

Perceptual learning, motor learning, and automaticity

Statistically optimal perception and learning: from behavior to neural representations

József Fiser^{1,2}, Pietro Berkes¹, Gergő Orbán^{1,3} and Máté Lengyel⁴

¹ National Volen Center for Complex Systems, Brandeis University, Volen 208/MS 013, Waltham, MA 02454, USA

² Department of Psychology and the Neuroscience Program, Brandeis University, 415 South Street, Waltham, MA 02453, USA

³ Department of Biophysics, Research Institute for Particle and Nuclear Physics, Hungarian Academy of Sciences, Konkoly Thege Miklós út 29–33, H-1121, Budapest, Hungary

⁴ Computational and Biological Learning Lab, Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, United Kingdom

Human perception has recently been characterized as statistical inference based on noisy and ambiguous sensory inputs. Moreover, suitable neural representations of uncertainty have been identified that could underlie such probabilistic computations. In this review, we argue that learning an internal model of the sensory environment is another key aspect of the same statistical inference procedure and thus perception and learning need to be treated jointly. We review evidence for statistically optimal learning in humans and animals, and re-evaluate possible neural representations of uncertainty based on their potential to support statistically optimal learning. We propose that spontaneous activity can have a functional role in such representations leading to a new, sampling-based, framework of how the cortex represents information and uncertainty.

Probabilistic perception, learning and representation of uncertainty: in need of a unifying approach

One of the longstanding computational principles in neuroscience is that the nervous system of animals and humans is adapted to the statistical properties of the environment [1]. This principle is reflected across all organizational levels, ranging from the activity of single neurons to networks and behavior, and it has been identified as key to the survival of organisms [2]. Such adaptation takes place on at least two distinct behaviorally relevant time scales: on the time scale of immediate inferences, as a moment-by-moment processing of sensory input (perception), and on a longer time scale by learning from experience. Although statistically optimal perception and learning have most often been considered in isolation, here we promote them as two facets of the same underlying principle and treat them together under a unified approach.

Although there is considerable behavioral evidence that humans and animals represent, infer and learn about the statistical properties of their environment efficiently [3], and there is also converging theoretical and neurophysiological work on potential neural mechanisms

of statistically optimal perception [4], there is a notable lack of convergence from physiological and theoretical studies explaining whether and how statistically optimal learning might occur in the brain. Moreover, there is a missing link between perception and learning: there exists virtually no crosstalk between these two lines of research focusing on common principles and on a unified framework down to the level of neural implementation. With recent advances in understanding the bases of probabilistic coding and the accumulating evidence supporting probabilistic computations in the cortex, it is now possible to take a closer look at both the basis of probabilistic learning and its relation to probabilistic perception.

We first provide a brief overview of the theoretical framework as well as behavioral and neural evidence for representing uncertainty in perceptual processes. To highlight the parallels between probabilistic perception and learning, we then revisit in more detail the same issues with regard to learning. We argue that a main challenge is to pinpoint representational schemes that enable neural circuits to represent uncertainty for both perception and learning, and compare and critically evaluate existing proposals for such representational schemes. Finally, we review a seemingly disparate set of findings regarding variability of evoked neural responses and spontaneous activity in the cortex and suggest that these phenomena can be interpreted as part of a representational framework that supports statistically optimal inference and learning.

Probabilistic perception: representing uncertainty, behavioral and neural evidence

At the level of immediate processing, perception has long been characterized as unconscious inference, where incoming sensory stimuli are interpreted in terms of the objects and features that gave rise to them [5]. Traditional approaches treated perception as a series of classical signal processing operations, by which each sensory stimulus should give rise to a single perceptual interpretation [6]. However, because sensory input in general is noisy and ambiguous, there is usually a range of different possible

Corresponding author: Fiser, J. (fiser@brandeis.edu).

Glossary

Expected utility: the average expected reward associated with a particular decision, α , when the state of the environment, y , is unknown. It can be computed by calculating the average of the utility function, $U(\alpha, y)$, describing the amount of reward obtained when making decision α if the true state of the environment is y , with regard to the *posterior distribution*, $p(y|x)$, describing the degree of belief about the state of the environment given some sensory input, x : $R(\alpha) = \int U(\alpha, y) p(y|x) dy$.

Likelihood: the function specifying the probability $p(x|y, M)$ of observing a particular stimulus x for each possible state of the environment, y , under a probabilistic model of the environment, M .

Marginalization: the process by which the distribution of a subset of variables, y_1 , is computed from the joint distribution of a larger set of variables, $\{y_1, y_2\}$: $p(y_1) = \int p(y_1, y_2) dy_2$. (This could be important if, for example, different decisions rely on different subsets of the same set of variables.) Importantly, in a sampling-based representation, in which different neurons represent these different subsets of variables, simply "reading" (e.g. by a downstream brain area) the activities of only those neurons that represent y_1 automatically implements such a marginalization operation.

