

Acetylcholine in Cortical Inference

Angela J Yu Peter Dayan

Gatsby Computational Neuroscience Unit
University College London

Acknowledgements

We are very grateful to Mike Hasselmo, David Heeger, Sham Kakade, Szabolcs Káli, Martin Sarter, and especially Yair Manor for helpful comments. Funding was provided by the Gatsby Charitable Foundation and the NSF Graduate Research Fellowship Program.

Address for correspondence:

Angela Yu
Gatsby Computational Neuroscience Unit
17 Queen Square
London WC1N 3AR
United Kingdom
phone: +44 (0) 20 7679 1195
fax: +44 (0) 20 7679 1173
email: feraina@gatsby.ucl.ac.uk

Running title: *ACh in Cortical Inference*

Section: **Neuroscience and Neuropsychology**

Acetylcholine in Cortical Inference

Abstract

Acetylcholine (ACh) plays an important role in a wide variety of cognitive tasks, such as perception, selective attention, associative learning, and memory. Extensive experimental and theoretical work in tasks involving learning and memory has suggested that ACh reports on unfamiliarity and controls plasticity and effective network connectivity. Based on these computational and mechanistic insights, we develop a theory of cholinergic modulation in perceptual inference. We propose that ACh levels reflect the *uncertainty* associated with top-down information, and have the effect of modulating the interaction between top-down and bottom-up processing in determining the appropriate neural representations for inputs. We illustrate our proposal by means of an hierarchical hidden Markov model, showing that cholinergic modulation of contextual information leads to appropriate perceptual inference.

keywords: acetylcholine, perception, neuromodulation, representational inference, hidden Markov model, temporal context, selective attention, memory

1 Introduction

Neuromodulators such as acetylcholine (ACh), serotonin, dopamine, norepinephrine, and histamine play two characteristic roles. One, most studied in vertebrate systems, concerns the control of plasticity. The other, most studied in invertebrate systems, concerns the control of network responses. For instance, a single, recurrently connected, assembly of neurons can exhibit multiple dynamical modes (Pflüger, 1999), as neuromodulators alter the excitabilities of individual neurons and the amplitudes of synaptic potentials (Marder, 1998). These two roles have also been brought together, notably in the theoretical and experimental studies of Hasselmo and his colleagues (Hasselmo & Bower, 1993; Hasselmo, 1995) into the neuromodulatory control of plasticity in recurrently connected neural networks. This work sits with that on dopamine (*eg* Schultz, Dayan & Montague, 1997) in proposing *computationally* specific roles for neuromodulation.

A significant proportion of Hasselmo's work has concerned ACh. This neuromodulator is delivered to the cortex from a small number of nuclei in the basal forebrain (BF): medial septum (MS), diagonal band of Broca (DBB), and nucleus basalis (NBM). Physiological studies on ACh indicate that its neuromodulatory effects at the cellular level are diverse, causing direct hyperpolarization and depolarization as well as synaptic facilitation and suppression, all within the same cortical area (Kimura, Fukuda, & Tsumoto, 1999). ACh is involved in behavioral tasks designed to test a wide variety of cognitive functions, such as perception, selective attention, associative learning, and memory (Everitt & Robbins, 1997; Hasselmo, 1995; Holland, 1997).

Hasselmo and his colleagues (Hasselmo & Bower, 1993; Hasselmo, 1995) focused on neuromodulatory influences over learning and memory in the hippocampus and cortex. They proposed that cholinergic (and perhaps other) neuromodulation controls read-in to, and read-out from, recurrently-connected, attractor-like memories, such as that in area CA3 of the hippocampus. Such attractor networks (Amit, 1989) fail if the recurrent connections are operational during storage, since new memories lose their specific identity by being forced to map onto existing memories retrieved through the recurrent dynamics. Hasselmo and colleagues suggested that cholinergic neuromodulation during storage could selectively suppress the recurrent connections (and perhaps the perforant path connections) onto CA3 cells and selectively boost the feedforward mossy fiber inputs from the dentate gyrus. During recall, the recurrent connections should play a fuller part, allowing associative retrieval. The degree of ACh release would reflect the unfamiliarity of the input, and thereby act as a gate to learning. This mechanism has been widely adopted, for instance in our own work, by Káli & Dayan (2000) to understand how spatial place cells in CA3 might result from a learned surface attractor network.

Hasselmo and colleagues demonstrated physiologically how septal ACh might achieve this putative modulatory function. During the learning of a new memory, ACh prevents retrograde interference by selectively suppressing recurrent and feedback synaptic transmission, and enhances response to feedforward inputs by decreasing pyramidal cell adaptation. During recall, a lower ACh level boosts recurrent synaptic transmission and suppresses recurrent synaptic plasticity, thus allowing the network to settle into a stored pattern. Hasselmo and his col-

leagues have also demonstrated that ACh has similar physiological and functional effects in the piriform cortex (Linster & Hasselmo, 2001), which is important for olfactory memory.

Lesions studies in classical conditioning tasks provide additional insight into ACh's role in the cortex. Animals are known to learn faster about stimuli about whose consequences the animals are uncertain (Pearce & Hall, 1980). Through an extensive series of selective lesion experiments in rats, Holland and his colleagues have demonstrated that the cholinergic projection from the nucleus basalis magnocellularis (nucleus basalis of Meynert in primates) to the parietal cortex is essential for this sort of faster learning (Holland, 1997; Holland & Gallagher, 1999). These data have been interpreted, using the theoretical viewpoint of statistical learning models, as implying that the ACh signal reports the *unfamiliarity* of the stimuli, or the *uncertainty* in its predictions (Dayan, Kakade, & Montague, 2000).

In this paper, we present a theory of cortical cholinergic function in perceptual inference based on combining the physiological evidence that ACh can differentially modulate synaptic transmission to control states of cortical dynamics, together with theoretical ideas about the information carried by the ACh signal. Crudely speaking, perception involves the inference of the most appropriate representation for sensory inputs. This inference is influenced by both top-down inputs, providing contextual information, and bottom-up inputs from sensory processing. We propose that ACh reports on the uncertainty associated with top-down information, and has the effect of modulating the relative strengths of these two input sources. Many cognitive functions affected by ACh levels can be recast in the conceptual framework of representational inference.

In section 2, we present a simple hierarchical HMM model that casts sensory perception in the theoretical framework of representational inference. As we demonstrate in section 3, approximate inference in such a model could be mediated by cortical cholinergic innervation. A summary of relevant experimental data and proposals for new experiments is presented in section 4.

2 Hidden Markov Models and Perceptual Inference

Inferring appropriate representations for the constant stream of sensory inputs is a formidable task, largely because of the inherent ambiguity and noise in the sensory input. A vital source of information that helps resolve ambiguities comes from temporal and spatial context, and thus a key issue for perceptual inference is updating and maintaining this top-down contextual information, and using it correctly in concert with bottom-up information from the sensory input (Helmholtz, 1896; Neisser, 1967; Grenander, 1995).

