Published in final edited form as:
J Neurosci. 2013 May 8; 33(19): 8227–8236. doi:10.1523/JNEUROSCI.4255-12.2013.

Free energy, precision and learning: the role of cholinergic neuromodulation

Rosalyn J. Moran[*,1,2], **Pablo Campo**[1,3], **Mkael Symmonds**[1], **Klaas E. Stephan**[1,4,5], **Raymond J. Dolan**[1], and **Karl J. Friston**[1]

[1]Wellcome Trust Centre for Neuroimaging, Institute of Neurology, University College London, 12 Queen Square, London WC1N 3BG, UK [2]Virginia Tech Carilion Research Institute & Bradley Department of Electrical and Computer Engineering, Virginia Tech, Roanoke, Virginia, USA [3]Department of Basic Psychology, Autonoma University of Madrid, Madrid, Spain [4]Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich & Swiss Federal Institute of Technology (ETH Zurich) [5]Laboratory for Social and Neural Systems Research, Dept. of Economics, University of Zurich, Switzerland

Abstract

Acetylcholine (ACh) is a neuromodulatory transmitter implicated in perception and learning under uncertainty. This study combined computational simulations and pharmaco-electroencephalography in humans, to test a formulation of perceptual inference based upon the free energy principle. This formulation suggests that acetylcholine enhances the precision of bottom-up synaptic transmission in cortical hierarchies by optimising the gain of supragranular pyramidal cells. Simulations of a mismatch negativity paradigm predicted a rapid trial-by-trial suppression of evoked sensory prediction error (PE) responses that is attenuated by cholinergic neuromodulation. We confirmed this prediction empirically with a placebo-controlled study of cholinesterase inhibition. Furthermore – using dynamic causal modelling – we found that drug-induced differences in PE responses could be explained by gain modulation in supragranular pyramidal cells in primary sensory cortex. This suggests that acetylcholine adaptively enhances sensory precision by boosting bottom-up signalling when stimuli are predictable, enabling the brain to respond optimally under different levels of environmental uncertainty.

Keywords

Free Energy Principle; Predictive Coding; Neuromodulation; Acetylcholine; Galantamine; Oddball Response; Precision; Dynamic Causal Modelling

[*]Rosalyn Moran, Virginia Tech Carilion Research Institute, 2 Riverside Circle, Roanoke Virginia 24016, Tel (1) 540-526-2136, rosalynj@vtc.vt.edu.

Software Note

All the simulations reported in this paper are available in the DEM toolbox of SPM http://www.fil.ion.ucl.ac.uk/spm/. The inversions can be simulated for different drug effects by changing DEM_demo_MMN.m {Placebo: M(1).hE = 4; M(1).hC = exp(−4)} or {Galantamine: M(1).hE = 6; M(1).hC = exp(−6)}.

Europe PMC Funders Author Manuscripts

Europe PMC Funders Author Manuscripts

Introduction

The Helmholtzian notion of the brain as a statistical inference machine (Hemholtz 1866, von Helmholtz 1962) can be realised – under free energy formulations of perceptual inference – as a neurobiologically plausible implementation of predictive coding. In this framework, the brain entails a generative model of its environment that reproduces the hierarchical and dynamic generation of sensory input. Neuronal responses are proposed to reflect inference on the (hidden) causes of sensory stimuli that enable the brain to predict its sensations (Mumford 1992; Rao and Ballard 1999: Friston 2010).

The implicit Bayes-optimal recognition requires statistical distributions over hidden states and the environment to be encoded by neuronal activity in cortical hierarchies; where these representations are updated until they mirror real-world statistics through accurate predictions of sensory input. To weight hierarchical predictions with greater or lesser importance, optimal inference assigns greater or lesser precision (inverse variance) to different hierarchical levels.

One candidate mechanism, for signalling precision is neuromodulation: acetylcholine (ACh) is well equipped to influence low-level auditory processing through modulating postsynaptic gain. Traditional views of ACh as a non-specific modulator of arousal that responds – through extra-synaptic volume-transmission – to enhanced sensory and task demands (Giovannini, Rakovska et al. 2001, Hasselmo and McGaughy 2004, Pepeu and Giovannini 2004, Dani and Bertrand 2007) are being rethought on the basis of improved (multi-region) microdialysis recordings in animal studies. These data suggest a selective cholinergic modulation of sensory cortex (Nelson, Sarter et al. 2005, Sarter, Hasselmo et al. 2005, Benarroch 2010, Hasselmo and Sarter 2010, Fadel 2011).

Under the free energy formulation, see Figure 1, ACh and other neuromodulators encode changes in the precision of (certainty about) prediction errors in sensory cortical hierarchies (Friston 2008). Each level of a processing hierarchy sends predictions to the level below, which reciprocate bottom-up signals. These signals are prediction errors that report discrepancies between top-down predictions and representations at each level (Kiebel, Daunizeau et al. 2009). This recurrent message passing continues until prediction errors are minimised throughout the hierarchy. The ensuing Bayes optimal perception rests on optimising precision at each level of the hierarchy that is commensurate with the environmental statistics they represent. Put simply, to infer the causes of sensory input, the brain has to recognise when sensory information is noisy or uncertain and down-weight it suitably in relation to top-down predictions (Yu and Dayan 2005, Deco and Thiele 2009). In this work, we address cholinergic neuromodulation as a candidate for optimising precision: first we simulate (Bayes-optimal) neural responses to sequences of repeated stimuli, under different levels of sensory precision (modulatory gain). We then compare simulated neuronal responses to empirical event related responses from an EEG study, using an oddball paradigm and a pharmacological (cholinergic) manipulation. The simulations and empirical results suggest that ACh biases perception towards bottom-up sensory processing – by boosting prediction errors at the lower (sensory) levels of the hierarchy (see Figure 1).

Europe PMC Funders Author Manuscripts

The empirical responses – and associated changes in postsynaptic gain – provide strong support for a specific role of ACh during perceptual inference under uncertainty.

Materials and Methods

Simulating Brain Responses using Variational Free Energy Minimisation

Predictive coding schemes (Rao and Ballard 1999) invoke recurrent neural message passing to simultaneously predict sensory stimuli and report prediction errors (e.g., to account for 'pop-out', figure-surround responses in visual cortex; (Lamme 1995, Zipser, Lamme et al. 1996)). The free energy framework places predictive coding in a more general setting, using dynamic and hierarchical generative models of hidden environmental states (causes of sensory inputs) and an approximate posterior density on the states and parameters of the generative model, encoded by synaptic activity and weights respectively. This approximate posterior (recognition) density is updated to maximise the evidence for the generative model (Kiebel, Daunizeau et al. 2009). Mathematically, this optimisation involves minimising a variational free energy bound on log evidence (henceforth free energy) which, when minimised, provides a recognition density that approximates the true posterior – this is also known as approximate Bayesian inference. Crucially, the process of free energy minimisation – in networks implementing predictive coding – can be used to predict the neuronal dynamics one would expect to measure using EEG – see Figure 1 (Friston and Kiebel 2009). Here, we will use dynamic causal modelling to test the theoretical assumptions about the neurobiological mechanisms that support Bayes-optimal inference in an oddball task.

To produce quantitative predictions about this inference, we generated auditory stimuli in blocks of different frequencies. We then simulated how the brain would respond to the stimuli if it minimised free energy using a neuronally plausible predictive coding scheme. In brief, this scheme encodes beliefs about the world by associating synaptic activity with expected states of the world, causal structure (parameters controlling transitions among states) with synaptic connection strengths and the precision (inverse variance) of random fluctuations of hidden states with synaptic gain: see (Friston 2008) for details. By assuming these beliefs are encoded with a Gaussian density, the expected states, parameters and precision can be optimised in a fairly straightforward way using a gradient descent on free energy – which is also known as Bayesian filtering (c.f., Kalman filtering). This Bayesian filtering or predictive coding can be formulated in terms of Bayesian updates that are driven by precision-weighted prediction errors. Mathematically, precision increases the influence of prediction errors, when they are precise. In neurobiological terms, it is generally assumed that prediction errors are encoded by pyramidal cells in the superficial layers of cortex that receive top-down predictions from deeper laminae in higher regions (Mumford 1992, Cauller 1995, Bar 2003). The precision weighting of these prediction errors units depends on their postsynaptic gain (c.f., Kalman gain). By changing the precision (synaptic gain) at a particular level of cortical hierarchy one can simulate the effects of manipulating cholinergic neuromodulation at this level, on the ensuing inference (Figure 1).

In more detail, we assume that the hierarchical nature of cortical organisation (Felleman and Van Essen 1991) implies a hierarchical model of the sensorium that can be described by a set of stochastic differential equations of the following general form:

$$\begin{aligned} y= &\ g\left(x^{(1)}, v^{(1)}, \theta\right) + \omega_v^{(1)} \\ \dot{x}^{(1)}= &\ f\left(x^{(1)}, v^{(1)}, \theta\right) + \omega_x^{(1)} \\ &\vdots \\ v^{(i-1)}= &\ g\left(x^{(i)}, v^{(i)}, \theta\right) + \omega_v^{(i)} \\ \dot{x}^{(i)}= &\ f\left(x^{(i)}, v^{(i)}, \theta\right) + \omega_x^{(i)} \\ &\vdots \\ \omega_x^{(i)} \sim &\ N\left(0, \Pi_x^{(i)-1}\right) \\ \omega_v^{(i)} \sim &\ N\left(0, \Pi_x^{(i)-1}\right) \\ \Pi_x^{(i)}= &\ exp\left(\gamma_x^{(i)}\right) \cdot I \\ \Pi_v^{(i)}= &\ exp\left(\gamma_v^{(i)}\right) \cdot I \end{aligned} \tag{1}$$

These equations model how sensory data of generated causes *v*, from level *i* enter as inputs to lower levels, eliciting changes in hidden states at that level $x^{(i-1)}$. These hidden states then produce further hidden causes that cascade through lower hidden states to the bottom of the hierarchy that generates sensory input $y(t)$. These hidden variables (causes and states) could correspond to things like the motion of objects in the visual field. The functions *f* and *g* that determine how hidden variables conspire to produce sensations have parameters θ that represent causal structure in the world (Friston 2005). These variables and their motion are subject to Gaussian random fluctuations $\omega^{(i)}$ with log-precision $\gamma^{(i)}$.

