state space should consist of sequences of stimuli with and without the critical stimulus. The ignition response is then instead explained by the update that occurs at temporally deep levels of the model when the stimulus is seen as opposed to unseen.

### **5.3 Brief Note on Phenomenology**

Although not the focus of this paper, it is worth briefly clarifying how phenomenology does and does not situate plausibly within our model. Specifically, we would like to avoid implicitly conveying that the phenomenological contents of consciousness correspond to the contents of second-level states. A major problem with this is that

the contents of second-level states appear to operate on timescales that are too slow to match the moment-to-moment sensory flux.

However, there is also a problem with identifying phenomenological contents with the contents of lower-level states. To see why, consider that, as is the case in our model, there is no explicit representation of absence at the higher level generating reports. The model will always represent and report an experience of something (e.g., either lines or a square) based on the posterior distribution over second-level states. If this architecture is representative of human cognition, it highlights an interesting change in perspective. Specifically, the question about phenomenology being separable from access (c.f. Block, 2005) changes to a question about the possibility of inconsistencies between phenomenology and *what* was accessed (i.e., the states represented at the second level). For example, if the agent reported currently seeing lines (and not a square), and yet a square stimulus was present, a strong distinction between phenomenology and access would not merely entail that “square” phenomenology was present but not accessed. Instead it would mean that the agent’s confident self-reported phenomenology of experiencing lines (i.e., what was represented at the second level) was inconsistent with their “true” phenomenology (in other words, they would be wrong about what they believed they were currently experiencing or had just experienced). Taken to the extreme, this would entail that any honestly reported phenomenology could be problematically different from true phenomenology (e.g., a person honestly reporting experiencing a loud screeching sound could have the “true” phenomenology of a hearing a piece of classical music).

Another way to highlight this problem formally is by considering that one could manipulate the second-level likelihood mapping (**A** matrix) in our model while leaving the first-level likelihood mapping unchanged. If so, first-level states would still reliably track presented stimuli (e.g., square stimuli would activate first-level square states), but those states could now update the second-level in an entirely different way. For example, with the right second-level likelihood mapping, red circle and line representations at the first level could be specified so as to activate representations of, and promote self-reported phenomenology of, a black circle and a square (which would also obviously be problematic for the empirical study of conscious experience).

To avoid this uncomfortable conclusion, while also keeping sensory phenomenology at the correct timescale, we suggest that phenomenology in our model is most plausibly situated *at the point in which (and based on the manner in which) lower-level representations update the content of higher-level states*. In other words, at the point where the higher level assesses or “decodes” the contents of the lower level through the second-level likelihood mapping (**A** matrix). These updates occur at the fast timescales associated with the sensory flux; yet, the nature of each fast-timescale update (i.e., the nature of the influence the first level has on the second level at each time point) will necessarily be correlated with self-reported beliefs about what type of phenomenology was experienced – preventing the possibility of strong disagreements between “true” and self-reported beliefs about phenomenology (for more on this line of argument see Smith, 2016).

## 5.4 Limitations

To make the PGNW testable, we have deliberately limited the scope of the model to experimental settings where visual consciousness is operationalised via report. We follow Baars (1988) in taking report, or rather the availability of information for report, as the epistemic foundation of the scientific study of consciousness. However, we acknowledge that report paradigms come with methodological difficulties (Tsuchiya et al., 2015). By the same token, in limiting the scope of the PGNW to the visual modality, we are aware that we are reifying the pervasive bias in consciousness science of primarily studying vision. That said, it is crucial to emphasise that the model structure is sufficiently general that it can be straightforwardly applied to other modalities (e.g., first-level observations could fairly easily be understood as auditory as opposed to visual). A somewhat similar 2-level architecture was also recently employed to simulate emotional awareness based on interoceptive stimuli (Smith, et al., 2019a).