For simplicity, we consider the most basic form of top-down contextual information, namely that coming from the recent past. That is, we consider a series of sensory inputs whose internal representations are individually ambiguous. Disambiguation comes via top-down information based on a slowly-changing overall state of the environment. Here, only temporal

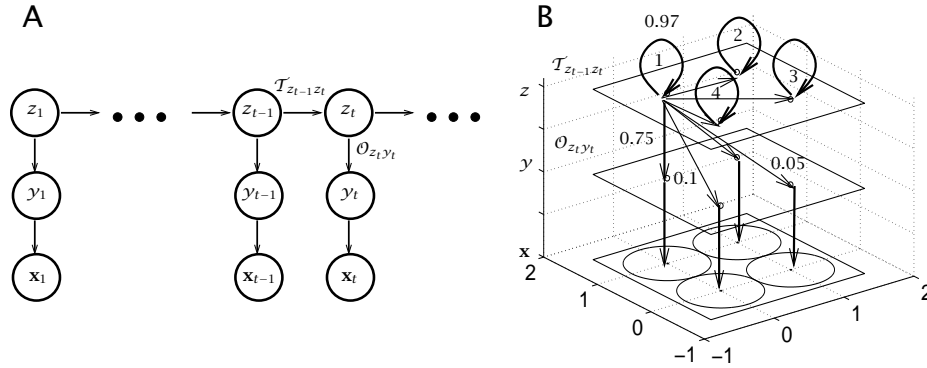


Figure 1: Hierarchical HMM. A) Three-layer model, with two hidden layers, z and y , and one observed layer, x . The temporal dynamics are captured by the transition matrix $\mathcal{T}_{z_{t-1}z_t}$ in the z layer, and the observations x are generated from y and, indirectly, from z . B) Example parameter settings: $z \in \{1 - 4\} \Rightarrow y \in \{1 - 4\} \Rightarrow x \in \mathcal{R}^2$ with dynamics (\mathcal{T}) in the z layer ($P[z_t = z_{t-1}] = 0.97$), a probabilistic mapping (\mathcal{O}) from $z \rightarrow y$ ($P[y_t = z_t | z_t] = 0.75$), and a Gaussian model $p[x|y]$ with means at the corners of the unit square and standard deviation $\sigma = 0.5$ in each direction. Only some of the links are shown to reduce clutter.

context is relevant; there is no spatial context. The resulting model (see also Becker, 1999) is a form of Hidden Markov Model (HMM). The HMM captures the way that sensory inputs are generated or synthesized. We consider the inferential task of recognition or analysis in which the representation for each input is determined. We compare an approximate model based on cholinergic neuromodulation with the exact model, which, in this case, is computationally tractable (Rabiner, 1989).

Our HMM (figure 1A;B) consists of three pieces. One, z_t , is the overall state of the environment at time t , which we also call the context. Changes to z_t are stochastically controlled by a *transition matrix* $\mathcal{T}_{z_{t-1}z_t}$, whose entries ensure that the context changes rarely. The second piece is y_t , which is determined stochastically on each time-step, in a way that depends on the current state of the environment. The third piece, the observed input x_t , depends stochastically on y_t . The inferential task is to represent inputs x in terms of the y values that were responsible for them. However, the relationship between y_t and x_t is such that this is ambiguous, so top-down information from the likely states of z_t , *ie* the likely context, is important to find the correct representation for x_t .

Figure 1A shows the probabilistic contingencies among the variables. Figure 1B shows the same contingencies in a different way, and specifies the particular setting of parameters used to generate the examples found in the remainder of the paper.

More formally, the context is a discrete, hidden, random variable z_t , whose stochastic temporal dynamics are described by a Markov chain with transition matrix $\mathcal{T}_{z_{t-1}z_t}$, where

$$P[z_t | z_{t-1}] \equiv \mathcal{T}_{z_{t-1}z_t} = \begin{cases} y & \text{if } z_t = z_{t-1} \\ \frac{1-y}{n_z-1} & \text{otherwise} \end{cases} \quad (1)$$




Figure 2: Generative model. A sample sequence involving 400 time steps, generated from the model shown in Figure 1B. Note the slow dynamics in z , the stochastic mapping into y , and substantial overlap in x 's generated from the different y 's (different symbols correspond to different Gaussians shown in Figure 1B).

where n_z is the number of all possible states of z , and y is the probability of persisting in one context. When y is close to 1, as is the case in the example of Figure 1B, the context tends to remain the same for a long time. When y is close to 0, the context tends to switch among the different states of z rapidly and randomly. The state of the second hidden layer, y is generated from z with the mapping $\mathcal{O}_{z_t y_t}$, which specifies $P[y_t | z_t]$, and controls which of a set of circular two dimensional Gaussians is used to generate the observations \mathbf{x}_t via the densities $p[\mathbf{x} | y]$. The y_t that was actually involved in generating \mathbf{x}_t is also called the model's (true) *representation* of \mathbf{x}_t . The means of the Gaussians $p[\mathbf{x} | y]$ are at the corners of the unit square, as shown in Figure 1B, and the variances of these Gaussians are $\sigma^2 I$. The parameters in the model are the prior distribution of z , its temporal dynamics $\mathcal{T}_{z_{t-1} z_t}$, the conditional distributions $\mathcal{O}_{z y}$, and the emission densities $p[\mathbf{x} | y]$. It is assumed that all the parameters have already been correctly learned at the outset of the inference problem.

Figure 2 shows an example of a sequence of 400 states generated from the model. The state in the z layer stays the same for an average of about 30 time steps, and then switches to one of the other states, chosen with equal probability. The key inference problem is to determine the posterior distribution over y_t , that best explains the observation \mathbf{x}_t , given the past experiences $\mathcal{D}_{t-1} = \{\mathbf{x}_1, \dots, \mathbf{x}_{t-1}\}$.

Inference of the true posterior distribution, $P[y_t | \mathbf{x}_t, \mathcal{D}_{t-1}] = p[y_t | \mathcal{D}_t]$, uses temporal contextual information, consisting of existing knowledge built up from past observations, as well as the new observation \mathbf{x}_t . Figure 3A shows the structure of the standard HMM inference model, where the posterior distributions $P[y_t | \mathcal{D}_t]$ and $P[z_t | \mathcal{D}_t]$ can be computed using a procedure

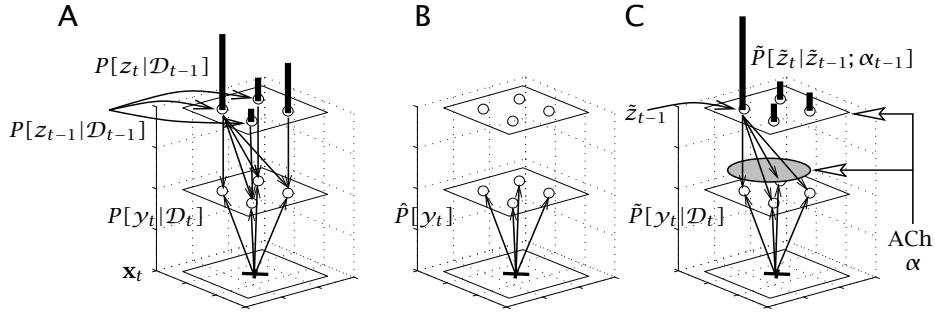


Figure 3: Recognition models. A) Exact recognition model. $P[z_{t-1}|\mathcal{D}_{t-1}]$ is propagated to provide the prior $P[z_t|\mathcal{D}_{t-1}]$ (shown by the lengths of the thick vertical bars), and thus the prior $P[y_t|\mathcal{D}_{t-1}]$. This is combined with the likelihood term from the data \mathbf{x}_t to give the true $P[y_t|\mathcal{D}_t]$. B) Bottom-up recognition model uses only a generic prior over y_t , which conveys no information, so the likelihood term dominates. C) ACh model. A single estimated state \tilde{z}_{t-1} is used, in conjunction with its certainty α_{t-1} , reported by cholinergic activity, to produce an approximate prior $\tilde{P}[\tilde{z}_t|\tilde{z}_{t-1};\alpha_{t-1}]$ over z_t (which is a mixture of a delta function and a uniform), and thus an approximate prior over y_t . This is combined with the likelihood to give an approximate $\tilde{P}[y_t|\mathcal{D}_t]$, and a new cholinergic signal α_t is calculated.

that is equivalent to the forward part of the forwards-backwards algorithm (Rabiner, 1989). The adaptation to include the y layer is straightforward.