Given the form of the generative model (Equation 1) one can write down the differential equations for predictive coding or Bayesian filtering in terms of prediction errors on the hidden variables.

$$\begin{aligned} \dot{\tilde{\mu}}_x^{(i)}= &\ D\tilde{\mu}_x^{(i)} + \frac{\partial \tilde{g}^{(i)}}{\partial \tilde{\mu}_x^{(i)}} \cdot \Pi_v^{(i)} \tilde{\varepsilon}_v^{(i)} + \frac{\partial \tilde{f}^{(i)}}{\partial \tilde{\mu}_x^{(i)}} \cdot \Pi_x^{(i)} \tilde{\varepsilon}_x^{(i)} - D \cdot \Pi_x^{(i)} \tilde{\varepsilon}_x^{(i)} \\ \dot{\tilde{\mu}}_v^{(i)}= &\ D\tilde{\mu}_v^{(i)} + \frac{\partial \tilde{g}^{(i)}}{\partial \tilde{\mu}_v^{(i)}} \cdot \Pi_v^{(i)} \tilde{\varepsilon}_v^{(i)} + \frac{\partial \tilde{f}^{(i)^T}}{\partial \tilde{\mu}_v^{(i)}} \cdot \Pi_x^{(i)} \tilde{\varepsilon}_x^{(i)} - \Pi_v^{(i+1)} \tilde{\varepsilon}_v^{(i+1)} \\ \tilde{\varepsilon}_x^{(i)}= &\ D\tilde{\mu}_x^{(i)} - \tilde{f}^{(i)}\left(\tilde{\mu}_x^{(i)}, \tilde{\mu}_v^{(i)}\right) \\ \tilde{\varepsilon}_v^{(i)}= &\ \tilde{\mu}_v^{(i-1)} - \tilde{g}^{(i)}\left(\tilde{\mu}_x^{(i)}, \tilde{\mu}_v^{(i)}\right) \end{aligned} \tag{2}$$

Equation (2) can be derived fairly easily by computing the free energy for the hierarchical model in Equation (1) and using its gradients, perform a generalised descent. In neural network terms, Equation (2) says that error-units $\tilde{\varepsilon}^{(i)}$ receive predictions from the same hierarchical level $\tilde{\mu}_v^{(i-1)}$ and the level above $\tilde{\mu}_v^{(i)}$. Conversely, conditional expectations (encoded by the activity of state units) are driven by prediction errors from the same level $\tilde{\varepsilon}_v^{(i+1)}$ and the level below $\tilde{\varepsilon}_v^{(i)}$. This is the essence of recurrent message passing between hierarchical levels to suppress free energy or prediction error: see Friston (2008) for a more detailed discussion. Equation (2) shows how precision $\Pi^{(i)}$ plays an important role in weighting the influence of prediction errors at any level of the hierarchy. In other words, by

changing the precision on the prediction errors, we can bias inference towards sensory information or top-down (empirical) prior beliefs.

To simulate perceptual inference one simply integrates Equation (2), to see how the predictions and prediction errors respond to sensory input. To simulate an oddball paradigm, we generated a stream of auditory tones according to a roving design (cf. Haenschel et al. 2005; Figure 2): this comprised mini-blocks of pure tones (up to 10 repetitions in length). The onset of a tone is caused by a Gaussian bump function, that peaked at 100 ms. This hidden cause $(v)^{(1)}$ perturbed two dynamic hidden states $x^{(1)}=\left[x_1^{(1)}, x_2^{(1)}\right]^T$ controlling the amplitude and frequency of auditory stimuli. The hidden states were mixed using parameters θ_f to drive their motion and were mapped, using parameters θ_g to the amplitude and frequency of auditory stimuli. This model was used to simulate tones of different frequencies (by changing the parameter *C* in Equation 3). We simulated a change from a high tone to a low tone, with ten repetitions: d1-d10. State noise was assumed to be zero mean with a log-precision of 16. This generative model can be written in the form of Equation 1 as

$$\begin{aligned} y(t) &= \theta_g x^{(1)} + \omega_v^{(1)} \\ \dot{x}^{(1)} &= \theta_f x^{(1)} + \begin{bmatrix} v^{(1)} \\ 0 \end{bmatrix} + \omega_x^{(1)} \\ v^{(1)} &= \eta(t) + \omega_v^{(2)} \\ \theta_g &= \begin{bmatrix} 0 & 1 \\ C & 0 \end{bmatrix} \quad \theta_f = \begin{bmatrix} -1/16 & 4/16 \\ -2/16 & -1/16 \end{bmatrix} \end{aligned} \tag{3}$$

To simulate neuronal responses, we assumed that the brain was equipped with the same hierarchical model used to generate the stimuli, but did not know the parameter controlling frequency modulation over peristimulus time (the *C* parameter above). The brain must therefore learn this parameter over successive repetitions. This learning also corresponds to a gradient descent on free energy and – under the simple linear model above – reduces to associative plasticity (see Friston 2008 for details). Using the generative model in Equation (3), we simulated neuronal responses using Equation (2). This effectively updates posterior beliefs about the current stimulus using prediction errors (Figure 1). Crucially, these Bayesian updates involve not only the parameters responsible for generating sensory predictions but also the precision encoding the uncertainty about those predictions.

We used a veridical prior for the presentation of a high frequency tone (*C* = 4), then presented ten tones at a frequency using *C=1* (d1-d10). After the scheme had observed each tone, the prior expectation about the unknown parameter was updated to the posterior expectation – similarly for the priors on the precision of sensory noise. We simulated the neuronal responses using the precision weighted prediction errors of tone frequency ($\xi_v^{(1)}=\Pi_v^{(i)}\tilde{\varepsilon}_v^{(i)}$ in Figure 1; upper panel). Assuming that prediction errors are reported by superficial pyramidal cells – that are the principal contributors to electrophysiological measurements – we can treat the precision weighted prediction errors as event related potentials (ERPs).

To model a placebo condition we used a Gaussian prior on the sensory log precision (Figure 1): $\Pi_v^{(1)} = exp(\gamma) \cdot I$, with $p(\gamma) = N(\mu_{t-1}, \exp(-4))$ where μ_{t-1} is the posterior expectation of the sensory log-precision following the previous trial (starting with a log precision of 4). To model the effects of Galantamine we used a prior with a higher mean and log-precision: $p(\gamma) = N(\mu_{t-1} + 2, \exp(-4+2))$). In other words, we assumed that Galantamine increases both the brain's prior expectation about the precision of sensory input and the confidence about that expectation (Figure 1; upper panel). In terms of the underlying neurobiology, we propose that this increase in sensory precision is mediated by an associated neuromodulatory gain effect (Figure 1; lower panel); where greater ACh levels enhance postsynaptic responses in the superficial layers of primary auditory cortex.

As a point of comparison for the empirical data, we also simulated two alternative ways in which Galantamine might affect sensory precision: the first "straw-man" alternative assumed Galantamine *reduced* sensory precision; i.e., $p(\gamma) = N(\mu_{t-1}-2, \exp(-(4-2)))$. The second alternative assumed that Galantamine had no effect on precision updating. In this case, we assumed that the log precision remained at its initial value, i.e., $p(\gamma) = N(4, \exp(-4))$. Having produced predictions for ERPs in this oddball paradigm, under different assumptions about the effect of Galantamine, we then measured real ERPs using the same paradigm:

Subjects and pharmacological manipulation

We studied 13 right-handed, healthy volunteers (5 female, aged 25 ± 7 yrs), using a within-subject crossover placebo-controlled double-blind design. Subjects attended on two sessions, exactly one week apart. Prior to drug administration, a baseline electrocardiogram (ECG) was performed to exclude cardiac conduction abnormalities. One hour prior to the EEG study, subjects were given either a tablet containing 8 mg of Galantamine or a multivitamin placebo. Tasks began ~45 min after EEG preparation. Galantamine increases the availability of ACh in cholinergic synapses by competitive inhibition of acetylcholinesterase – the enzyme responsible for its breakdown. Galantamine also enhances cholinergic neurotransmission by sensitizing nicotinic receptors (it both increases the probability of channel opening induced by acetylcholine and slows down receptor desensitization) (Coyle and Kershaw 2001, Samochocki, Höffle et al. 2003). Following an oral dose, the peak plasma concentration is attained within 2 hours, and declines with a half-life of about 7 hours (Huang, Lasseter et al. 2002).

Subjects were paid for their participation and consented to all procedures, which were conducted in accordance with the Declaration of Helsinki (1991) and approved by the local Ethics Committee.