Maximum a posteriori (or MAP) estimate: in the context of *probabilistic inference*, it is an approximation by which instead of representing the full *posterior* distribution, only a single value of y is considered that has the highest probability under the *posterior*. (Formally, the full *posterior* is approximated by a Dirac-delta distribution, an infinitely narrow Gaussian, located at its maximum.) As a consequence, *uncertainty* about y is no longer represented.

Maximum likelihood estimate: as the *MAP estimate*, it is also an approximation, but the full *posterior* is approximated by the single value of y which has the highest *likelihood*.

Posterior: the probability distribution $p(y|x, M)$ produced by *probabilistic inference* according to a particular probabilistic model of the environment, M , giving the probability that the environment is in any of its possible states, y , when stimulus x is being observed.

Prior: the probability distribution $p(y|M)$ defining the expectation about the environment being in any of its possible states, y , before any observation is available according to a probabilistic model of the environment, M .

Probabilistic inference: the process by which the posterior is computed. It requires a probabilistic model, M , of stimuli x and states of the environment y , containing a *prior* and a *likelihood*. It is necessary when environmental states are not directly available to the observer: they can only be inferred from stimuli through inverting the relationship between y and x through Bayes' rule: $p(y|x, M) = p(x|y, M) p(y|M) / Z$, where Z is a factor independent of y , ensuring that the *posterior* is a well-defined probability distribution. Note, that the *posterior* is a full probability distribution, rather than a single estimate over environmental states, y . In contrast with approximate inference methods, such as *maximum likelihood* or *maximum a posteriori* that compute single best estimates of y , the *posterior* fully represents the uncertainty about the inferred variables.

Probabilistic learning: the process of finding a suitable model for *probabilistic inference*. This itself can be viewed as a problem of probabilistic inference at a higher level, where the unobserved quantity is the model, M , including its parameters and structure. Thus, the complete description of the results of probabilistic learning is a posterior distribution, $p(M|X)$, over possible models given all stimuli observed so far, X . Even though approximate versions, such as *maximum likelihood* or *MAP*, compute only a single best estimates of M , they still need to rely on representing uncertainty about the states of the environment, y . The effect of learning is usually a gradual change in the *posterior* (or estimate) as more and more stimuli are observed, reflecting the incremental nature of learning.

interpretations compatible with any given input. A well-known example is the ambiguity caused by the formation of a two-dimensional image on our retina by objects that are three-dimensional in reality (Figure 1a). When such multiple interpretations arise, the mathematically appropriate way to describe them is to assign a probability value to each of them that expresses how much one believes that a particular interpretation might reflect the true state of the world [7], such as the true three-dimensional shape of an object in Figure 1a. Although the principles of probability theory have been established and applied to studying economic decision-making for centuries [8], only recently has their relevance to perception been appreciated, causing a quiet paradigm shift from signal processing

to probabilistic inference as the appropriate theoretical framework for studying perception.

Evidence has been steadily growing in recent years that the nervous system represents its uncertainty about the true state of the world in a probabilistically appropriate way and uses such representations in two cognitively relevant domains: information fusion and perceptual decision-making. When information about the same object needs to be fused from several sources, inferences about the object should rely on these sources commensurate with their associated uncertainty. That is, more uncertain sources should be relied upon less (Figure 1b). Such probabilistically optimal fusion has been demonstrated in multisensory integration [9,10] when the different sources are different sensory modalities, and also between information coming from the senses and being stored in memory [11,12].

Probabilistic representations are also key to decision-making under risk and uncertainty [13]. Making a well-informed decision requires knowledge about the true state of the world. When there is uncertainty about the true state, then the decision expected to yield the most reward (or utility) can be computed by weighing the reward associated with each possible state of the world with their probabilities (Figure 1c). Indeed, several studies demonstrated that, in simple sensory and motor tasks, humans and animals do take into account their uncertainty in such a way [14].

The compelling evidence at the level of behavior that humans and animals represent uncertainty during perceptual processes initiated intense research into the neural underpinnings of such probabilistic representations. The two main questions that have been addressed is how sensory stimuli are represented in a probabilistic manner by neural cell populations (i.e. how neural activities encode probability distributions over possible states of the sensory world), and how the dynamics of neural circuits implement appropriate probabilistic inference with these representations [15,16]. As a result, there has been a recent surge of converging experimental and theoretical work on the neural bases of statistically optimal inference. This work has shown that the activity of groups of neurons in particular decision tasks can be related to probabilistic representations and that dynamical changes in neural activities are consistent with probabilistic computations with the represented variables [4,17].