In each time step t , the top-down information is communicated by z_t , while the bottom-up information is carried by x_t . The prior distribution over z_t

$$P[z_t|\mathcal{D}_{t-1}] = \sum_{z_{t-1}} P[z_{t-1}|\mathcal{D}_{t-1}] \mathcal{T}_{z_{t-1}z_t} \quad (2)$$

distills the contextual information from past experiences \mathcal{D}_{t-1} . This information is propagated to the representational units y by

$$P[z_t, y_t|\mathcal{D}_{t-1}] = P[z_t|\mathcal{D}_{t-1}] \mathcal{O}_{z_t y_t} . \quad (3)$$

The bottom-up information, $P[y_t|\mathbf{x}_t]$, is proportional to the likelihood, $p[\mathbf{x}_t|y_t]$, and interacts with the top-down information, $P[z_t, y_t|\mathcal{D}_{t-1}]$ in the conditioning step:

$$P[z_t, y_t|\mathcal{D}_t] \propto P[z_t, y_t|\mathcal{D}_{t-1}] p[\mathbf{x}_t|y_t] \quad (4)$$

where the constant of proportionality normalizes the full conditional distribution. From the joint posterior distribution over z and y , we can then compute the marginalized posterior distribution of y_t , which gives the relative belief in each of the states of y_t having generated the current observation x_t , in the context of past experiences:

$$P[y_t|\mathcal{D}_t] = \sum_{z_t} P[z_t, y_t|\mathcal{D}_t] , \quad (5)$$

This distribution, henceforth referred to as the exact posterior, is the fullest possible repre-

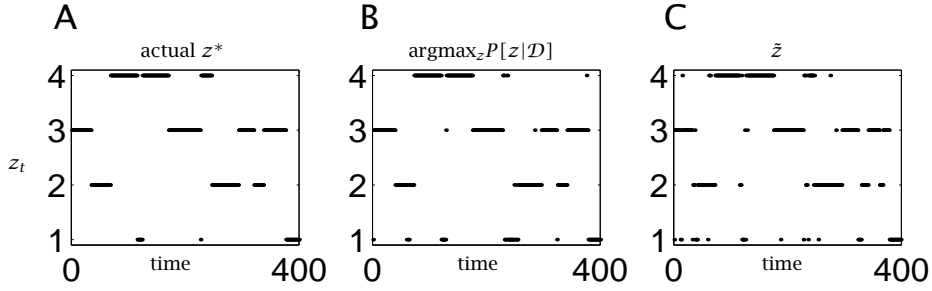


Figure 4: Contextual representation in exact inference. A) Actual z . B) Highest probability z state from the exact posterior distribution. C) Single \tilde{z} state from the ACh-mediated approximate inference model.

sensation of \mathbf{x}_t . One can also create the updated contextual information:

$$P[z_t | \mathcal{D}_t] = \sum_{y_t} P[z_t, y_t | \mathcal{D}_t]. \quad (6)$$

Figures 4 and 5 show various aspects of inference in the HMM for a particular run. The “true” contextual states $\{z_1^*, z_2^*, \dots\}$, the true representational states $\{y_1^*, y_2^*, \dots\}$, and the observations $\{\mathbf{x}_1, \mathbf{x}_2, \dots\}$ are generated from the model with the parameters given in figure 1B. The posterior distributions over z_t and y_t given \mathcal{D}_t , that is all the observations up to, and including time t , are computed at each time step using the algorithm detailed above. If the algorithm is working properly, then we would expect to see a high correspondence between the “true” contextual state z_t^* and the inferred, most likely state $z_t^\circ = \operatorname{argmax}_{z_t} P[z_t | \mathcal{D}_t]$. Figure 4A;B shows that z_t° mostly replicates z_t^* faithfully. One quirk of inference in HMMs is that these individually most likely states z_t° do not form a most likely state *sequence* as, for instance, found by the Viterbi algorithm.

Figure 5 shows normalized histograms of the representational posterior probabilities of the true states y_t^* (figure 5A) and the other states \bar{y}_t^* (figure 5B). As one might hope, the former are generally large, and the latter generally small.

The exact inference algorithm that we have described achieves good performance. However, one may well ask whether it is computational feasible for the brain to perform the complete, exact inference in all its mathematical complexity. Viewed abstractly, the most critical problem seems that of maintaining and manipulating simultaneously the information about all possible contexts ($P[z_t | \mathcal{D}_t]$). This is particularly difficult in the face of population coding, for which the activity pattern of one or a few populations of units in relevant cortical areas are used to represent all possible contexts. Of course, in our simple example, there are only four possible contexts. However, in general, there are potentially as many contexts as known visual environments, a huge number.

A “naive” solution to the complexity problem is to use only the likelihood term, $p[\mathbf{x}_t | y_t]$, in the inference about the current representational states y_t , and ignore the top-down contextual information altogether. This is actually the traditional model of inference for unsupervised




Figure 5: Quality of exact representational inference. Normalized histograms of A) the exact posterior distribution $P[y_t | \mathcal{D}]$ over the actual state y_t^* (upper) and B) the other possible states $y_t \neq y_t^*$ (lower, written $P[\bar{y}_t^*]$). The x-axis is divided into bins of $P[y_t | \mathcal{D}]$ ranging from 0 to 1, and the y-axis refers to the relative frequency of a particular state of y_t being in each of the binned probability intervals in the posterior distribution. This is an indication of the quality of exact representational inference.

analysis-by-synthesis models (eg Hinton & Ghahramani, 1997). Figure 3B shows the structure of a purely bottom-up model, where the approximate posterior is computed by $\hat{P}[y_t | \mathbf{x}_t] = p[\mathbf{x}_t | y_t] / \mathcal{Z}$, where \mathcal{Z} is a normalization factor. Purely bottom-up inference solves the problem of high computational costs: there is no need to carry any information from one time step to the next. However, the performance of this algorithm is likely to be poor, whenever the probability distribution of generating \mathbf{x} for the different values of y overlap substantially, as is the case in our example. This is just the ambiguity problem described above.

Figure 6A shows the representational performance of this model, through a scatter-plot of $\hat{P}[y_t | \mathbf{x}_t]$ against the exact posterior $P[y_t | \mathcal{D}_t]$. If bottom-up inference was perfectly correct, then all the points would lie on the diagonal line of equality. The bow-shape shows that purely bottom-up inference is relatively poor. The particularly concentrated upper and lower boundaries indicate that when the true posterior distribution assigns a very high or very low probability to a state of y , the corresponding distribution inferred from bottom-up information alone tends to assign a much more neutral probability to that state. This tendency highlights the loss of the contribution of the disambiguating top-down signal in the bottom-up model. With only the bottom-up information, it rarely happens that one can say with confidence that a state of y_t is either definitely the one, or definitely not the one, that generated \mathbf{x}_t . The exact shape of the envelope is determined by the extent of overlap in the densities $p[\mathbf{x} | y]$ for the various values of y , though we have yet to analyze this relationship in detail.




Figure 6: Representational performance. Comparison of (A) the purely bottom-up $\hat{P}[\mathcal{Y}|\mathcal{D}]$ and (B) the ACh-based approximation $\tilde{P}[\mathcal{Y}|\mathcal{D}]$ with the true $P[\mathcal{Y}|\mathcal{D}]$ across all values of \mathcal{Y} . The ACh-based approximation is substantially more accurate.