Task and EEG acquisition

EEG recordings were made in a quiet dimly-lit room using a 128-channel *Biosemi* system. Electrical signals were digitized at a sampling rate of 512 Hz. Auditory stimuli were presented binaurally over headphones. The stimuli comprised a structured sequence of pure sinusoidal tones, with a roving or sporadically changing frequency (Figure 2). This roving-oddball paradigm (Garrido, Friston et al. 2008) comprises mini-blocks of 6 – 10 tone

repetitions at frequencies of 500, 550, 600, 650, 700, 750 or 800 Hz. At the end of each mini-block the frequency changed pseudo-randomly to another frequency. The duration of each tone was 70 ms, with 5 ms rise and fall times. On average, 180 stimuli for each repetition order were presented to each subject. Subjects were instructed to fixate on a central cross presented on a computer screen and perform an incidental visual task – in which they were required to make a button-press whenever the central fixation cross changed from grey to white or white to grey: this occurred pseudo-randomly every 2 – 5 seconds. Subjects were not asked to attend to the auditory stimuli, since this type of auditory oddball paradigm elicits mismatch potentials that are automatic or pre-attentive (Garrido, Kilner et al. 2009, Näätänen, Astikainen et al. 2010).

Data pre-processing and statistical analysis

EEG data were epoched offline to obtain 400 ms epoch's corresponding to −100 ms to 300 ms peristimulus time. The epoched data were bandpass filtered from 1 – 40 Hz, down-sampled to 200 Hz and re-referenced to the nasion. The data were corrected for artefacts by thresholding (at 100 mV) and averaged according to repetition position in the mini-block. The first presentation, d1, (the deviant) the second presentation d2 and so on, to the last possible presentation d10 (the standard) were averaged over stimulus frequency. The event related responses were subsequently baseline corrected (the analysis routines used for the present study are available in the academic freeware SPM8;http://www.fil.ion.ucl.ac.uk/spm/).

For our sensor level analysis, one electrode was selected from a fronto-central region ((Garrido, Kilner et al. 2007) channel C21). Event related potentials corresponding to trials d1, d2 and d10 were averaged over subjects for each drug condition (Figure 3c) and plotted with the standard error of the mean. A mismatch negativity was seen in the grand-averaged data – most prominently in trial d1 at 150 ms for both placebo and Galantamine conditions. The repetition suppression of the MMN was characterised with a repeated measures ANOVA, using the mean evoked response from 140 to 160 ms as the repeated measure.

Source reconstruction

To define the prior location of electromagnetic sources for a subsequent DCM analysis, multiple sparse priors were used to estimate the cortical source of the mismatch negativity (Friston, Harrison et al. 2008). A tessellated cortical mesh template surface in canonical Montreal Neurological Institute (MNI) anatomical space (http://www.bic.mni.mcgill.ca/brainweb) as implemented in SPM8, served as a brain model to estimate the current source distribution (Mattout, Henson et al. 2007). This dipole mesh was used to calculate the forward solution using a spherical head model. The inverse solution was calculated for deviant (d1) and standard (d10) trials for conditions of placebo and Galantamine separately. Multiple sparse priors employ several hundred patches of activation that are iteratively reduced until an optimal number and location of active patches are found, using a greedy Bayesian search (Friston, Harrison et al. 2008). Source activity measures were then interpolated into MNI voxel space and analysed using statistical parametric mapping – at the between subject level – in the usual way (Kilner, Kiebel et al. 2005): A contrast of evoked signal strength, for standard relative to deviant auditory responses (d10 – d1) was computed

for the placebo condition. This simple main effect of repetition (deviance), was computed at $p < 0.05$ uncorrected – using a two sample *t* test across subjects – and used as an inclusive mask (the MMN mask) for a test of the orthogonal interaction between repetition and drug condition ($[d10 - d1]_{placebo} - [d10 - d1]_{Galantamine}$). Note that we report uncorrected results at $p < 0.05$ because these analyses only served to identify candidate sources in the subsequent DCM analysis. We observed a significant drug effect within the MMN mask in the middle temporal gyri bilaterally (Figure 4b).

Dynamic Causal Modelling

In DCM, event related potentials are modelled as the response of a dynamic input-output system to exogenous (experimental) inputs (David, Kiebel et al. 2006, Kiebel, David et al. 2006). The DCM generates a predicted ERP as the response of a network of coupled sources to sensory (thalamic) input – where each source corresponds to a neural mass model of several neuronal subpopulations. For each source, we used a model based on the canonical microcircuit; Figure 4c (CMC model in DCM for ERPs:http://www.fil.ion.ucl.ac.uk/spm/). This ensemble model is based on a mesoscopic circuit description of cat primary visual cortex, which represents the cardinal features of functional macrocolumns throughout the cortical sheet (Douglas and Martin 2004, Bastos, Usrey et al. 2012). Each source includes four neuronal subpopulations: infra- and supragranular pyramidal cell populations, a granular layer of stellate cells, and an inhibitory interneuron population. Each subpopulation has its own (intrinsic) dynamics modelled by standard neural mass equations. The sources and their interactions are fully specified by these equations that are formally related to neural mass models in the simulation literature (Breakspear, Roberts et al. 2006, Rodrigues, Chizhov et al. 2010). The differential equations describe the evolution of hidden neuronal states (membrane potentials and currents) in the subpopulations comprising each source. The form of these equations can be found in previous papers (David, Kiebel et al. 2006, Garrido, Kilner et al. 2007, Garrido, Friston et al. 2008, Kiebel, Garrido et al. 2008, Garrido, Kilner et al. 2009). In brief, a convolution operation transforms the average density of pre-synaptic input arriving at the population into an average postsynaptic membrane potential; in turn, a sigmoidal function transforms the average membrane potential of the population into its average firing rate. The sensitivity or gain of each subpopulation to its afferents is controlled by a lumped parameter which induces self-inhibition – this parameter is a lumped representation of the effect of several currents, including voltage-gated potassium currents, calcium-gated potassium channels and slow recovery from inactivation of the fast sodium current. These currents cause the firing rate to fall for a given membrane potential mediating a reduction in postsynaptic gain.

These models have been validated using animal electrophysiological recordings (Moran, Stephan et al. 2008, Moran, Jung et al. 2011). Crucially the parameters of the model can be inferred from empirical data, using standard Bayesian inversion techniques (Kiebel at al 2006). This involves supplementing the neuronal model with a conventional electromagnetic forward model to map source activity to sensors (Kiebel, David et al. 2006). Given priors over the source locations, non-invasive ERPs can then be used to estimate the parameters of the neuronal model (Garrido, Kilner et al. 2007). This enables one to test – using pharmacological manipulations – which synaptic or coupling parameters are affected when

changing the levels of specific neurotransmitters, such as increasing the level of acetylcholine with Galantamine in this study. Extrinsic connections in these models follow the same rules on which predictive coding schemes are based: forward connections from superficial pyramidal cells synapse on granular layer IV and deep cell layers, while backward connections from deep pyramidal cells terminate outside layer IV (Felleman and Van Essen 1991).

This CMC model (Figure 4c) was used to model explicitly the laminar specificity of forward and backward connections that are central to predictive coding – where forward connections convey prediction errors to higher levels and backward connections send predictions back to lower areas. The CMC thus allows us to test formal hypotheses, or models, regarding the effects of Galantamine. Namely, the drug could alter (i) the influence of prediction errors on their targets; instantiated as a change in forward connection strengths, or (ii) the precision of prediction errors; instantiated as a change in the gain of superficial pyramidal cells encoding prediction errors or (iii) the influence of predictions; instantiated as a change in the strength of backward connections. We specified a model space along these dimensions, resulting in a model space with 10 competing hypotheses (figure 5A), each of which could potentially explain the differences in the ERP to the largest oddball response (d1) for Galantamine relative to placebo. In comparing these DCMs – using Bayesian model comparison – we hoped to show that Galantamine increased the gain of superficial pyramidal cells early in the auditory hierarchy.

For computational expediency, DCMs were fitted to sensor data following dimensionality reduction to eight channel mixtures or spatial modes. These were the eight principal modes of a singular value decomposition of the channel data between 0 and 250 ms, over trial types of interest. Interestingly, DCMs are inverted using exactly the same variational scheme (variational Laplace) that underlies predictive coding (Friston et al 2007). After model inversion or optimisation, the variational free energy can then used as an approximate log model evidence for Bayesian model comparison.

Results

In brief, we found a remarkable correspondence between simulations of Bayes optimal sensory learning and the empirical responses – at the level of scalp EEG and within the network of sources generating EEG signals. In what follows, we describe simulated (Bayes optimal) responses that, we suppose, underlie mismatch negativity (MMN) responses elicited during the roving auditory oddball paradigm (Haenschel, Vernon et al. 2005). We simulated EEG responses to unexpected or deviant auditory tones that gradually become “standard” tones, under placebo and Galantamine. We then report the empirical findings, in light of the simulations, focusing on the repetition suppression of responses over time – in sensor space – and the underlying synaptic mechanisms – in source space.

Simulating Perception under Placebo and Galantamine

To simulate the processing of auditory stimuli in the roving oddball paradigm, we first generated auditory stimuli, using a model of the auditory environment with two levels of causal dynamics (Equation 3). We then applied Bayesian filtering (Equation 2) to these

sensory data and harvested the precision-weighted prediction errors generated over peristimulus time, for successive tone presentations.