In addition to these big picture limitations, our model has a number of more specific limitations that apply strictly to the study of visual consciousness. Principally, our model only has two levels and we treat the entire visual system as a singular and discrete level instead of modelling it for what it is – a continuous and multi-level system. The need for multiple levels brings us to the next limitation. As Kouider et al (2010) argue, people are often only partially aware of a visual scene in the sense that they may be aware of an object's colour but not its identity. This requires information to skip levels of the hierarchy, which is also not possible in the present model. A more complete model would therefore allow both shallower and deeper

representations to selectively update the workspace (e.g., allowing separable awareness of representations of an eye vs. a face vs. a person's identity).

## 6. Conclusions and Future Directions

This paper introduced a formal extension of the global neuronal workspace – the predictive global neuronal workspace – implemented within a deep active inference architecture. In addition to explaining and unifying otherwise disparate findings in the neural correlates of visual consciousness literature, the predictive global neuronal workspace generates several predictions and mechanistic explanations about the relationship of the P3 to subjective report, the neurobiological implementation of the inferential machinery underlying conscious access, and the role of expectation in visual consciousness.

In future work, we hope to build on the wealth of existing active inference models (e.g. Allen et al., 2019; Parr et al., 2019c; Smith et al., 2019a; 2019b), to extend the PGNW to other sensory modalities and more sophisticated experimental paradigms.

**Software note:** The generative model detailed in this paper used a generic belief updating scheme (**spm\_MDP\_VB\_X.m**) implemented in Matlab code using the freely available SPM academic software: <https://www.fil.ion.ucl.ac.uk/spm/>. The scripts used to produce the simulations are included in the supplementary material.



## References

- Allen, M., Levy, A., Parr, T., & Friston, K. J. (2019). In the Body's Eye: The Computational Anatomy of Interoceptive Inference. *BioRxiv*, 603928.
- Andersen, L. M., Overgaard, M., & Tong, F. (2019). Visual expectations change subjective experience without changing performance. *Consciousness and cognition*, 71, 59-69.
- Aru, J., Bachmann, T., Singer, W., & Melloni, L. (2012). Distilling the neural correlates of consciousness. *Neuroscience & Biobehavioral Reviews*, 36(2), 737-746.
- Aru, J., Rutiku, R., Wibral, M., Singer, W., & Melloni, L. (2016). Early effects of previous experience on conscious perception. *Neuroscience of consciousness*, 2016(1), niw004.
- Aru, J., Tulver, K., & Bachmann, T. (2018). It's all in your head: Expectations create illusory perception in a dual-task setup. *Consciousness and cognition*, 65, 197-208.
- Baars, B.J. (1988). *A Cognitive Theory of Consciousness*. Cambridge, UK: Cambridge University Press.
- Bannister, A. P. (2005). Inter-and intra-laminar connections of pyramidal cells in the neocortex. *Neuroscience research*, 53(2), 95-103.
- Bisenius, S., Trapp, S., Neumann, J., & Schroeter, M. L. (2015). Identifying neural correlates of visual consciousness with ALE meta-analyses. *Neuroimage*, 122, 177-187.
- Block, N. (2005). Two neural correlates of consciousness. *Trends in cognitive sciences*, 9(2), 46-52.

Boly, M., Massimini, M., Tsuchiya, N., Postle, B. R., Koch, C., & Tononi, G. (2017).

Are the neural correlates of consciousness in the front or in the back of the cerebral cortex? Clinical and neuroimaging evidence. *Journal of Neuroscience*, 37(40), 9603-9613.

Bogacz, R. (2017). A tutorial on the free-energy framework for modelling

perception and learning. *Journal of mathematical psychology*, 76, 198-211.

Buckley, C. L., Kim, C. S., McGregor, S., & Seth, A. K. (2017). The free energy

principle for action and perception: A mathematical review. *Journal of Mathematical Psychology*, 81, 55-79.