3 ACh-Mediated Approximate Inference

A natural compromise between the exact inference model, which is representationally and computationally expensive, and the naive inference model, which has poor performance, is to use a model that captures useful top-down information at a realistic computational cost. The intuition we gain from exact inference is that top-down expectations can resolve bottom-up ambiguities, permitting better processing. However, in the face of contextual uncertainty, top-down information is just generic. Thus, we consider a model in which just a single contextual state is represented in the activity of contextual units (presumably in pre-frontal areas), and ACh is used to report on the uncertainty of this contextual state and to control the balance between bottom-up and top-down inference. In exact inference, the notion of uncertainty is captured in the (entropy of) the posterior distribution of the contextual state $P[z_t|\mathcal{D}_{t-1}]$ in equation 2. This uncertainty determines the relative strength of the top-down information, $P[z_t, \mathcal{Y}_t|\mathcal{D}_{t-1}]$, compared with the information from the likelihood $p[\mathbf{x}_t|\mathcal{Y}_t]$, in equation 4.

More formally, in our ACh-mediated approximate inference model, only two quantities of information about the context are maintained over time: \tilde{z}_{t-1} , the most likely contextual *state* having seen \mathcal{D}_{t-1} , and α_{t-1} , the measure of *uncertainty* associated with that state. The idea is that α_{t-1} is reported by the level of ACh, and is used to control the extent to which top-down information based on \tilde{z}_{t-1} is used to influence inference about \mathcal{Y}_t .

Figure 3C shows a schematic diagram of the proposed approximate inference model. If we were given the full, exact posterior distribution $P[z_{t-1}, \mathcal{Y}_{t-1}|\mathcal{D}_{t-1}]$, then one natural definition for this ACh signal would be the uncertainty in the most likely contextual state

$$\alpha_{t-1} = 1 - \max_z P[z_{t-1} = z|\mathcal{D}_{t-1}] = 1 - P[z_{t-1}^\circ|\mathcal{D}_{t-1}] \quad (7)$$

Figure 7B shows the resulting ACh signal for one run with the actual sequence of contextual states shown in Figure 7A. As expected, ACh level is generally high at times when the true state

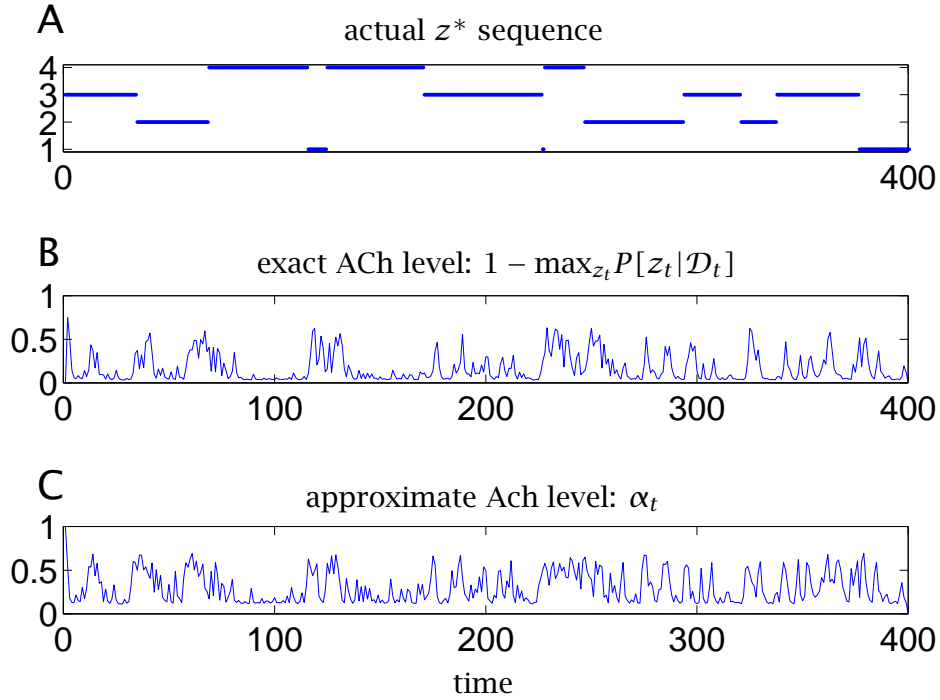


Figure 7: ACh model. A) Actual sequence of contextual states z for one run. B) ACh level from the exact posterior in the same run. C) ACh level α_t from the approximate model. Note the coarse similarity between A and B.

z_t^* is changing, and decreases during the periods that z_t^* is constant. During times of change, top-down information is confusing or potentially incorrect, and so the current context should be abandoned while a new context is gradually built up from a period of perception that is mainly dominated by bottom-up input. This switch in inferential strategy is just the putative inferential effect of ACh.

The ACh signal of Figure 7B was calculated assuming knowledge of the true posterior. This is, of course, unreasonable. The model of figure 3C includes the key approximation that the only information from \mathcal{D}_{t-1} about the state of z , besides uncertainty signaled by ACh, is in the single choice of context variable \tilde{z}_{t-1} . As in the full inference model, the first computation at each time step t is to compute the prior distribution over z_t . However, since the full posterior distribution of z_{t-1} is no longer available, it has to be reconstructed from what is believed to be the most likely contextual state \tilde{z}_{t-1}° , together with its associated uncertainty α_{t-1} :

$$\tilde{P}[\tilde{z}_{t-1}; \alpha_{t-1}] = \begin{cases} 1 - \alpha_{t-1}, & \text{if } \tilde{z}_{t-1} = \tilde{z}_{t-1}^\circ \\ \frac{\alpha_{t-1}}{n_z - 1}, & \text{otherwise} \end{cases} \quad (8)$$

The approximation made is that all the non-explicitly modeled states of z_{t-1} equally share a fraction of the probability that \tilde{z}_{t-1} was not the “true” context. As before, the contextual

information is then propagated to z_t and y_t :

$$\tilde{P}[\tilde{z}_t; \alpha_{t-1}] = \sum_{z_{t-1}} \tilde{P}[\tilde{z}_{t-1}; \alpha_{t-1}] \mathcal{T}_{z_{t-1}z_t} \quad (9)$$

$$\tilde{P}[y_t, \tilde{z}_t; \alpha_{t-1}] = \tilde{P}[\tilde{z}_t | \tilde{z}_{t-1}; \alpha_{t-1}] \mathcal{O}_{\tilde{z}_t y_t} \quad (10)$$

and the new observation is incorporated into the inference in the conditioning step:

$$\tilde{P}[y_t, \tilde{z}_t | \mathcal{D}_t] \propto \tilde{P}[y_t, \tilde{z}_t | \tilde{z}_{t-1}; \alpha_{t-1}] p[\mathbf{x}_t | y_t] \quad (11)$$

The new posterior distributions are computed, as before, by marginalizing the joint posterior distribution:

$$\tilde{P}[y_t | \mathcal{D}_t] = \sum_{z_t} \tilde{P}[y_t, \tilde{z}_t | \mathcal{D}_t] \quad (12)$$

$$\tilde{P}[\tilde{z}_t | \mathcal{D}_t] = \sum_{y_t} \tilde{P}[y_t, \tilde{z}_t | \mathcal{D}_t] \quad (13)$$

In addition, we compute the contextual information that is to be propagated to the next time step:

$$\tilde{z}_t^\circ = \operatorname{argmax}_{z_t} \tilde{P}[\tilde{z}_t | \mathcal{D}_t] \quad \text{most likely contextual state} \quad (14)$$

$$\alpha_t = 1 - \tilde{P}[\tilde{z}_t^\circ | \mathcal{D}_t] \quad \text{ACh level} \quad (15)$$

The crucial differences between this approximate inferential algorithm and the exact one detailed before are the use of ACh as a scalar measure of uncertainty (equation 15) and the reconstruction of an estimate of the posterior distribution from this scalar estimate (equation 8). If α_{t-1} (ie the ACh level) is high, then the input stimulus-bound likelihood term dominates in the conditioning process (equation 11); if α_{t-1} (ie the ACh level) is low, then the temporal context (\tilde{z}_{t-1}) and likelihood terms are appropriately balanced. These computations are all local and straightforward, except for the representation and normalization of the joint distribution over y_t and \tilde{z}_t .