Figure 3a (left) shows the (precision weighted) prediction error in the placebo condition for the first deviant. A large precision-weighted prediction error is observed with a late deflection at around 150 ms (blue curve). This almost disappears on the second trial (green curve), which has a form similar to d10 (red curve). In other words, the first deviant stimulus produces large amplitude prediction errors that are boosted by an inappropriately high precision – established by the preceding sequence of predictable standards. This produces the mismatch negativity. After the first deviant, sensory precision is reduced (by the high prediction errors in the previous trial), rendering subsequent prediction errors less precise and thereby attenuating neuronal responses. These simulation results suggest a type of one-shot learning or repetition suppression – where the brain responds to a new tone with a precise prediction error and then reduces the precision of subsequent error-related responses – as precision is updated and new parameters of the generative model are learned. Under Galantamine (Figure 3a right), the high *a priori* precision on the sensory prediction error increases the response elicited by the first presentation of a new tone (blue curve), and attenuates the reduction of precision on subsequent presentations – leading to larger responses at the second deviant (green curve) and a more protracted repetition suppression of subsequent responses.

To quantify this repetition suppression, we summed the prediction error over 140-160 ms (where MMNs are typically observed) of peristimulus time for each trial and computed the difference for each repetition, compared to d10 (defined as the standard). These differences correspond to a simulated MMN with arbitrary units (a.u.) (Figure 3b). For the placebo MMN, when comparing trial d10 and trial d1, we observe a MMN of 8.1 a.u. that reverses polarity and reduces in size to −1 a.u. on trial d2. The MMN effect then gradually increases again to baseline (−0.1 a.u.) for d9. In contrast, Galantamine increases the amplitude of the MMN and slows its decline (Figure 3b): Comparing trial d10 and trial d1, we observe a difference of magnitude of 34 a.u. which remains positive at 2.5 a.u. on the subsequent trial, d2. The MMN effect decreases slowly to a baseline for d9. It is this quantitative difference in the MMN at the first and subsequent repetitions of a deviant stimulus that we hoped to observe empirically:

Allowing Galantamine to change precision in other ways produced very different response profiles. Our first alternate hypothesis was that Galantamine reduced – rather than increased – sensory precision. This assumption resulted in a reduced MMN response of 2.9 a.u. (d10−d1) that reduces polarity on d2 to −0.5. Our null hypothesis was that Galantamine precluded precision updates. This produced a very different profile, driven by slower changes in extrinsic connections that underlie learning. Here the MMN on the second trial d2 (7.2 a.u.) was greater than on d1 (6.7 a.u.).

Empirical Evoked Brain Responses and the MMN under Placebo and Galantamine

To compare our simulation-based predictions to empirical MMN responses, we examined sensor level evoked responses. We selected a fronto-central electrode (channel C21) typically found to express the MMN and computed the grand-averaged responses for the

placebo and Galantamine conditions (Figure 3c). These responses exhibited the same form of repetition suppression that we saw in the simulations – with trial d2 (green curve; Figure 1c) more similar to d10 (red curves) and a deflection at around 150 ms for the first oddball tone (blue curves). We used the MMN effect (between 140 and 160 ms; defined here for the real data as the deviant – standard) to test for statistical differences between the two drug conditions:

As shown in Figure 3c, we observed an MMN in response to the onset of a deviant tone at a time commensurate with previous studies – at ~150 msec. For each of the 13 subjects, we then quantified nine subject-specific MMNs (averaged over 140 – 160 ms) for each drug state by computing the difference between d10 and d1 to d9 (Figure 3d) as in the simulations. We entered these differences into a repeated-measures ANOVA and found both a significant effect of repetition ($p < 0.001$) and a significant effect of drug ($p < 0.05$). Although we did not observe a significant drug × repetition interaction, we did find a significantly greater MMN under drug compared to placebo on the first and second presentation of a new tone (d1 and d2 relative to d10; $p < 0.05$, two-sample, one-tailed *t* test; Figure 3d). Furthermore, we observed an overall profile for the drug condition that is very similar to the simulated MMN when Galantamine enhanced precision. Our empirical results show an enhanced MMN on Galantamine at d1 compared to placebo that persists at d2, while the placebo MMN reverses in polarity after the first oddball (Figure 3d). These empirical response profiles are inconsistent with either of the alternative models of Galantamine effects (Figure 3b).

Optimal changes in synaptic efficacy and gain

Under our biological implementation of Bayes-optional perception, model parameters (like *C* in Equation 2) are encoded by the strength of extrinsic connections between brain regions; while log-precisions are encoded by the gain of prediction error units (i.e. intrinsic excitation of supragranular pyramidal cells). Our precision estimates show biphasic changes, where the initial oddball coincides with a high precision estimate (d1 = 5.1 a.u.) that undergoes a large reduction (d2 = 2.1 a.u.) and a subsequent rebound (d10 = 4.7 a.u.); Figure 4a (left). From Figure 3a, we see that the corresponding evoked responses reflect these estimates of precision, because they are based on precision-weighted prediction errors. In the Galantamine simulations (Figure 4a – right), however, there was a profound difference in the log-precision updates, compared to placebo. The log-precision at the first deviant was higher (8.1 a.u.) and remained relatively high after the second tone repetition (d2 = 6.1 a.u.), reaching a minimum at d3 (Figure 4a). This difference in the optimisation of (cholinergic) precision or gain is reflected in the simulated evoked responses (Figure 3a), with enhanced and prolonged precision-weighted prediction errors over repetitions, compared to placebo. Neurobiologically, this would correspond to a qualitative difference in the changes in gain of superficial pyramidal cells encoding prediction error – a prediction we assessed empirically using dynamic causal modelling:

DCM results and their relationship to theoretical predictions

In this section, we focus on the synaptic mechanisms producing the repetition suppression in sensor space described above. To address this, we model the generation of evoked responses

in source space using DCM and parameterise this model in terms of coupling strengths and postsynaptic gain that have a direct interpretation in relation to the above simulations.

The locations of electromagnetic (dipolar) sources in our DCM were based on a contrast of source-localised event related potentials that tested for drug effects. After source localisation on a canonical mesh, we tested for a simple main effect of repetition (d10 – d1) in the placebo condition and found bilateral effects in middle temporal gyri (left MTG: $x = -66, y = -30, z = -4$; right MTG: $x = 64, y = -16, z = -18$) and inferior frontal gyri (left IFG: $x = -46, y = 40, z = 2$; right IFG: $x = 36, y = 48, z = -2$), (Figure 4b). Using the results of this contrast as an inclusive mask, we then tested for an interaction of repetition and drug – identifying a significant interaction in middle temporal gyri bilaterally (left MTG: $x = -60, y = -14, z = -18$, right MTG: $x = 62, y = -14, z = -18$; Figure 3b). We used the local maxima of the simple main effect in bilateral frontal cortex and the local maxima of the interaction in bilateral middle temporal gyrus as our prior source locations for the DCM. We also included primary auditory cortex bilaterally (left A1: $x = -42, y = -22, z = 7$; right A1: $x = 46, y = -14, z = 8$), as per Garrido et al. (Garrido, Friston et al. 2008). This ensured that all levels of auditory network processing tones were included (Figure 4). This three-level hierarchy is consistent with previous models of the MMN (Garrido, Kilner et al. 2007, Garrido, Kilner et al. 2007, Garrido, Friston et al. 2008, Garrido, Kilner et al. 2009, Garrido, Kilner et al. 2009). To test for the predicted differences in placebo and Galantamine, we focused on the first presentation of the deviant (d1), because this is not confounded by sensory learning – under the different drug conditions – in subsequent repetitions. Our hypotheses (models) allowed for the Galantamine to exert its effects at different locations in the predictive coding hierarchy: (i) forward connections from superficial pyramidal cells reporting prediction errors; (ii) postsynaptic gain at superficial pyramidal cells encoding the precision of prediction errors and (iii) backward connections from deep pyramidal cells reporting predictions (Figure 5a). For completeness we included models in which Galantamine increased the gain of deep pyramidal cells (blue in Figure 5a). In detail, we compared ten models of grand-averaged ERPs (average seven subjects), where Galantamine changed (bilaterally): the gain of deep pyramidal cells in primary auditory cortex (model 1); the gain of superficial pyramidal cells in primary auditory cortex (model 2); the gain of deep pyramidal cells in MTG (model 3); the gain of superficial pyramidal cells in MTG (model 4); the gain of deep pyramidal cells in IFG (model 5); the gain of superficial pyramidal cells in IFG (model 6); forward connections to MTG (model 7); forward connections to IFG (model 8); backward connections to MTG (model 9) and backward connections to primary auditory cortex (model 10) – see Figure 5a. Figure 5b shows the results of a Bayesian model comparison – showing that the evidence for model 2 is greatest, with very strong evidence relative to the next best model – model 6 (log Bayes Factor = 153). This means that the effect of Galantamine can be best explained as a change in the gain of superficial pyramidal cells in primary auditory cortex. The direction of this effect is shown in Figure 5c and suggests that Galantamine enhances the gain in these bilateral sensory sources with a posterior probability of nearly 100%. The fits to the empirical data under the best model are presented in Figure 5d.

In summary, there is a remarkable convergence between the predictions of Bayes-optimal inference and learning under uncertainty and the empirical results obtained with a roving mismatch negativity paradigm. This convergence was both phenomenologically, in terms of repetition suppression at the sensor level, and mechanistically, in terms of underlying changes in synaptic gain at the source level (as inferred by dynamic causal modelling). These results are consistent with a key role for acetylcholine in modulating the gain of superficial pyramidal cells in early sensory regions that may encode prediction error in predictive coding formulations of hierarchical inference.