In summary, the study of perception as probabilistic inference, for which the key is the representation of uncertainty, provides an exemplary synthesis of a sound theoretical background with behavioral as well as neural evidence.

Probabilistic learning: representing uncertainty

In contrast to immediate processing during perception, the probabilistic approach to learning has been less explored in a neurobiological context. This is surprising given the fact that, from a computational standpoint, probabilistic inference based on sensory input is always made according to a model of the sensory environment which typically needs to be acquired by learning (Figure 2). Thus, the goal of probabilistic learning can be defined as acquiring appropriate models for inference based on past experience.

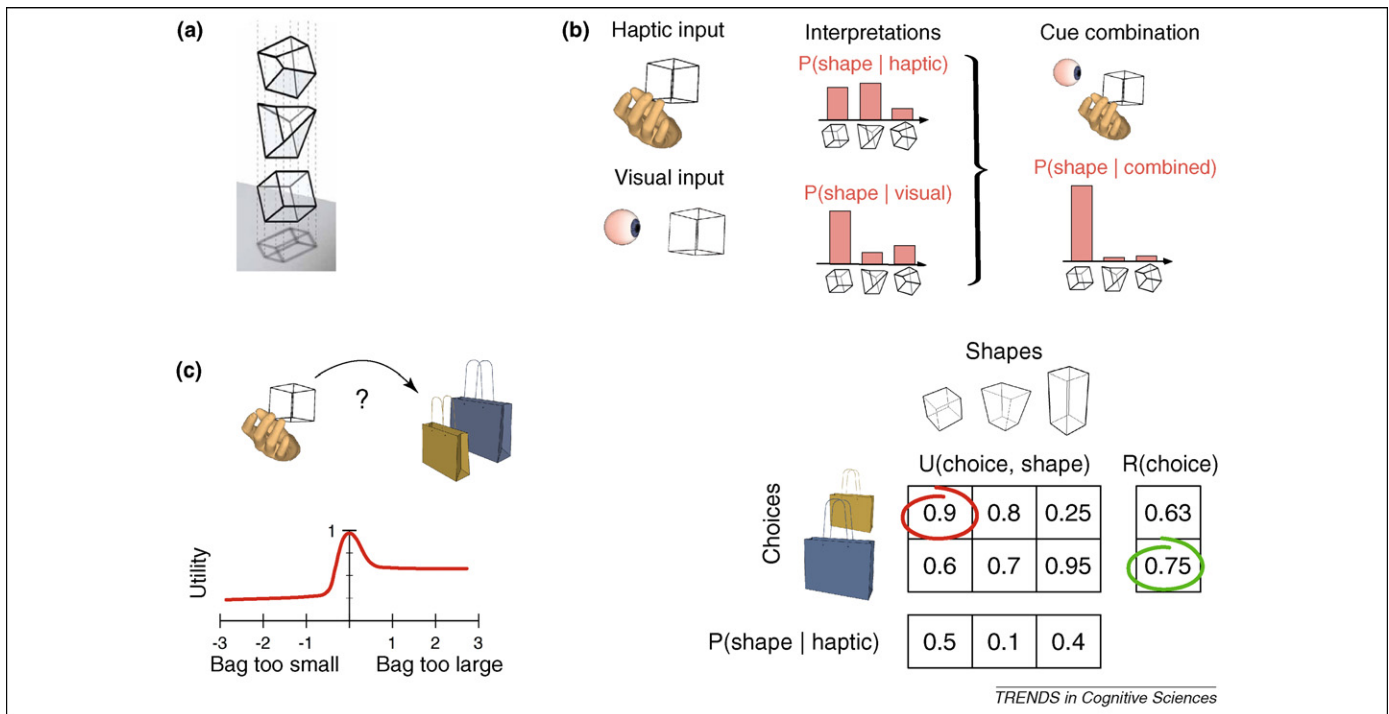


Figure 1. Representation of uncertainty and its benefits. (a) Sensory information is inherently ambiguous. Given a two-dimensional projection on a surface (e.g. a retina), it is impossible to determine which of the three different three-dimensional wire frame objects above cast the image (adapted with permission from [96]). (b) Cue integration. Independent visual and haptic measurements (left) support to different degrees the three possible interpretations of object identity (middle). Integrating these sources of information according to their respective uncertainties provides an optimal probabilistic estimate of the correct object (right). (c) Decision-making. When the task is to choose the bag with the right size for storing an object, uncertain haptic information needs to be utilized probabilistically for optimal choice (top left). In the example shown, the utility function expresses the degree to which a combination of object and bag size is preferable: for example, if the bag is too small, the object will not fit in, if it is too large, we are wasting valuable bag space (bottom left, top right). In this case, rather than inferring the most probable object based on haptic cues and then choosing the bag optimal for that object (in the example, the small bag for the cube), the probability of each possible object needs to be weighted by its utility and the combination with the highest expected utility (R) has to be selected (in the example, the large bag has the highest expected utility). Evidence shows that human performance in cue combination and decision-making tasks is close to optimal [10,97].