One potentially dangerous aspect of this inference procedure is that it might get unreasonably committed to a single state: $\tilde{z}_{t-1} = \tilde{z}_t = \dots$. Because the probabilities accorded to the other possible values of z_{t-1} given \mathcal{D}_{t-1} are not explicitly represented from one time step to the next, there is little chance for uncertainties about a context to build up, a condition necessary for inducing a context switch. A natural way to avoid this is to bound the ACh level from below by a constant, φ , making approximate inference slightly more stimulus-bound than exact inference. Thus, in practice, rather than using equation 15, we use

$$\alpha_t = \varphi + (1 - \varphi)(1 - \tilde{P}[\tilde{z}_t^\circ | \mathcal{D}_t]) \quad (16)$$

Larger values of φ lead to larger *guaranteed* contribution of the bottom-up, stimulus-bound likelihood term to inference.

Figure 7C shows the approximate ACh level for the same case as in Figure 7A;B, using $\varphi = 0.1$. Although the detailed value of this signal over time is clearly different from that arising from an exact knowledge of the posterior probabilities in Figure 7B, the gross movements are quite




Figure 8: Representational cost. Solid: the mean extra representational cost for the true state y_t^* over that in the exact posterior using the ACh model as a function of the minimum allowed ACh level φ . Dashed: the same quantity for the pure bottom-up model (which is equivalent to the approximate model for $\varphi = 1$), denoted $\Delta\langle\log \hat{P}[y^*]\rangle$ here. Errorbars show standard errors of the means over 1000 trials.

similar. Note the effect of φ in preventing the ACh level from dropping to 0. Figure 6B shows that the ACh-based approximate posterior values $\tilde{P}[y|\mathcal{D}]$ are much closer to the true values than for the purely bottom-up model, particularly for values of $P[y_t|\mathcal{D}_t]$ near 0 and 1, where most data lie. Figure 4C shows that inference about z_t^* is noisy, but the pattern of true values is certainly visible.

Figure 8 shows the effects of different φ on the quality of inference about the true states y_t^* . What is plotted is the difference between approximate and exact log probabilities of the true states y_t^* , averaged over 1000 cases. The average log likelihood for the exact model is -210 . If $\varphi = 1$, then inference is completely stimulus-bound, just like the purely bottom-up model. Note the poor performance for this case. For values of φ slightly less than 0.2, the approximate inference model does well, both for the particular setting of parameters described in Figure 1B and for a range of other values (not shown here). An upper bound on the performance of approximate inference can be calculated in three steps by: i) using the exact posterior to work out \tilde{z}_t and α_t , ii) using these values to approximate $P[\tilde{z}_t; \alpha_t]$ as in equation 8, and iii) using this approximate distribution in equation 10 and the remaining equations. The average resulting cost (*ie* the average resulting difference from the log probability under exact inference) is -3.5 log units. Thus, the ACh-based approximation performs well, and much better than purely bottom-up inference.

4 Experimental Data

We have argued that accurate reporting of top-down uncertainty, signaled by ACh level, is important for achieving an *optimal* balance between top-down and bottom-up processing in perceptual inference. To validate this statement, two main properties of cholinergic modulation need to be verified. One is the modulatory effect of ACh on the relative processing of top-down and bottom-up information. The other is that it is uncertainty that controls ACh release in relevant areas of the cortex. We know of no experiment that directly tests our theory - however, data from various behavioral and physiological studies lend some support to its elements. For example, abnormal ACh levels, due to either pharmacological manipulations or neurological diseases, lead to characteristic deficits in attentional perceptual tasks (Sarter & Bruno, 1998) and general behavioral symptoms such as hallucination (Perry & Perry, 1995). Also, physiological data indicate that the cellular and network effects of ACh activation on the processing of sensory stimuli are *facilitatory*, as we would expect, although there is not yet data showing what specific effects ACh has on top-down information. There is also a scarcity of data on the *drive* underlying ACh activation. After summarizing the relevant, existing data, we will propose some experiments to investigate aspects of ACh-mediated perception that are less well understood.

4.1 General Behavioral Effects of ACh

Tasks involving sustained attention appear to involve ACh. Sustained attention refers to a prolonged state of readiness to respond to rarely and unpredictably occurring signals (Sarter, Givens, & Bruno, 2001). Uncertainty associated with contextual information in these tasks can be induced by introducing variability in stimulus presentation time, location, stimulus luminance or duration, or alternating between signal and non-signal trials (Parasuraman, 1986; Parasuraman, Warm, & Dember, 1987). Manipulations of ACh release pattern or downstream effects, according to our theory, would lead to performance impairment. In particular, stimuli typically occur rarely in these tasks, so there is a strong top-down expectation of detecting nothing. Thus, an abnormally low level of cortical ACh might lead to excessive confidence in not detecting a stimulus when one is present, but only a small effect when no stimulus is present. Correspondingly, when ACh level is abnormally high, bottom-up information is excessively processed, perhaps lending undue credibility to signals that arise from irrelevant noise. Thus, we might expect to see an increase in false alarms with no effects on hits. Data from experiments in which cortical ACh levels are pharmacologically manipulated (Holley, Turchi, Apple, & Sarter, 1995; Turchi & Sarter, 2001; McGaughy, Kaiser, & Sarter, 1996) corroborate these hypotheses.

Weaker evidence comes from tasks involving selective spatial attention (see Kramer, Coles, & Logan, 1996). In a version that has been used to study neuromodulatory effects, one of two stimulus locations is cued and then a delay period introduced, before the target stimulus appears and the subject is required to respond (Muir, Dunnett, Robbins, & Everitt, 1992; Jackson, Marrocco, & Posner, 1994). The spatial cue provides the obvious source of top-down infor-

mation, and uncertainty associated with that information can be induced by presenting the target stimulus at the unexpected location. Abnormal ACh levels, induced by pharmacological manipulations (Muir et al., 1992), basal forebrain lesions (Muir, Everitt, & Robbins, 1994), or neurological diseases (Parasuraman, Greenwood, Haxby, & Grady, 1992), result in longer reaction times and, in more severe cases, lower accuracy. In addition, cortical administration of physostigmine, an ACh re-uptake inhibitor, can eliminate behavioral deficits in the basal forebrain-lesioned animals at the appropriate dosage, indicating the deficit is cholinergic in nature (Muir et al., 1992). Moreover, dosages greater or less than optimal both lead to worsening in performance, indicating that it is not merely the presence of ACh in the cortex, but the *appropriate* level of ACh that enables optimal perceptual inference.