Discussion

In this work, we have used a psychopharmacological study to test specific and formal predictions about the effect of cholinergic manipulations on event-related responses and their repetition suppression. These predictions rest on a Bayes-optimal formulation of perceptual inference and learning that can be implemented in a neuronally plausible way. In brief, our results are consistent with a role for ACh in setting the gain of prediction error units implied by predictive coding in the brain. Neurobiologically, our DCMs represent postsynaptic gain via a parameter that determines the intrinsic excitability of a subpopulation (Moran, Kiebel et al. 2007). In predictive coding formulations of free energy minimisation, this gain corresponds to the precision of prediction errors. By assigning ACh the role of signalling sensory precision, we have shown how its augmentation can reduce adaptation to surprising stimuli in sensory cortices. Our simulations and empirical results suggest that when stimuli are repeated in a predictable way, the statistical regularity leads to an adaptive increase in the expected precision of sensory prediction errors. This could be mediated (in part) by increased cholinergic neurotransmission that is suppressed after an unpredicted or surprising (deviant) stimulus. This withdrawal of a neuromodulatory boost leads to a rapid repetition suppression during successive presentations of the same (initially surprising and then increasingly predictable) stimulus. Several intracellular processes may underlie such repetition suppression, including slow after-hyperpolarizing currents (sI_{AHP}) that are mediated by calcium dependent potassium channels and decrease neuronal excitability and firing rate: for a review, see (Faber and Sah 2003). Importantly, acetylcholine modulates this process; e.g., by activation of muscarinic receptors suppresses sI_{AHP} through a cGMP-dependent mechanism (Krause and Pedarzani 2000). This activation may maintain a high postsynaptic gain and facilitate oscillatory activity, with a shift towards higher frequencies (desynchronisation) in the power spectra of neuronal recordings (Liljenstrom and Hasselmo 1995).

In summary, pharmacologically augmenting cholinergic neuromodulation appears to boost event related responses to deviant stimuli and attenuate their subsequent repetition suppression. Using precision-weighted prediction errors from our Bayes optimal (predictive coding) scheme, we could predict suppression of the MMN response under the different drug conditions (Figure 3). Moreover, dynamic causal modelling of a temporo-frontal network suggested that the observed differences in event-related responses can be attributed to cholinergic gain modulation of supragranular pyramidal cells in early auditory cortex (Figure 5). In short, these findings suggest that acetylcholine mediates the representation of precision and acts to facilitate the bottom-up signalling of sensory information (through

precise prediction errors). In the long-term this type of analysis may be applied to empirical data alone, to inform the computational consequences of pathophysiological changes. For example, if a DCM analysis reveals particular changes in directed extrinsic connections or intrinsic gain, one may be able to map this to changes in perceptual prediction error processing; e.g., in neuropsychiatric disorders such as schizophrenia (Marr 2010, Adams, Perrinet et al. 2012).

Previous theoretical circuit level proposals have suggested that ACh supports learning and memory processes by switching the cortex from a 'read-out' to 'read-in' mode – providing an enhancement of sensory-evoked afferent responses and suppression of internal cortical processing (Hasselmo and McGaughy 2004). This proposal is consistent with the current study. Microdialysis measures of ACh in behaving animals support this view – with increases observed in modality-specific sensory cortex and hippocampus during the acquisition of novel, behaviourally-relevant stimuli. Indeed, it has been demonstrated in the auditory domain that a sensory "memory" relies upon cholinergic modulation (Miasnikov, Chen et al. 2008). Similarly, in human studies with cholinergic agonists, the specificity of perceptual learning for behaviourally-relevant features is enhanced (Rokem and Silver 2010), while validity effects in spatial cueing tasks are diminished (Bentley, Husain et al. 2004, Vossel, Thiel et al. 2007, Thiel and Fink 2008), indicative of selective boosting of stimulus-related information. Our analysis of the simulated and empirical responses recapitulates some of these established ideas about the role of cholinergic signalling in perception and attention. Our findings fit comfortably with previous theoretical and empirical accounts; namely, that ACh boosts bottom-up signals in response to uncertainty (Yu and Dayan 2002, Yu and Dayan 2003, Bentley, Husain et al. 2004). In short, there is a wealth of theoretical and empirical evidence for a computational role of ACh in promoting the influence of sensory evidence in perception and attention – which our study supplements with a functional anatomy at the systems level.

One of the key features of the responses elicited by the roving paradigm is a repetition suppression – following a deviant stimulus – that recovers with subsequent presentations (see Figure 3). This profile has been reported previously in intrinsic (within-source) DCM connections exhibiting biphasic changes that were much faster than changes in extrinsic (between-source) connections (Garrido, Kilner et al. 2009). We extend these findings to show that this short-term plasticity may be modulated by ascending neuromodulatory systems. Previous research has shown that MMN responses are enhanced following cholinergic stimulation, particularly nicotinic stimulation in healthy controls (Baldeweg, Wong et al. 2006) and patient populations (Engeland, Mahoney et al. 2002). We simulated these electrophysiological responses by weighting prediction errors – during recognition dynamics – more highly under the influence of Galantamine. This results in exaggerated and prolonged MMN like responses over repetitions of a new stimulus. Crucially, this effect can be simulated easily, when sensory precision remains higher for longer periods.

Given the ability to characterise the putative neuronal implementation of predictive coding schemes using DCM (Garrido, Friston et al. 2008), we tested whether the sensitivity or gain of neuronal populations encoding prediction error was increased in human subjects by pharmacological enhancement of ACh. ACh has long been known to enhance the firing of

cortical pyramidal neurons (Krnjevi and Phillis 1963), inducing a decrease in spike frequency adaptation via several mechanisms – including the reduction of hyperpolarizing potassium currents (Benardo and Prince 1982, Cole and Nicoll 1984, Huang, Morielli et al. 1993). This effect was modelled in our dynamic causal models via a gain parameter that enhanced the excitability of supragranular pyramidal cells (through decreased self inhibition).

Our DCM analysis helps to resolve previous ambiguities regarding the effects of ACh on top-down *vs*. bottom-up processing in human neuroimaging experiments. Using standard univariate analyses, the neural network responsible for these – possibly reciprocal – adjustments under ACh has remained unclear (Thiel and Fink 2008), given the difficulty of specifying how a “bottom-up” enhancement and “top-down” reduction would manifest in measurements of brain responses such as fMRI. Our DCM results suggest that under cholinesterase inhibition, sensory cortices respond to incoming stimuli with exaggerated and prolonged trial-by-trial responses at superficial pyramidal cells – that are the source of bottom-up or forward projections. This enhanced bottom-up effect is mediated via tonic increases in ACh in our simulations and experiments. It is possible that other brain regions, such as the prefrontal cortex or insula, modulate cholinergic afferents – in a phasic fashion – through inputs to the basal forebrain (Sarter, Parikh et al. 2009). This would represent top-down control of facilitation of bottom-up inputs. This type of precision modulation; i.e., state induced changes in sensory precision, can be included in generalised Bayesian schemes to provide a compelling metaphor for attentional modulation (Feldman and Friston 2010).

In general, hierarchical inference using empirical Bayes, as implemented in our simulations, provides a mechanism by which neural circuits can selectively up-weight and down-weight particular features in the sensorium. When making perceptual inferences about the states of the world and their causes, the brain should up-weight reliable sensory evidence by increasing the precision of signals from lower-level cortical regions; e.g., primary sensory areas. For people familiar with statistical procedures this is nothing more than weighted least squares, where noisy data are down weighted. The framework of Yu and Dayan (Yu and Dayan 2002) proposes that ACh modulates uncertainty about high-level representations. We show here an equivalent and complementary effect, where ACh boosts bottom-up sensory evidence in the auditory hierarchy. The resulting adaptive sensory bias may be implemented via gain control in a manner suggested by previous physiological accounts of cholinergic modulation. Of course, our EEG measurements do not allow us to characterise the source of cholinergic modulation *per se* – e.g. between prefrontal regions and basal forebrain cholinergic neurons, which may serve as a route through which enhanced precision is initiated.

In summary, the proposition that neuromodulators encode the precision or predictability of environmental states is supported here by empirical evidence that sensory regions respond with enhanced gain under cholinergic enhancement – producing exaggerated and prolonged mismatch responses that localise to superficial pyramidal populations encoding prediction error (Mumford 1992).

Acknowledgments

RJM was supported by a Max Planck award to RJD. KJF, RJD and MS were supported by the Wellcome Trust. KES was supported by Neurochoice (SystemsX.ch) and an endowed ETH Chair from the René and Susanne Braginsky Foundation. PC was supported by a Ramon y Cajal Fellowship from the Spanish Ministry of Science and Innovation (RYC-2010-05748).