Importantly, just as perception can be formalized as inferring hidden states, variables, of the environment from the limited information provided by sensory input (e.g. inferring the true three-dimensional shape and size of the seat of a chair from its two-dimensional projection on our retinae), learning can be formalized as inferring some more persistent hidden characteristics, parameters, of the environment based on limited experience. These inferences could target concrete physical parameters of objects, such as the typical height or width of a chair, or more abstract descriptors, such as the possible categories to which objects can belong (e.g. chairs and tables) (Figure 2).

There are two different ways in which representing uncertainty is important for learning. First, learning about our environment modifies the perceptual inferences we draw from a sensory stimulus. That is, the same stimulus gives rise to different uncertainties after learning. For example, having learned about the geometrical properties of chairs and tables allows us to increase our confidence that an unusually looking stool is really more of a chair than a table (Figure 2). At the neural level, this constrains learning mechanisms to change neural activity patterns such that they correctly encode the ensuing changes in perceptual uncertainty, thus keeping the neural representation of uncertainty self-consistent before and after learning. Second, representing uncertainty does not just constrain but also benefits learning. For example, if there is uncertainty as to whether an object is a chair or a table, our models for both of these categories should be updated,

rather than only updating the model of the most probable category (Figure 2). Crucially, the optimal magnitude of these updates depends directly (and inversely) on the uncertainty about the object belonging to each category: the model of the more probable category should be updated to a larger degree [18].

Thus, probabilistic perception implies that learning must also be probabilistic in nature. Therefore, we now examine behavioral and neural evidence for probabilistic learning.

Probabilistic learning: behavioral level

Evidence for humans and animals being sensitive to the probabilistic structure of the environment ranges from low-level perceptual mechanisms, such as visual grouping mechanisms conforming with the co-occurrence statistics of line edges in natural scenes [19], to high-level cognitive decisions such as humans' remarkably precise predictions about the expected life time of processes as diverse as cake baking or marriages [20]. A recent survey demonstrated how research in widely different areas ranging from classical forms of animal learning to human learning of sensorimotor tasks found evidence of probabilistic learning [21]. It has been found that configural learning in animals [22], causal learning in rats [23] as well as in human infants [24] and a vast array of inductive learning phenomena fit comfortably with a hierarchical probabilistic framework, in which probabilistic learning is performed at increasingly higher levels of abstraction [25].