Data from a third type of attentional task, divided attention with modality uncertainty, also implicate ACh in modulating perceptual processes (Turchi & Sarter, 1997). In one version of a divided attention task, the subject is trained to respond with different sets of response rules depending on the modality of the stimulus (*eg* visual *versus* auditory). Compared to the unimodal trials, the bimodal condition imposes an additional source of uncertainty. When ACh level is kept abnormally low due to a basal forebrain lesion (Turchi & Sarter, 1997), rats are observed to have longer reaction times than controls in the bimodal case, but not the unimodal case. On the basis of our theory, we might indeed expect that excessive processing of the top-down information (coming from a uniform distribution across modalities) would harm the integration of helpful bottom-up, modality information, potentially resulting in longer reaction times.

A further source of behavioral data on cholinergic modulation comes from patients with neurological diseases. Cortical cholinergic deficit is common among patients diagnosed with Lewy Body Dementia, Parkinson's Disease, and Alzheimer's Disease. A symptom common among these patients is hallucinations, or the imagined or distorted perception of sensory stimuli. In our model, such hallucinations might reflect an incorrect over-reliance on top-down information because of inadequate ACh. In such patients, the severity of the hallucinations appears to be correlated with cholinergic depletion. In Lewy Body Dementia patients, for instance, cholinergic enzyme activity in temporal and parietal cortex is reduced to below 20% of the normal in hallucinating patients, compared with around 50% in those not experiencing hallucinations (Perry et al., 1993). It is interesting to note that the majority of plants with identified hallucinogenic chemicals contain anti-muscarinic agents such as scopolamine and atropine (Schultes & Hofmann, 1992). Hallucinatory experiences induced by these chemicals are enhanced during eye closure and suppressed by visual input (Fisher, 1991). Many patients with Lewy Body Dementia and Alzheimer's Disease also exhibit paresthesias, or the discernment of images such as faces or animals in wallpaper, curtains, or clouds (Perry & Perry, 1995), a condition ameliorated by the administration of physostigmine, an ACh reuptake-inhibitor (Cummings, Gorman, & Shapira, 1993). Of course, patient data must be interpreted cautiously, in particular in this case, since there are many different forms of hallucinations, various of which are induced by non-cholinergic pharmacological factors.

4.2 Physiological effects of ACh

Physiological studies, though traditionally focusing only on the effects of ACh on bottom-up, stimulus-bound processing, have suggested a much more specific set of cholinergic effects than behavioral studies. There exists a large body of physiological data from both anesthetized and awake animals supporting the notion that basal forebrain ACh activation enhances stimulus processing across sensory cortices (Sillito & Murphy, 1987; Metherate, Ashe, & Weinberger, 1990; Tremblay, Warren, & Dykes, 1990). For example, tetanic stimulation in the nucleus basalis increases cortical responsiveness by facilitating the ability of synaptic potentials in thalamocortical connections to elicit action potentials in the rat auditory cortex (Metherate, Ashe, & Weinberger, 1993; Hars, Maho, Edeline, & Hennevin, 1993). This effect is blocked by the application of atropine, a muscarinic receptor antagonist. In the rat somatosensory cortex, iontophoretic application of ACh enhances sensory stimulus-evoked discharges (Donoghue & Carroll, 1987). In the cat visual cortex, iontophoretic injection of ACh induces a striking increase in stimulus-specific responses without concomitant loss in selectivity (Sillito & Kemp, 1983). In the cat somatosensory cortex, simultaneous iontophoretic application of ACh in single cells and tactile stimulation induce short-term potentiation in the majority of cells (Metherate, Tremblay, & Dykes, 1987).

At the network level, ACh seems selectively to promote the flow of information in the feedforward pathway over that in the top-down feedback pathway. Via nicotinic receptors, ACh appears selectively to enhance thalamocortical synapses without affecting the other synapses (Gil, Connors, & Amitai, 1997). In addition, it has been observed that ACh strongly suppresses intracortical connectivity in the visual cortex through presynaptic muscarinic receptors, but has a much reduced effect on thalamocortical afferents that arise from white matter (Kimura et al., 1999). Similarly, experiments in brain slice preparations of the rat somatosensory cortex indicate that ACh selectively suppresses synaptic potentials elicited by the stimulation of layer I, which contains a high percentage of feedback synapses, while having no effect on synaptic potentials elicited by the stimulation of layer IV, which has a high percentage of what we would consider as feedforward synapses (Hasselmo & Cekic, 1996). The overall effect of ACh activation appears to enable the stimulus-bound input to have a dominant effect in the sensory cortices.

4.3 Drive of ACh activation

Unfortunately, there is little detailed information on what drives ACh activation. It is known that endogenous, task-related release of ACh occurs shortly before the presentation of stimuli, as measured by microdialysis (Fadel, Sarter, & Bruno, 2001) and single-cell recordings (Richardson & DeLong, 1991), indicating at a minimum that ACh release cannot be a simple consequence of ongoing perceptual processing.

4.4 New Experiments

From the perspective of our model, the two areas in which data on cortical cholinergic modulation are particularly lacking are the effects of ACh on top-down processing relative to bottom-up processing and the control of ACh activation. Investigating these issues requires a perceptual task in which the contextual and sensory inputs are clearly distinct, and both they and their inferential implications can be carefully controlled.

One possible paradigm is that described in Ress *et al* (Ress, Backus, & Heeger, 2000), which uses fMRI techniques to measure visual cortical activity during a stimulus detection task. Their data show a strong, stimulus-independent, spatially selective (for the region in which the stimulus is likely to appear), and putatively top-down signal correlated with detection performance. In terms of our model, this top-down signal would correspond to something like $\tilde{P}[\mathcal{Y}_t, \tilde{z}_t; \alpha_{t-1}]$ or $\hat{P}[\mathcal{Y}_t; \alpha_{t-1}]$. To ascertain the effects of top-down uncertainty on this signal, it is possible to manipulate uncertainty by changing the average ratio of signal to non-signal trials in a given session. Or, more radically, something closer to the selective spatial attention tasks as described above could be introduced in the task. When uncertainty is high, the signal should be weak; when uncertainty is low, the signal should be strong. If this signal indeed varies as a function of uncertainty as expected, then a next step would be to investigate the relationship between ACh level, this fMRI signal, and the subject's behavioral performance. The ACh level can be manipulated with the administration of ACh receptor agonists/antagonists, ACh re-uptake inhibitors, or other drugs with known effects on cortical ACh. We predict that elevated ACh levels would lead to a strong top-down signal, and a selectively impaired performance with increased false alarms but not misses. In contrast, lowered ACh level would lead to a weaker top-down signal and an impairment in performance reflected mainly in the lengthening of reaction time.

A direct relationship between top-down uncertainty and ACh level is difficult to establish using the Ress paradigm, as the available techniques for measuring ACh level - single-cell recordings and microdialysis - are inappropriate for fMRI studies. However, this relationship can be separately investigated by adapting the same task to animal models. The top-down signal, presumably similar to that for humans, would no longer be monitored. However, it would be possible to measure the level of ACh using either single-cell recordings in the nucleus basalis, the main source of cortical ACh, or microdialysis in the visual cortex itself. Top-down uncertainty could again be manipulated by varying the average frequency of signal trials relative to non-signal trials. Note that, due to the lack of data, it is difficult to hypothesize *a priori* what the precise relationship is between uncertainty and the concentration of ACh. In our mathematical model, we implicitly assumed a linear relationship and ignored the fact that ACh release and effects have various components at different time scales (Sarter & Bruno, 1997; Hasselmo, 1995). It will be very interesting to find out the exact relationship between uncertainty and ACh level, and how this relationship differs for the tonic and phasic components of ACh signal.