Chang, A. Y. C., Kanai, R., & Seth, A. K. (2015). Cross-modal prediction changes

the timing of conscious access during the motion-induced blindness. *Consciousness and cognition*, 31, 139-147.

Cohen, M. A., Ortego, K., Kyroudis, A., & Pitts, M. (2020). Distinguishing the neural

correlates of perceptual awareness and post-perceptual processing. *bioRxiv*.

Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V., & Friston, K. (2020a). Active

inference on discrete state-spaces: a synthesis. *arXiv preprint arXiv:2001.07203*.

Da Costa, L., Parr, T., Sengupta, B., & Friston, K. (2020b). Natural selection finds

natural gradient. *arXiv preprint arXiv:2001.08028*.

Dehaene, S. (2014). Consciousness and the brain: Deciphering how the brain

codes our thoughts. Penguin.

Dehaene, S., Changeux, J. P., & Naccache, L. (2011). The global neuronal

workspace model of conscious access: from neuronal architectures to clinical applications. In *Characterizing consciousness: from cognition to the clinic?* (pp. 55-84). Springer Berlin Heidelberg.

Dehaene, S., & Changeux, J. P. (2011). Experimental and theoretical approaches to Conscious processing. *Neuron*, 70(2), 200-227.

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in cognitive sciences*, 10(5), 204-211.

Dehaene, S., Charles, L., King, J. R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current opinion in neurobiology*, 25, 76-84.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79(1), 1-37.

Dehaene, S., Naccache, L., Cohen, L., LeBihan, D., Mangin, J. F., Poline, J. B., & Rivière, D. (2001). Cerebral mechanisms of word masking and unconscious Repetition priming. *Nature neuroscience*, 4(7), 752-758.

Dehaene, S., Sergent, C., & Changeux, J. P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences*, 100(14), 8520-8525.

DelCul, A., Baillet, S., & Dehaene, S. (2007). Brain dynamics underlying the non-linear threshold for access to consciousness. *PLoSbiology*, 5(10), e260.

Del Cul, A., Dehaene, S., Reyes, P., Bravo, E., & Slachevsky, A. (2009). Causal role of prefrontal cortex in the threshold for access to consciousness. *Brain*, 132(9), 2531-2540.

Denison, R. N., Piazza, E. A., & Silver, M. A. (2011). Predictive context influences perceptual selection during binocular rivalry. *Frontiers in human neuroscience*, 5, 166.

- Dong, X., Zhang, M., Dong, B., Jiang, Y., & Bao, M. (2019). Reward produces learning of a consciously inaccessible feature. *bioRxiv*.
- Friston, K. (2005). A theory of cortical responses. *Philosophical transactions of the Royal Society B: Biological sciences*, 360(1456), 815-836.
- Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10(86), 20130475.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017a). Active inference: a process theory. *Neural computation*, 29(1), 1-49.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2016). Active inference and learning. *Neuroscience & Biobehavioral Reviews*, 68,862-879.
- Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521), 1211-1221.
- Friston, K. J., Rosch, R., Parr, T., Price, C., & Bowman, H. (2018). Deep temporal models and active inference. *Neuroscience & Biobehavioral Reviews*, 90, 486-501.
- Friston, K. J., Parr, T., & de Vries, B. (2017b). The graphical brain: belief propagation and active inference. *Network Neuroscience*, 1(4), 381-414
- Garcia-Cabezas, M. A., Zikopoulos, B., & Barbas, H. (2019). The Structural Model: a theory linking connections, plasticity, pathology, development and evolution of the cerebral cortex. *Brain Structure and Function*, 224(3), 985-1008.
- Gordon, N., Koenig-Robert, R., Tsuchiya, N., Van Boxtel, J. J., & Hohwy, J. (2017). Neural markers of predictive coding under perceptual uncertainty revealed with Hierarchical Frequency Tagging. *Elife*, 6, e22749.