References

Adams RA, Perrinet LU, Friston K. Smooth Pursuit and Visual Occlusion: Active Inference and Oculomotor Control in Schizophrenia. PloS one. 2012; 7(10):e47502. [PubMed: 23110076]

Baldeweg T, Wong D, Stephan KE. Nicotinic modulation of human auditory sensory memory: evidence from mismatch negativity potentials. International Journal of Psychophysiology. 2006; 59(1):49–58. [PubMed: 16313986]

Bar M. A cortical mechanism for triggering top-down facilitation in visual object recognition. Journal of Cognitive Neuroscience. 2003; 15(4):600–609. [PubMed: 12803970]

Bastos AM, Usrey WM, Adams RA, Mangun GR, Fries P, Friston KJ. Canonical Microcircuits for Predictive Coding. Neuron. 2012; 76(4):695–711. [PubMed: 23177956]

Benardo LS, Prince DA. Cholinergic excitation of mammalian hippocampal pyramidal cells. Brain research. 1982; 249(2):315–331. [PubMed: 6291715]

Benarroch EE. Acetylcholine in the cerebral cortex. Neurology. 2010; 75(7):659–665. [PubMed: 20713954]

Bentley P, Husain M, Dolan R. Effects of cholinergic enhancement on visual stimulation, spatial attention, and spatial working memory. Neuron. 2004; 41(6):969–982. [PubMed: 15046728]

Breakspear M, Roberts JA, Terry JR, Rodrigues S, Mahant N, Robinson PA. A unifying explanation of primary generalized seizures through nonlinear brain modeling and bifurcation analysis. Cereb Cortex. 2006; 16(9):1296–1313. [PubMed: 16280462]

Cauller L. Layer I of primary sensory neocortex: where top-down converges upon bottom-up. Behavioural brain research. 1995; 71(1):163–170. [PubMed: 8747184]

Cole AE, Nicoll RA. The pharmacology of cholinergic excitatory responses in hippocampal pyramidal cells. Brain research. 1984; 305(2):283–290. [PubMed: 6331600]

Coyle J, Kershaw P. Galantamine, a cholinesterase inhibitor that allosterically modulates nicotinic receptors: effects on the course of Alzheimer's disease. Biological psychiatry. 2001; 49(3):289–299. [PubMed: 11230880]

Dani JA, Bertrand D. Nicotinic acetylcholine receptors and nicotinic cholinergic mechanisms of the central nervous system. Annu. Rev. Pharmacol. Toxicol. 2007; 47:699–729. [PubMed: 17009926]

David O, Kiebel S, Harrison L, Mattout J, Kilner J, Friston K. Dynamic causal modeling of evoked responses in EEG and MEG. NeuroImage. 2006; 30(4):1255–1272. [PubMed: 16473023]

Deco G, Thiele A. Attention–oscillations and neuropharmacology. European journal of neuroscience. 2009; 30(3):347–354. [PubMed: 19614749]

Douglas RJ, Martin KAC. Neuronal circuits of the neocortex. Annu. Rev. Neurosci. 2004; 27:419–451. [PubMed: 15217339]

Engeland C, Mahoney C, Mohr E, Ilivitsky V, Knott VJ. Acute nicotine effects on auditory sensory memory in tacrine-treated and nontreated patients with Alzheimer's disease: an event-related potential study. Pharmacology Biochemistry and Behavior. 2002; 72(1):457–464.

Faber ESL, Sah P. Calcium-activated potassium channels: multiple contributions to neuronal function. The Neuroscientist. 2003; 9(3):181–194. [PubMed: 15065814]

Fadel JR. Regulation of cortical acetylcholine release: insights from in vivo microdialysis studies. Behavioural Brain Research. 2011; 221(2):527–536. [PubMed: 20170686]

Feldman H, Friston KJ. Attention, uncertainty, and free-energy. Frontiers in Human Neuroscience. 2010; 4

Felleman DJ, Van Essen DC. Distributed hierarchical processing in the primate cerebral cortex. Cerebral Cortex. 1991; 1(1):1. [PubMed: 1822724]

Friston K. A theory of cortical responses. Philosophical Transactions of the Royal Society B: Biological Sciences. 2005; 360(1456):815–836.

Friston K. Hierarchical models in the brain. PLoS computational biology. 2008; 4(11):e1000211. [PubMed: 18989391]

Friston K. The free-energy principle: a rough guide to the brain? Trends in cognitive sciences. 2009; 13(7):293–301. [PubMed: 19559644]

Friston K. The free-energy principle: a unified brain theory? Nature Reviews Neuroscience. 2010; 11(2):127–138.

Friston K,L, Harrison J, Daunizeau S, Kiebel C, Phillips N, Trujillo-Barreto R, Henson G. Flandin, Mattout J. Multiple sparse priors for the M/EEG inverse problem. NeuroImage. 2008; 39(3):1104–1120. [PubMed: 17997111]

Friston K, Kiebel S. Predictive coding under the free-energy principle. Philosophical Transactions of the Royal Society B: Biological Sciences. 2009; 364(1521):1211.

Garrido M, Friston K, Kiebel S, Stephan K, Baldeweg T, Kilner J. The functional anatomy of the MMN: a DCM study of the roving paradigm. NeuroImage. 2008; 42(2):936–944. [PubMed: 18602841]

Garrido M, Kilner J, Kiebel S, Friston K. Evoked brain responses are generated by feedback loops. Proceedings of the National Academy of Sciences. 2007; 104(52):20961.

Garrido M, Kilner J, Stephan K, Friston K. The mismatch negativity: a review of underlying mechanisms. Clinical Neurophysiology. 2009; 120(3):453–463. [PubMed: 19181570]

Garrido MI, Kilner JM, Kiebel SJ, Friston KJ. Evoked brain responses are generated by feedback loops. Proceedings of the National Academy of Sciences. 2007; 104(52):20961–20966.

Garrido MI, Kilner JM, Kiebel SJ, Stephan KE, Baldeweg T, Friston KJ. Repetition suppression and plasticity in the human brain. NeuroImage. 2009; 48(1):269–279. [PubMed: 19540921]

Garrido MI, Kilner JM, Kiebel SJ, Stephan KE, Friston KJ. Dynamic causal modelling of evoked potentials: a reproducibility study. NeuroImage. 2007; 36(3):571–580. [PubMed: 17478106]

Ghahramani Z, Beal MJ. Variational inference for Bayesian mixtures of factor analysers. Advances in neural information processing systems. 2000; 12:449–455.

Giovannini M, Rakovska A, Benton R, Pazzagli M, Bianchi L, Pepeu G. Effects of novelty and habituation on acetylcholine, GABA, and glutamate release from the frontal cortex and hippocampus of freely moving rats. Neuroscience. 2001; 106(1):43–53. [PubMed: 11564415]

Haenschel C, Vernon DJ, Dwivedi P, Gruzelier JH, Baldeweg T. Event-related brain potential correlates of human auditory sensory memory-trace formation. The journal of Neuroscience. 2005; 25(45):10494–10501. [PubMed: 16280587]

Hasselmo ME, McGaughy J. High acetylcholine levels set circuit dynamics for attention and encoding and low acetylcholine levels set dynamics for consolidation. Progress in brain research. 2004; 145:207–231. [PubMed: 14650918]

Hasselmo ME, Sarter M. Modes and models of forebrain cholinergic neuromodulation of cognition. Neuropsychopharmacology. 2010; 36(1):52–73. [PubMed: 20668433]

Hemholtz, H. v. Handbuch der physiologische Optik III. Leipzig: 1866. Die Lehre von den Gesichtswahrnehmungen.

Huang F, Lasseter KC, Janssens L, Verhaeghe T, Lau H, Zhao Q. Pharmacokinetic and safety assessments of galantamine and risperidone after the two drugs are administered alone and together. The Journal of Clinical Pharmacology. 2002; 42(12):1341.

Huang XY, Morielli AD, Peralta EG. Tyrosine kinase-dependent suppression of a potassium channel by the G protein-coupled m1 muscarinic acetylcholine receptor. Cell. 1993; 75(6):1145–1156. [PubMed: 8261514]

Kiebel S, David O, Friston K. Dynamic causal modelling of evoked responses in EEG/MEG with lead field parameterization. NeuroImage. 2006; 30(4):1273–1284. [PubMed: 16490364]

Kiebel S, Garrido M, Moran R, Friston K. Dynamic causal modelling for EEG and MEG. Cognitive Neurodynamics. 2008; 2(2):121–136. [PubMed: 19003479]

Kiebel SJ, Daunizeau J, Friston KJ. Perception and hierarchical dynamics. Frontiers in neuroinformatics. 2009; 3

Kiebel SJ, David O, Friston KJ. Dynamic causal modelling of evoked responses in EEG/MEG with lead field parameterization. NeuroImage. 2006; 30(4):1273–1284. [PubMed: 16490364]

Kilner JM, Kiebel SJ, Friston KJ. Applications of random field theory to electrophysiology. Neuroscience letters. 2005; 374(3):174–178. [PubMed: 15663957]

Krause M, Pedarzani P. A protein phosphatase is involved in the cholinergic suppression of the Ca 2+-activated K+ current s IAHP in hippocampal pyramidal neurons. Neuropharmacology. 2000; 39(7):1274–1283. [PubMed: 10760369]

Krnjevi K, Phillis J. Acetylcholine-sensitive cells in the cerebral cortex. The Journal of physiology. 1963; 166(2):296. [PubMed: 14035889]

Lamme VA. The neurophysiology of figure-ground segregation in primary visual cortex. The Journal of Neuroscience. 1995; 15(2):1605–1615. [PubMed: 7869121]

Liljenstrom H, Hasselmo ME. Cholinergic modulation of cortical oscillatory dynamics. Journal of neurophysiology. 1995; 74(1):288–297. [PubMed: 7472331]

Marr D. Vision: A computational investigation into the human representation and processing of visual information. 2010

Mattout J, Henson R, Friston K. Canonical source reconstruction for MEG. Computational Intelligence and Neuroscience. 2007; 17 2007.