5 Discussion

We have suggested that one role for ACh in cortical processing is to report contextual uncertainty in order to control the balance between stimulus-bound, bottom-up processing, and context-bound, top-down processing. This is an extension of computational and mechanistic ideas about the involvement of cholinergic modulation in learning in attractor networks. We used the example of a hierarchical HMM, in which representational inference for a middle layer correctly reflects such a balance, and showed that a simple model of the drive and effects of ACh leads to competent inference.

The mathematical model we have used to illustrate our general theory on ACh is overly simple in several respects. For example, it uses a localist representation for the state z , so that exact inference is feasible. Also, only a two-level hierarchy was modeled, whereas sensory systems in the brain are known to involve many levels of processing. It would be more biologically realistic to consider distributed representations at each of many levels in a hierarchy, and in which only one or a very few contexts, presumably stored in the prefrontal cortex, could be entertained at once. Also, it is necessary to modify the steps in Equations 10 and 11, since it would be difficult to represent the joint uncertainty over representations at multiple levels in the hierarchy. Further, despite the limited physiological evidence mentioned above, exactly which sets of cortical connections should be modulated by ACh is not completely clear (see also Dayan, 1999). For instance, to what extent should recurrent connections within a cortical column, or within a cortical hypercolumn be influenced in the same way as long range horizontal connections between hypercolumns, or top-down connections from cortical areas higher in the cortical hierarchy?

A strong assumption made in using a hidden Markov model is that the sequence of contextual states obeys the Markov property: the context at any particular time step only depends on the context in the preceding step and not on any of the previous ones. However, perceptual inference in real systems has the potential of using top-down information from arbitrarily distant past, stored in long-term memory. A more sophisticated mathematical model would be needed to capture the contribution of multiple and longer term temporal dependencies.

In the HMM example, we also did not consider sources of top-down information that are distinct from temporal context. Clearly, processes such as inter-modality interactions and spatial contextual integration could also exert a top-down influence on perceptual inference. By comparison with multiple timescales, it would be relatively straightforward to build a more complete mathematical model that captures these other sources of information too.

In this work, we have discussed ACh in the context of perceptual inference in isolation, independent of other processes modulated by ACh. However, ample data at both the systems (*eg* Hasselmo & Bower, 1993) and cellular (*eg* Sillito & Kemp, 1983) levels indicate ACh plays an important role in cortical learning (Hasselmo & Bower, 1993). Studies of cholinergic modulation in conditioning (Holland, 1997; Holland & Gallagher, 1999; Dayan et al., 2000) suggest there are strong interactions between learning and inference, both of which are modulated by ACh. As in the work that inspired ours, the association of ACh with uncertainty or unfamiliarity

makes it an ideal signal for controlling learning.

We took from the animal conditioning studies that ACh might report uncertainty. However, we have not modeled a particularly interesting aspect of those studies, namely that cholinergic modulation might differentially affect learning for different stimuli that are present simultaneously. If the same applies for different possible contexts, this could significantly enrich the model.

Another important aspect of cholinergic modulation we have omitted is non-basal forebrain cholinergic modulation, *ie* mainly subcortical innervation by the pedunculopontine nucleus, the cuneiform nucleus, and the laterodorsal tegmental nucleus. ACh released by these nuclei has been implicated in modulating REM sleep (Jewett & Norton, 1986; Velazquez-Moctezuma, Shiromani, & Gillin, 1990; Lavie, Pratt, Scharf, Peled, & Brown, 1984) and saccadic eye movement (Aizawa, Kobayashi, Yamamoto, & Isa, 1999), among other processes. It is not yet clear what, if any, similarities or interactions exist in the drive and effects of cortical ACh released by the basal forebrain and by the other sources.

A final important aspect of cholinergic modulation that we have not yet addressed is the interaction between ACh and other neuromodulators. For example, there is evidence that dopaminergic afferents from the nucleus accumbens modulates the activity of cholinergic neurons in the basal forebrain, and furthermore, that this dopaminergic modulation underlies the cholinergic impairments in schizophrenics (Sarter & Bruno, 1998). It has also been suggested that ACh and norepinephrine play complementary roles in cortical developmental plasticity (Bear & Singer, 1986; Kirkwood, Rozas, Kirkwood, & Perez, 1999).

References

- Aizawa, H., Kobayashi, Y., Yamamoto, M., & Isa, T. (1999). Injection of nicotine into the superior colliculus facilitates occurrence of express saccades in monkeys. *Journal of Neurophysiology*, *82*(3), 1642-6.
- Amit, D. J. (1989). *Modelling brain function: the world of attractor neural networks*. Cambridge University Press.
- Bear, M. F., & Singer, W. (1986). Modulation of visual cortical plasticity by acetylcholine and noradrenaline. *Nature*, *320*(6058), 172-6.
- Becker, S. (1999). Implicit learning in 3d object recognition: the importance of temporal context. *Neural Computation*, *11*(2), 347-74.
- Cummings, J. L., Gorman, D. G., & Shapira, J. (1993). Physostigmine ameliorates the delusions of Alzheimer's disease. *Biological Psychiatry*, *33*, 536-541.
- Dayan, P. (1999). Recurrent sampling models for the helmholtz machine. *Neural Computation*, *11*, 653-677.
- Dayan, P., Kakade, S., & Montague, P. R. (2000). Learning and selective attention. *Nature Neuroscience*, *3*, 1218-1223.
- Donoghue, J. P., & Carroll, K. L. (1987). Cholinergic modulation of sensory responses in rat primary somatic sensory cortex. *Brain Research*, *408*, 367-71.
- Everitt, B. J., & Robbins, T. W. (1997). Central cholinergic systems and cognition. *Annual Review in Psychology*, *48*, 649-84.
- Fadel, J., Sarter, M., & Bruno, J. P. (2001). Basal forebrain glutamatergic modulation of cortical acetylcholine release. *Synapse*, *39*(3), 201-12.
- Fisher, C. M. (1991). Visual hallucinations on eye closure associated with atropine toxicity. *Canadian Journal of Neurological Sciences*, *18*, 18-27.
- Gil, Z., Connors, B. W., & Amitai, Y. (1997). Differential regulation of neocortical synapses by neuromodulators and activity. *Neuron*, *19*(3), 679-86.
- Grenander, U. (1995). *Elements of pattern theory*. Baltimore, MD: Johns Hopkins University Press.
- Hars, B., Maho, C., Edeline, J. M., & Hennevin, E. (1993). Basal forebrain stimulation facilitates tone-evoked responses in the auditory cortex of awake rat. *Neuroscience*, *56*, 61-74.
- Hasselmo, M. E. (1995). Neuromodulation and cortical function: modeling the physiological basis of behavior. *Behavioural Brain Research*, *67*(1), 1-26.
- Hasselmo, M. E., & Bower, J. M. (1993). Acetylcholine and memory. *Trends in Neuroscience*, *16*(6), 218-22.