- Gordon, N., Tsuchiya, N., Koenig-Robert, R., & Hohwy, J. (2019). Expectation and attention increase the integration of top
- Gottlieb, J. (2007). From thought to action: the parietal cortex as a bridge between perception, action, and cognition. *Neuron*, 53(1), 9-16.
- Kapoor, V., Besserve, M., Logothetis, N. K., & Panagiotaropoulos, T. I. (2018). Parallel and functionally segregated processing of task phase and conscious content in the prefrontal cortex. *Communications biology*, 1(1), 215.
- King, J. R., Pescetelli, N., & Dehaene, S. (2016a). Brain mechanisms underlying the brief maintenance of seen and unseen sensory information. *Neuron*, 92(5), 1122-1134.
- King, J. R., & Dehaene, S. (2014). A model of subjective report and objective discrimination as categorical decisions in a vast representational space. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1641), 20130204
- Kouider, S., De Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in cognitive sciences*, 14(7), 301-307.
- Lau, H. C. (2007). A higher order Bayesian decision theory of consciousness. *Progress in brain research*, 168, 35-48.
- Lau, H. (2019). Consciousness, metacognition, & perceptual reality monitoring.
- Lakatos I. Falsification and the methodology of scientific research programmes. Criticism and the growth of knowledge. 1970 Sep 2;4:91-196.
- Lamy, D., Carmel, T., & Peremen, Z. (2017). Prior conscious experience enhances conscious perception but does not affect response priming☆. *Cognition*, 160, 62-81

- Maniscalco, B., & Lau, H. (2016). The signal processing architecture underlying subjective reports of sensory awareness. *Neuroscience of Consciousness*, 2016(1).
- Marti, S., & Dehaene, S. (2017). Discrete and continuous mechanisms of temporal selection in rapid visual streams. *Nature communications*, 8(1), 1955.
- Marvan, Tomáš., & Havlík, Marek (2020) Is Predictive Processing a Theory of Consciousness?. Preprint. URL: <http://philsciarchive.pitt.edu/id/eprint/16890> (accessed 2020-02-07).
- Marx, S., & Einhäuser, W. (2015). Reward modulates perception in binocular rivalry. *Journal of vision*, 15(1), 11-11.
- Mayer, A., Schwiedrzik, C. M., Wibral, M., Singer, W., & Melloni, L. (2015). Expecting to see a letter: alpha oscillations as carriers of top-down sensory predictions. *Cerebral Cortex*, 26(7), 3146-3160.
- Meijs, E. L., Slagter, H. A., de Lange, F. P., & van Gaal, S. (2018). Dynamic Interactions between top-down expectations and conscious awareness. *Journal of Neuroscience*, 38(9), 2318-2327.
- Melloni, L., Schwiedrzik, C. M., Müller, N., Rodriguez, E., & Singer, W. (2011). Expectations change the signatures and timing of electrophysiological correlates of perceptual awareness. *Journal of Neuroscience*, 31(4), 1386- 1396.
- Mesulam, M. M. (1998). From sensation to cognition. *Brain: a journal of neurology*, 121(6), 1013-1052.
- Michel, M., Fleming, S. M., Lau, H., Lee, A. L., Martinez-Conde, S., Passingham, R. E., ... & Liu, K. (2018). An informal internet survey on the current state of consciousness science. *Frontiers in psychology*, 9, 2134.