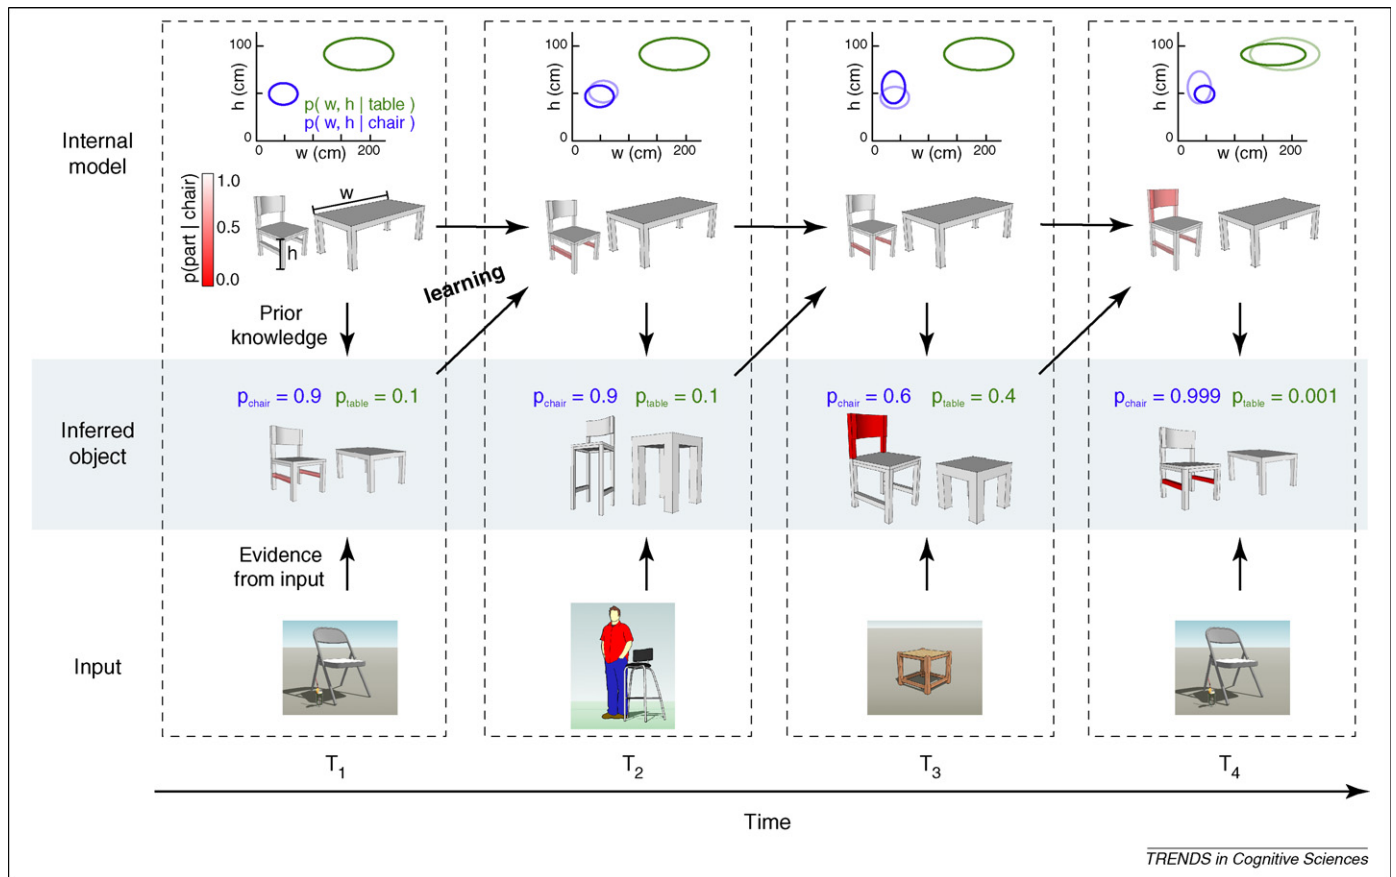


Figure 2. The link between probabilistic inference and learning. (Top row) Developing internal models of chairs and tables. The plot shows the distribution of parameters (two-dimensional Gaussians, represented by ellipses) and object shapes for the two categories. (Middle row) Inferences about the currently viewed object based on the input and the internal model. (Bottom row) Actual sensory input. Red color code represents the probability of a particular object part being present (see color scale on top left). T1–T4, four successive illustrative iterations of the inference–learning cycle. (T1) The interpretation of a natural scene requires combining information from the sensory input (bottom) and the internal model (top). Based on the internal models of chairs and tables, the input is interpreted with high probability ($p = 0.9$) as a chair with a typical size but missing crossbars (middle). (T2) The internal model of the world is updated based on the cumulative experience of previous inferences (top). The chair in T1, being a typical example of a chair, requires minimal adjustments to the internal model. Experience with more unusual instances, as for example the high chair in T2, provokes more substantial changes (T3, top). (T3) The representation of uncertainty allows to update the internal model taking into account all possible interpretations of the input. In T3, the stimulus is ambiguous as it could be interpreted as a stool, or a square table. The internal model needs to be updated by taking into account the relative probability of the two interpretations: that there exist tables with a more square shape or that some chairs miss the upper part. Since both probabilities are relatively high, both internal models will be modified substantially during learning (see the change of both ellipses). (T4) After learning, the same input as in T1 elicits different responses owing to changes in the internal model. In T4, the input is interpreted as a chair with significantly higher confidence, as experience has shown that chairs often lack the bottom crossbars.

A particularly direct line of evidence for humans learning complex, high-dimensional distributions of many variables by performing higher-order probabilistic learning, not just naïve frequency-based learning, comes from the domain of visual statistical learning (Box 1). An analysis of a series of visual statistical learning experiments showed that beyond the simplest results, recursive pairwise associative learning is inadequate for replicating human performance, whereas Bayesian probabilistic learning not only accurately replicates these results but it makes correct predictions about human performance in new experiments [26].

These examples suggest a common core representational and learning strategy for animals and humans that shows remarkable statistical efficiency. However, such behavioral studies provide no insights as to how these strategies might be implemented in the neural circuitry of the cortex.

Probabilistic learning in the cortex: neural level

Although psychophysical evidence has been steadily growing, there is little direct electrophysiological evidence showing that learning and development in neural systems is optimal in a statistical sense even though the effect of

learning on cortical representations has been investigated extensively [27,28]. One of the main reasons for this is that there have been very few plausible computational models proposed for a neural implementation of probabilistic learning that would provide easily testable predictions (but see Refs [29,30]). Here, we give a brief overview of the computational approaches developed to capture probabilistic learning in neural systems and discuss why they are unsuitable in their current form for being tested in electrophysiological experiments.