- Hasselmo, M. E., & Cekic, M. (1996). Suppression of synaptic transmission may allow combination of associative feedback and self-organizing feedforward connections in the neocortex. *Behavioural Brain Research*, *79*, 153-61.
- Helmholtz, H. von. (1896). Visual perception: Essential readings. *Physiological optics*, *3*(26), 1-36.
- Hinton, G. E., & Ghahramani, Z. (1997). Generative models for discovering sparse distributed representations. *Philosophical transactions of the royal society of london B*, *352*, 1177-90.
- Holland, P. (1997). Brain mechanisms for changes in processing of conditioned stimuli in pavlovian conditioning: implications for behavior theory. *Animal Learning and Behavior*, *25*(4), 373-399.
- Holland, P. C., & Gallagher, M. (1999). Amygdala circuitry in attentional and representational processes. *Trends in Cognitive Sciences*, *3*(2), 65-74.
- Holley, L. A., Turchi, J., Apple, C., & Sarter, M. (1995). Dissociation between the attentional effects of infusions of a benzodiazepine receptor agonist and an inverse agonist into the basal forebrain. *Psychopharmacology*, *120*, 99-108.
- Jackson, S. R., Marrocco, R., & Posner, M. I. (1994). Networks of anatomical areas controlling visuospatial attention. *Neural Networks*, *7*(6-7), 925-944.
- Jewett, R. E., & Nortan, S. (1986). Effect of some stimulant and depressant drugs on sleep cycles of cats. *Experimental Neurology*, *15*, 463-74.
- Kimura, F., Fukuda, M., & Tsumoto, T. (1999). Acetylcholine suppresses the spread of excitation in the visual cortex revealed by optical recording: possible differential effect depending on the source of input. *European Journal of Neuroscience*, *11*(10), 3597-3609.
- Kirkwood, J., Rozas, C., Kirkwood, J., & Perez, F. (1999). Modulation of long-term synaptic depression in visual cortex by acetylcholine and norepinephrine. *Journal of Neuroscience*, *19*(5), 1599-1609.
- Kramer, A. F., Coles, M. G. H., & Logan, G. D. (1996). *Converging operations in the study of visual selective attention*. Washington, DC: APA.
- Lavie, P., Pratt, H., Scharf, B., Peled, R., & Brown, J. (1984). Localized pontine lesion; nearly total absence of REM sleep. *Neurology*, *34*, 118-20.
- Linster, C., & Hasselmo, H. E. (2001). Neuromodulation and the functional dynamics of piriform cortex. *Chemical Senses*, *26*(5), 585-94.
- Marder, E. (1998). From biophysics to models of network function. *Annual Review in Neuroscience*, *21*, 25-45.
- McGaughy, J., Kaiser, T., & Sarter, M. (1996). Behavioral vigilance following infusions of 192 IgG-saporin into the basal forebrain: selectivity of the behavioral impairment and relation to cortical AChE-positive fiber density. *Behavioral Neuroscience*, *110*(2), 247-265.

- Metherate, R., Asche, J. H., & Weinberger, N. M. (1990). Acetylcholine modifies neuronal acoustic rate-level functions in guinea pig auditory cortex by an action at muscarinic receptors. *Synapse*, 6, 364-8.
- Metherate, R., Asche, J. H., & Weinberger, N. M. (1993). Nucleus basalis stimulation facilitates thalamocortical synaptic transmission in the rat auditory cortex. *Synapse*, 132-143.
- Metherate, R., Tremblay, N., & Dykes, R. W. (1987). Acetylcholine permits long-term enhancement of neuronal responsiveness in cat primary somatosensory cortex. *Neuroscience*, 22(1), 75-81.
- Muir, J. L., Dunnett, S. B., Robbins, T. W., & Everitt, B. J. (1992). Attentional functions of the forebrain cholinergic system: effects of intraventricular hemicholinium, physostigmine, basal forebrain lesions and intracortical grafts on a multiple-choice serial reaction time task. *Experimental Brain Research*, 89(3), 611-22.
- Muir, J. L., Everitt, B. J., & Robbins, T. W. (1994). AMPA-induced excitotoxic lesions of the basal forebrain: a significant role for the cortical cholinergic system in attentional function. *Journal of Neuroscience*, 14, 2313-26.
- Neisser, U. (1967). *Cognitive psychology*. New York, NY: Appleton-Century-Croft.
- Parasuraman, R. (1986). Vigilance, monitoring, and search. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance, cognitive processes and performance* (Vol. 2, p. 43.1-43.39). Wiley.
- Parasuraman, R., Greenwood, P. M., Haxby, J. V., & Grady, C. L. (1992). Visuospatial attention in dementia of the alzheimer type. *Brain*, 115, 711-33.
- Parasuraman, R., Warm, J. S., & Dember, W. N. (1987). Vigilance: taxonomy and utility. In L. S. Mark, J. Warm, & R. L. Huston (Eds.), *Ergonomics and human factors* (p. 11-39). Springer Verlag.
- Pearce, J. M., & Hall, G. (1980). A model for pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87(6), 532-52.
- Perry, E. K., Marshall, E., Thompson, P., McKeith, I. G., Collerton, D., Fairbairn, A. F., Ferrier, I. N., Irving, D., & Perry, R. H. (1993). Monoaminergic activities in Lewy body dementia: relation to hallucinosis and extrapyramidal features. *Journal of Neural Transmission*, 6(3), 167-77.
- Perry, E. K., & Perry, R. H. (1995). Acetylcholine and hallucinations: disease-related compared to drug-induced alterations in human consciousness. *Brain and Cognition*, 28, 240-258.
- Pflüger, H. J. (1999). Neuromodulation during motor development and behavior. *Current opinion in Neurobiology*, 9, 683-89.
- Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77, 257-86.

- Ress, D., Backus, B. T., & Heeger, D. (2000). Activity in primary visual cortex predicts performance in a visual detection task. *Nature Neuroscience*, 3(9), 940-945.
- Richardson, R. T., & DeLong, M. R. (1991). Electrophysiological studies of the functions of the nucleus basalis in primates. *Advances in Experimental Medical Biology*, 295, 233-52.
- Sarter, M., & Bruno, J. P. (1997). Cognitive functions of cortical acetylcholine: toward a unifying hypothesis. *Brain Research Reviews*, 23, 28-46.
- Sarter, M., & Bruno, J. P. (1998). Cortical acetylcholine, reality distortion, schizophrenia, and Lewy body dementia: too much or too little cortical acetylcholine? *Brain and Cognition*, 38, 297-316.
- Sarter, M., Givens, B., & Bruno, J. P. (2001). The cognitive neuroscience of sustained attention: where top-down meets bottom-up. *Brain Research Reviews*, 35, 146-160.
- Schultes, R. E., & Hofmann, A. (1992). *Plants of the gods*. Healing Arts Press.
- Sillito, A. M., & Kemp, J. A. (1983). Cholinergic modulation of the functional organization of the cat visual cortex. *Brain Research*, 289, 143-55.
- Sillito, A. M., & Murphy, P. C. (1987). The cholinergic modulation of cortical function. In E. G. Jones & A. Peters (Eds.), *Cerebral cortex*. Plenum Press.
- Tremblay, N., Warren, R. A., & Dykes, R. W. (1990). Electrophysiological studies of acetylcholine and the role of the basal forebrain in the somatosensory cortex of the cat. II. cortical neurons excited by somatic stimuli. *Journal of Neurophysiology*, 64(4), 1212-22.
- Turchi, J., & Sarter, M. (1997). Cortical acetylcholine and processing capacity: effects of cortical cholinergic deafferentation on crossmodal divided attention in rats. *Brain Research Cognitive Brain Research*, 6(2), 147-58.
- Turchi, J., & Sarter, M. (2001). Bidirectional modulation of basal forebrain N-methyl-D-aspartate receptor function differentially affects visual attention but not visual discrimination performance. *Neuroscience*, 104(2), 407-17.
- Velazquez-Moctezuma, J., Shiromani, P., & Gillin, J. C. (1990). Acetylcholine and acetylcholine receptor subtypes in REM sleep generation. *Progress in Brain Research*, 84, 407-413.