- Panagiotaropoulos, T. I., Deco, G., Kapoor, V., & Logothetis, N. K. (2012). Neuronal discharges and gamma oscillations explicitly reflect visual consciousness in the lateral prefrontal cortex. *Neuron*, 74(5), 924-935.
- Parr, T., Corcoran, A. W., Friston, K. J., & Hohwy, J. (2019c). Perceptual awareness and active inference. *Neuroscience of consciousness*, 2019(1), niz012.
- Parr, T., & Friston, K. J. (2017a). Working memory, attention, and salience in active inference. *Scientific reports*, 7(1), 14678.
- Parr, T., & Friston, K. J. (2017b). Uncertainty, epistemics and active inference. *Journal of The Royal Society Interface*, 14(136), 20170376.
- Parr, T., & Friston, K. J. (2018). The anatomy of inference: generative models and brain structure. *Frontiers in computational neuroscience*, 12, 90.
- Parr, T., Markovic, D., Kiebel, S. J., & Friston, K. J. (2019a). Neuronal message passing using Mean-field, Bethe, and Marginal approximations. *Scientific reports*, 9(1), 1889.
- Parr, T., Mirza, M. B., Caglan, H., & Friston, K. J. (2019b). Dynamic causal modelling of active vision. *Journal of Neuroscience*, 39(32), 6265-6275.
- Parr, T., Rikhye, R. V., Halassa, M. M., & Friston, K. J. (2019). Prefrontal computation as active inference. *Cerebral Cortex*.
- Peter, U. T., Martinez-Conde, S., Schlegel, A. A., & Macknik, S. L. (2005). Visibility, visual awareness, and visual masking of simple unattended targets are confined to areas in the occipital cortex beyond human V1/V2. *Proceedings of the National Academy of Sciences*, 102(47), 17178-17183.
- Peters, M. A., & Lau, H. (2015). Human observers have optimal introspective access to perceptual processes even for visually masked stimuli. *Elife*, 4, e09651.
- Pinto, Y., van Gaal, S., de Lange, F. P., Lamme, V. A., & Seth, A. K. (2015).

Expectations accelerate entry of visual stimuli into awareness. *Journal of Vision*, 15(8), 13-13.

Pitts, M. A., Padwal, J., Fennelly, D., Martínez, A., & Hillyard, S. A. (2014).

Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness. *Neuroimage*, 101, 337-350.

Rahnev, D., Lau, H., & De Lange, F. P. (2011). Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. *Journal of Neuroscience*, 31(29), 10741-10748.

Salti, M., Monto, S., Charles, L., King, J. R., Parkkonen, L., & Dehaene, S. (2015).

Distinct cortical codes and temporal dynamics for conscious and unconscious percepts. *Elife*, 4, e05652.

Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature neuroscience*, 8(10), 1391.

Smith, R. (2016). The relationship between consciousness, understanding, and rationality. *Philosophical Psychology*, 29(7), 943-957.

Smith, R., Lane, R. D., Parr, T., & Friston, K. J. (2019a). Neurocomputational mechanisms underlying emotional awareness: insights afforded by deep active inference and their potential clinical relevance. *Neuroscience & Biobehavioral Reviews*.

Smith, R., Parr, T., & Friston, K. J. (2019b). Simulating emotions: An active inference model of emotional state inference and emotion concept learning. *bioRxiv*, 640813.



Stein, T., & Peelen, M. V. (2015). Content-specific expectations enhance stimulus detectability by increasing perceptual sensitivity. *Journal of Experimental Psychology: General*, 144(6), 1089.

Sterzer, P., Frith, C., & Petrovic, P. (2008). Believing is seeing: expectations alter visual awareness. *Current Biology*, 18(16), R697-R698.

Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. (2015). No-report paradigms: extracting the true neural correlates of consciousness. *Trends in cognitive sciences*, 19(12), 757-770.

van den Heuvel, M. P., Kahn, R. S., Goñi, J., & Sporns, O. (2012). High-cost, high-capacity backbone for global brain communication. *Proceedings of the National Academy of Sciences*, 109(28), 11372-11377.

Valuch, C., & Kulke, L. (2019). Predictive context biases binocular rivalry onset in children and adults-down and bottom-up signals in perception through different pathways. *PLoS biology*, 17(4).

Whyte, Christopher J. Integrating the global neuronal workspace into the framework of predictive processing: Towards a working hypothesis. *Consciousness and cognition* 73 (2019): 102763.

Zeki, S. (2003). The disunity of consciousness. *Trends in cognitive sciences*, 7(5), 214-218.