Miasnikov AA, Chen JC, Weinberger NM. Specific auditory memory induced by nucleus basalis stimulation depends on intrinsic acetylcholine. Neurobiology of Learning and Memory. 2008; 90(2):443–454. [PubMed: 18573347]

Moran R, Kiebel S, Stephan K, Reilly R, Daunizeau J, Friston K. A neural mass model of spectral responses in electrophysiology. NeuroImage. 2007; 37(3):706–720. [PubMed: 17632015]

Moran R, Stephan K, Kiebel S, Rombach N, O'Connor W, Murphy K, Reilly R, Friston K. Bayesian estimation of synaptic physiology from the spectral responses of neural masses. NeuroImage. 2008; 42(1):272–284. [PubMed: 18515149]

Moran RJ, Jung F, Kumagai T, Endepols H, Graf R, Dolan RJ, Friston KJ, Stephan KE, Tittgemeyer M. Dynamic Causal Models and Physiological Inference: A Validation Study Using Isoflurane Anaesthesia in Rodents. PloS one. 2011; 6(8):e22790. [PubMed: 21829652]

Mumford D. On the computational architecture of the neocortex. Biological Cybernetics. 1992; 66(3): 241–251. [PubMed: 1540675]

Näätänen R, Astikainen P, Ruusuvirta T, Huotilainen M. Automatic auditory intelligence: An expression of the sensory–cognitive core of cognitive processes. Brain research reviews. 2010; 64(1):123–136. [PubMed: 20298716]

Nelson C, Sarter M, Bruno J. Prefrontal cortical modulation of acetylcholine release in posterior parietal cortex. Neuroscience. 2005; 132(2):347–359. [PubMed: 15802188]

Pepeu G, Giovannini MG. Changes in acetylcholine extracellular levels during cognitive processes. Learning & Memory. 2004; 11(1):21–27. [PubMed: 14747513]

Rao RPN, Ballard DH. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. nature neuroscience. 1999; 2:79–87.

Rodrigues S, Chizhov AV, Marten F, Terry JR. Mappings between a macroscopic neural-mass model and a reduced conductance-based model. Biol Cybern. 2010; 102(5):361–371. [PubMed: 20306202]

Rokem A, Silver MA. Cholinergic enhancement augments magnitude and specificity of visual perceptual learning in healthy humans. Current Biology. 2010; 20(19):1723–1728. [PubMed: 20850321]

Samochocki M, Höffle A, Fehrenbacher A, Jostock R, Ludwig J, Christner C, Radina M, Zerlin M, Ullmer C, Pereira EFR. Galantamine is an allosterically potentiating ligand of neuronal nicotinic but not of muscarinic acetylcholine receptors. Journal of Pharmacology and Experimental Therapeutics. 2003; 305(3):1024–1036. [PubMed: 12649296]

Sarter M, Hasselmo ME, Bruno JP, Givens B. Unraveling the attentional functions of cortical cholinergic inputs: interactions between signal-driven and cognitive modulation of signal detection. Brain Research Reviews. 2005; 48(1):98–111. [PubMed: 15708630]

Sarter M, Parikh V, Howe WM. Phasic acetylcholine release and the volume transmission hypothesis: time to move on. Nature Reviews Neuroscience. 2009; 10(5):383–390.

Thiel C, Fink G. Effects of the cholinergic agonist nicotine on reorienting of visual spatial attention and top-down attentional control. Neuroscience. 2008; 152(2):381–390. [PubMed: 18272290]

von Helmholtz, H. Helmholtz's treatise on physiological optics. The Optical Society of America; 1962.

Vossel S, Thiel CM, Fink GR. Behavioral and neural effects of nicotine on visuospatial attentional reorienting in non-smoking subjects. Neuropsychopharmacology. 2007; 33(4):731–738. [PubMed: 17551539]

Yu A, Dayan P. Expected and unexpected uncertainty: ACh and NE in the neocortex. Advances in neural information processing systems. 2003:173–180.

Yu AJ, Dayan P. Acetylcholine in cortical inference. Neural Networks. 2002; 15(4-6):719–730. [PubMed: 12371522]

Yu AJ, Dayan P. Uncertainty, neuromodulation, and attention. Neuron. 2005; 46(4):681–692. [PubMed: 15944135]

Zipser K, Lamme VAF, Schiller PH. Contextual modulation in primary visual cortex. The Journal of Neuroscience. 1996; 16(22):7376–7389. [PubMed: 8929444]

Europe PMC Funders Author Manuscripts

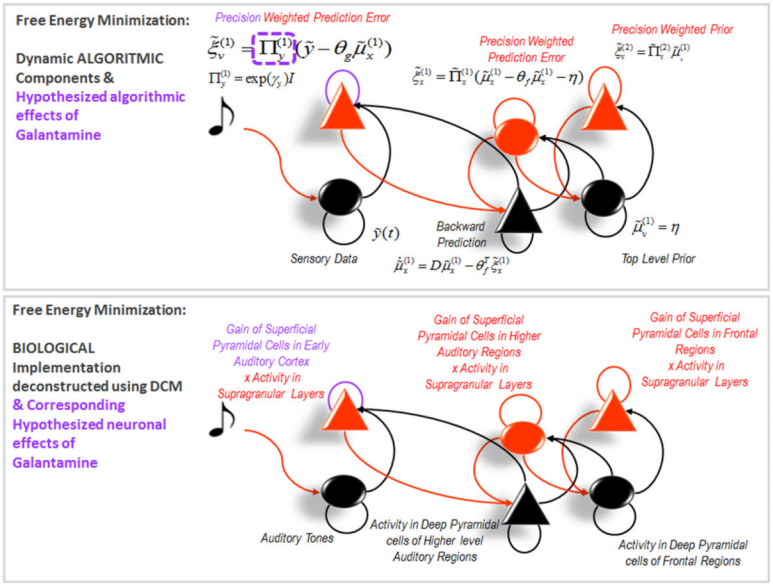

Figure 1.

Simulating predictive coding through free energy minimisation: Algorithmic Components & Proposed Neurobiological Implementation

Upper Panel: free energy, $F(y,\mu) = -\ln p(y|m) + D[q(x,v,\theta,\gamma|\mu)\ ||p(x,v,\theta,\gamma|y,m)]$ bounds the surprise about sensory inputs $y(t)$ and – when optimised with respect to some sufficient statistics $\mu(t)$ – yields a recognition density $q(x, v,\theta,\gamma|\mu)$ over unknown quantities that approximates the true posterior density $p(x,v,\theta,\gamma|y,m)$, under some generative model $m \supset \{f, g\}$. The sufficient statistics (generally posterior means) of unknown quantities are obtained in a relatively straightforward way using a (variational Bayes) coordinate descent where, optimising the unknown variables becomes (generalised) Bayesian filtering. This leads to update equations similar to variational expectation maximization (Ghahramani and Beal 2000) but equipped with a moving frame of reference for the time varying states and causes: see (Friston 2008) for details. Lower panel: the ensuing Bayesian filtering or predictive coding scheme suggests two distinct populations – corresponding to state units and error units, where changes in the activity of state units (encoding predictions) are linear functions of precision-weighted prediction errors. This linearity suggests that precision weighted prediction errors are conveyed by driving forward connections (e.g., mediated by AMPA receptors). Conversely, the prediction errors are nonlinear functions of top-down and lateral predictions from the state units, which implies modulatory backward connections (e.g., mediated by a mixture of NMDA, GABA and AMPA receptors) (Friston 2005). Crucially, the potency of ascending prediction errors depends upon the posterior expectations about precision, which act to control the postsynaptic gain of prediction error units. We will test whether cholinergic modulation by Galantamine enhances the precision of prediction errors in early sensory cortex. This proposition has been motivated theoretically in (Friston 2008, Friston 2009, Feldman and Friston 2010). The DCM used in this paper distinguishes activity in two subpopulations of pyramidal cells (supra and infragranular) and their respective excitability or gain.

Europe PMC Funders Author Manuscripts

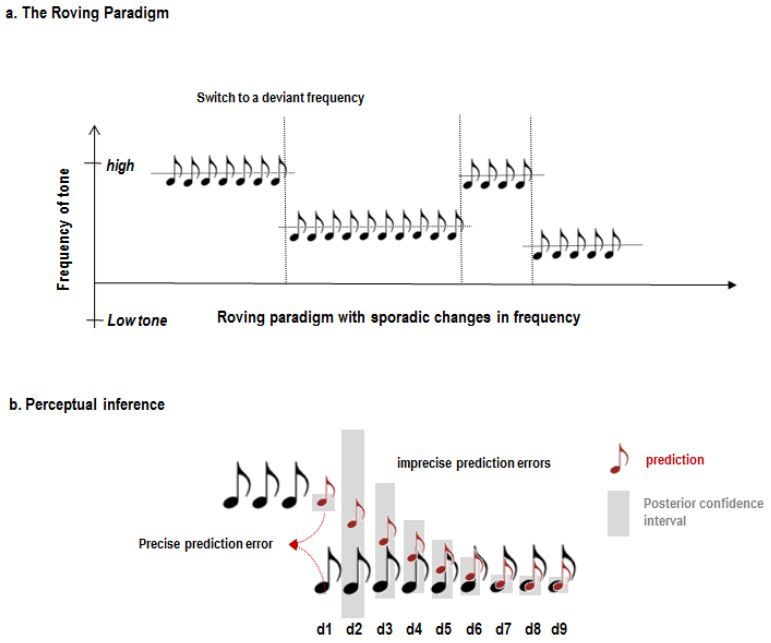

Figure 2.