Classical work on connectionist models aimed at devising neural networks with simplified neural-like units that could learn about the regularities hidden in a stimulus ensemble [31]. This line of research has been developed further substantially and demonstrated explicitly how dynamical interactions between neurons in these networks correspond to computing probabilistic inferences, and how the tuning of synaptic weights corresponds to learning the parameters of a probabilistic model of input stimuli [32–34]. A key common feature in these statistical neural networks is that inference and learning are inseparable: inference relies on the synaptic weights encoding a useful

Box 1. Visual statistical learning in humans

In experimental psychology, the term “statistical learning” refers to a particular type of implicit learning that emerged from investigating artificial grammar learning. It is fundamentally different from traditional perceptual learning and was first used in the domain of infant language acquisition [72]. The paradigm has been adapted from auditory to other sensory modalities such as touch and vision, to different species and various aspects of statistical learning have been explored such as multimodal interactions [73], effects of attention, interaction with abstract rule learning [74], together with its neural substrates [75]. The emerging consensus based on these studies is that statistical learning is a domain-general, fundamental learning ability of humans and animals that is probably a major component of the process by which internal representations of the environment are developed.

The basic idea of the statistical learning paradigm is to create an artificial mini-world by using a set of building blocks to generate several composite inputs that represent possible instances in this world. In the case of visual statistical learning (VSL), artificial visual scenes are composed from abstract shape elements where the building blocks are two or more such shapes always appearing in the same relative configuration (Figure 1). An implicit learning paradigm is used to test how internal visual representations emerge through passively observing a large number of such composite

scenes without any instruction as to what to pay attention to. After being exposed to the scenes, when subjects have to choose between two fragments of shape combinations based on familiarity, they reliably more often choose fragments that were true building blocks of the scenes compared to random combinations of shapes [76]. Similar results were found in 8-month-old infants [77,78] suggesting that humans from an early age can automatically extract the underlying structure of an unknown sensory data set based purely on statistical characteristics of the input.

Investigations of VSL provided evidence of increasingly sophisticated aspects of this learning, setting it apart from simple frequency-based naïve learning methods. Subjects not only automatically become sensitive to pairs of shapes that appear more frequently together, but also to pairs with elements that are more predictive of each other even when the co-occurrence of those elements is not particularly high [76]. Moreover, this learning highly depends on whether or not a pair of elements is a part of a larger building structure, such as a quadruple [79]. Thus, it appears that human statistical learning is a sophisticated mechanism that is not only superior to pairwise associative learning but also potentially capable to link appearance-based simple learning and higher-level “rule-learning” [26].