Paradigm and Perception

This schematic illustrates the roving oddball paradigm (**a**) used in both the simulations and empirical study. After the presentation of a deviant tone, the expected precision of sensory information falls due to the profound prediction errors elicited by an unpredicted (deviant) stimulus – d1. With subsequent repetitions of the same stimulus, the stimulus becomes more predictable and confidence or precision recovers. Posterior confidence is illustrated with the grey bars in the lower panel (**b**) and is determined by the sensory precision expected. Crucially, event related responses are assumed to report precision-weighted prediction errors. This means that the monotonically decreasing prediction error – with repetition – is modulated by precision to produce a characteristic mismatch negativity (precise prediction error) that almost disappears after presentation of the first deviant.

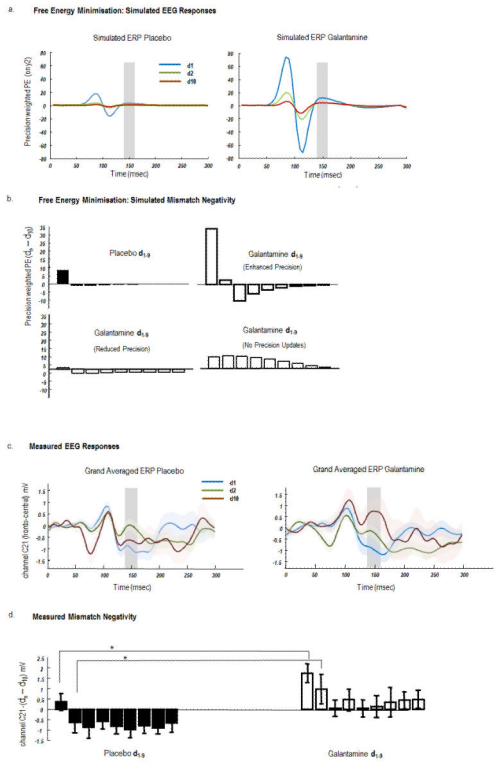

Figure 3.

EEG Responses: Simulated and Measured

a. Simulated ERPs evoked by tones whose frequency deviates from preceding tones (d1) and are repeated until the tenth presentation (d10). These ERPs are the precision weighted prediction errors (Figure 1) of sensory data encoding the frequency of the tone (y_2). Left: Simulated precision-weighted prediction error under placebo (illustrated are d1: blue, d2: green and d10: red). The agent learns the tone frequency over successive repetitions, resulting in a reduction in the size of the evoked responses. Right: A similar profile is observed under Galantamine, with elevated priors on sensory precision. Evoked responses are higher in magnitude and sustained for longer (d2). **b**. Top panel: Simulated mismatch response (d10 – dn), where d10 is set as the standard and d1-d9, a parametric deviant. The MMN is taken from the ERPs illustrated in **a**. It is simply the difference between simulated evoked responses (precision-weighted prediction errors) between standards and deviants summed from 140-160 msec (shaded areas). Left: the placebo MMN shows a rapid one-shot learning, with a smaller MMN on d1 and a reversal in MMN polarity on d2, which returns to close to 0 at d9. Right: Galantamine MMNs are prolonged and have greater magnitude for all trials. Bottom panel: Two alternate MMN effects under different Galantamine models: (left) where Galantamine reduced sensory precision and (right) where Galantamine prevents precision updating over trials. **c**. Scalp EEG measurements of auditory evoked potentials. Grand averaged waveforms from a single fronto-central electrode (C21), for the presentation of the first deviant tone in a sequence (d1), second tone (d2) and final tone (d10), averaged across tones of all frequencies, under placebo (left) and Galantamine (right). MMN effects are evident in both drug conditions ~150msec. **d**. The MMN effect across all 9 tone repetitions (d10 – dn). The MMN effect was significantly different across repetitions and

Europe PMC Funders Author Manuscripts

drug state. In particular both d1 and d2 induced MMNs were greater on Galantamine than on placebo.

Europe PMC Funders Author Manuscripts

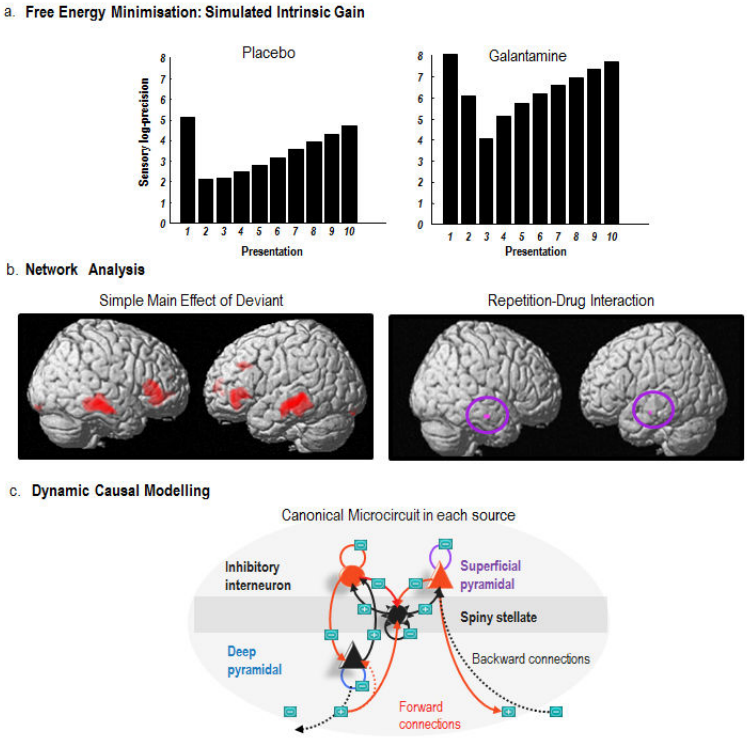

Figure 4.

Synaptic effects: Simulated and Measured

a. Simulated changes in precision parameter estimates. Left: the updates to precision parameters that encode sensory log-precision, exhibit a phasic learning profile, with high values on the first deviant presentation that are attenuated markedly by the second presentation. The third and successive presentations result in increased precision as the agent becomes more certain of its predictions. Right: In the Galantamine simulations, higher precisions are encoded with higher certainty and exhibit a different optimisation, whereby high precision during the first oddball drops on the second tone, but to a level greater than under placebo. **b**. Left: Source localised simple main effect of deviant (d10-d1), rendered onto a canonical cortical surface and thresholded at $p<0.05$ uncorrected. The local maxima in frontal vertices were used as prior source locations in the DCM analysis. The image served as a mask on tests for drug × repetition interactions. Right: bilateral activation was found in middle temporal gyrus ($p<0.05$ uncorrected within mask) for the interaction. **c.** The canonical microcircuit used to model intrinsic connections among subpopulations within each source in the neural mass model. Purple and blue highlight the intrinsic connections that were modulated by Galantamine in the dynamic causal modelling reported the next figure.

Europe PMC Funders Author Manuscripts

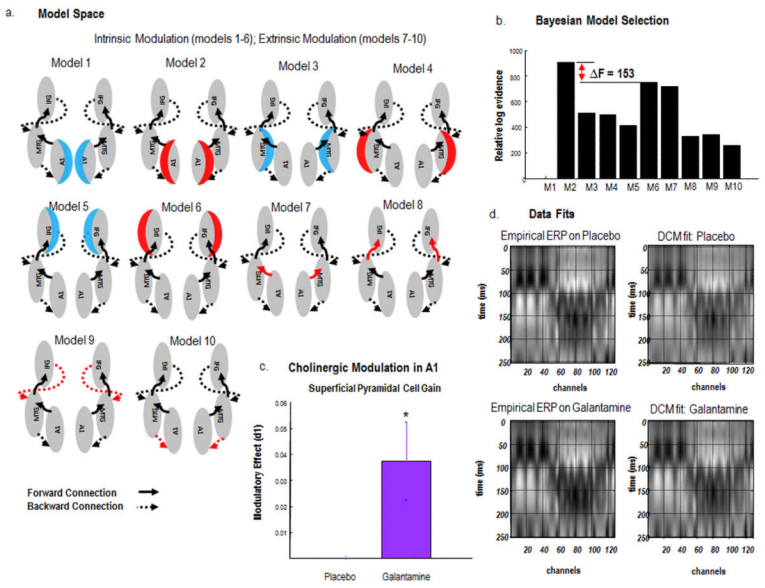

Figure 5.

Network Modulation by Galantamine

a. Models tested to discover where Galantamine boosted event related responses. These modelled an effect of Galantamine on either the gain of deep pyramidal cells (blue shading), the gain of superficial pyramidal cells (red shading) or extrinsic long-range connections (red arrows). **b**. Bayesian model comparison revealed very strong evidence in favour of model 2 (compared to second best performing model). Model 2 model the effects of Galantamine as a gain modulation of supragranular pyramidal cells in bilateral primary auditory cortex. **c**. The direction of the gain effect – under Galantamine relative to placebo – shows enhanced gain at supragranular pyramidal cells with a posterior probability = 1 (shown with 95% Bayesian confidence intervals). **d.** Fitting the DCM to empirical data shows a high correspondence across peristimulus time and channels for both placebo and galantamine responses.