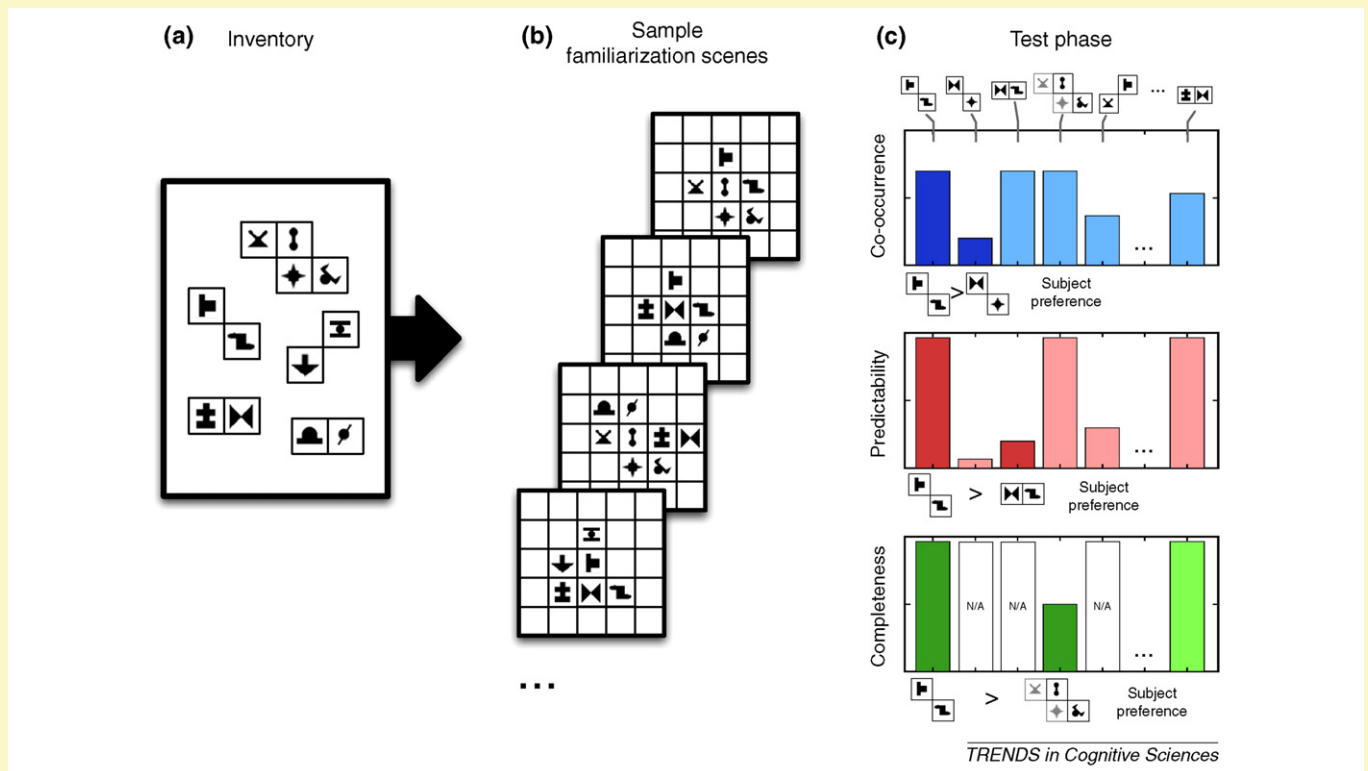


Figure 1. Visual statistical learning. (a) An inventory of visual chunks is defined as a set of two or more spatially adjacent shapes always co-occurring in scenes. (b) Sample artificial scenes composed of multiple chunks that are used in the familiarization phase. Note that there are no obvious low-level segmentation cues giving away the identity of the underlying chunks. (c) During the test phase, subjects are shown pairs of segments that are either parts of chunks or random combinations (segments on the top). The three histograms show different statistical conditions. (Top) There is a difference in co-occurrence frequency of elements between the two choices; (middle) co-occurrence is equated, but there is difference in predictability (the probability of one symbol given that the other is present) between the choices; (bottom) both co-occurrence and predictability are equated between the two choices, but the completeness statistics (what percentage of a chunk in the inventory is covered by the choice fragment) is different – one pair is a standalone chunk, the other is a part of a larger chunk. Subjects were able to use cues in any of these conditions, as indicated by the subject preferences below each panel. These observations can be accounted for by optimal probabilistic learning, but not by simpler alternatives such as pairwise associative learning (see text).

probabilistic model of the environment, whereas learning proceeds by using the inferences produced by the network (Figure 3a).

Although statistical neural networks have considerably grown in sophistication and algorithmic efficiency in recent years [35], providing cutting-edge performance in some challenging real-world machine learning tasks, much less

attention has been devoted to specifying their biological substrates. At the level of general insights, these models suggested ways in which internally generated spontaneous activity patterns (“fantasies”) that are representative of the probabilistic model encoded by the neural network can have important roles in the fine tuning of synapses during off-line periods of functioning. They also clarified that

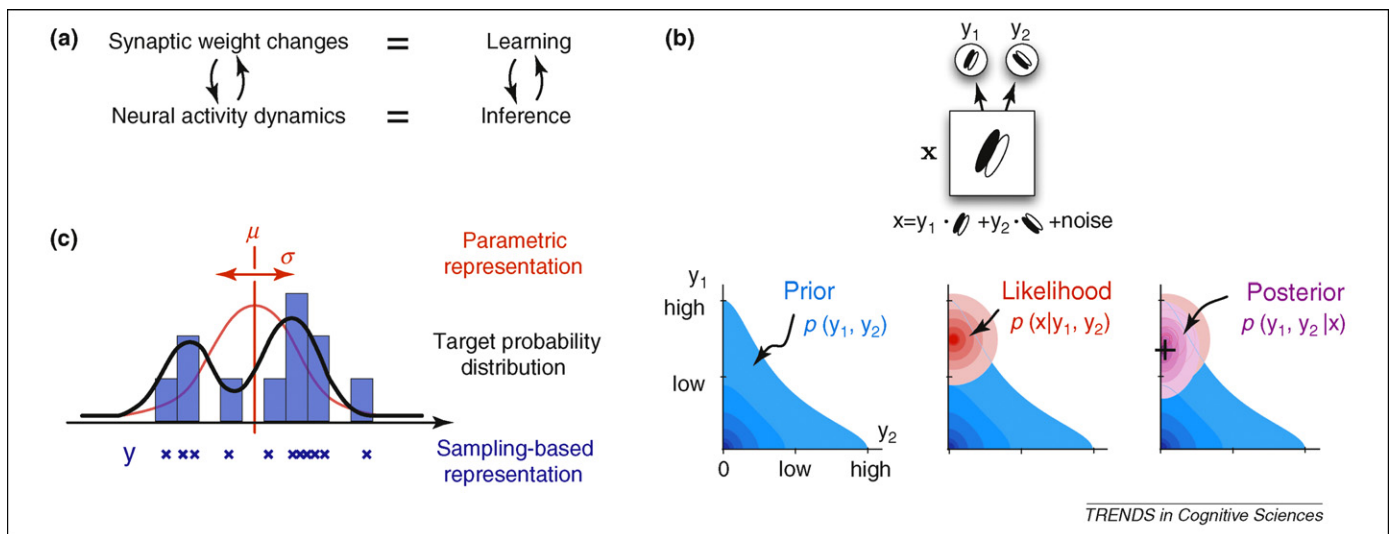


Figure 3. Neural substrates of probabilistic inference and learning. (a) Functional mapping of learning and inference onto neural substrates in the cortex. (b) Probabilistic inference for natural images. (Top) A toy model of the early visual system (based on Ref. [43]). The internal model of the environment assumes that visual stimuli, x , are generated by the noisy linear superposition of two oriented features with activation levels, y_1 and y_2 . The task of the visual system is to infer the activation levels, y_1 and y_2 , of these features from seeing only their superposition, x . (Bottom left) The *prior distribution* over the activation of these features, y_1 and y_2 , captures prior knowledge about how much they are typically (co-)activated in images experienced before. In this example, y_1 and y_2 are expected to be independent and sparse, which means that each feature appears rarely in visual scenes and independently of the other feature. (Bottom middle) The *likelihood* function represents the way the visual features are assumed to combine to form the visual input under our model of the environment. It is higher for feature combinations that are more likely to underlie the image we are seeing according to the equation on the top. (Bottom right) The goal of the visual system is to infer the *posterior distribution* over y_1 and y_2 . By Bayes' theorem, the posterior optimally combines the expectations from the prior with the evidence from the likelihood. *Maximum a posteriori* (MAP) estimate, used by some models [40,43,47], denoted by a + in the figure neglects uncertainty by using only the maximum value instead of the full distribution. (c) Simple demonstration of two probabilistic representational schemes. (Black curve) The probability distribution of variable y to be represented. (Red curve) Assumed distribution by the parametric representation. Only the two parameters of the distribution, the mean μ and variance σ are represented. (Blue "x"-s and bars) Samples and the histogram implied by the sampling-based representation.

useful learning rules for such tuning always include Hebbian as well as anti-Hebbian terms [32,33]. In addition, several insightful ideas about the roles of bottom-up, recurrent, and top-down connections for efficient inference and learning have also been put forward [36–38], but they were not specified at a level that would allow direct experimental tests.

Learning internal models of natural images has traditionally been one area where the biological relevance of statistical neural networks was investigated. As these studies aimed at explaining the properties of early sensory areas, the “objects” they learned to infer were simple localized and oriented filters assumed to interact mostly additively in creating images (Figure 3b). Although these representations are at a lower level than the “true” objects constituting our environment (such as chairs and tables) that typically interact in highly non-linear ways as they form images (owing to e.g. occlusion [39]), the same principles of probabilistic inference and learning also apply to this level. Indeed, several studies showed how probabilistic learning of natural scene statistics leads to representations that are similar to those found in simple and complex cells of the visual cortex [40–44]. Although some early studies were not formulated originally in a statistical framework [40,41,43], later theoretical developments showed that their learning algorithms were in fact special cases of probabilistic learning [45,46].

The general method of validation in these learning studies almost exclusively concentrated on comparing the “receptive field” properties of model units with those of sensory cortical neurons and showing a good match between the two. However, as the emphasis in many of these models is on learning, the details of the mapping of

neural dynamics to inference were left implicit (with some notable exceptions [44,47]). In cases where inference has been defined explicitly, neurons were usually assumed to represent single deterministic (so-called “*maximum a posteriori*”) estimates (Figure 3b). This failure to represent uncertainty is not only computationally harmful for inference, decision-making and learning (Figures 1–2) but it is also at odds with behavioral data showing that humans and animals are influenced by perceptual uncertainty. Moreover, this approach constrains predictions to be made only about receptive fields which often says little about trial-by-trial, on-line neural responses [48].

In summary, presently a main challenge in probabilistic neural computation is to pinpoint representational schemes that enable neural networks to represent uncertainty in a physiologically testable manner. Specifically, learning with such representations on naturalistic input should provide verifiable predictions about the cortical implementation of these schemes beyond receptive fields.

Probabilistic representations in the cortex for inference and learning

The conclusion of this review so far is that identifying the neural representation of uncertainty is key for understanding how the brain implements probabilistic inference and learning. Crucially, because inference and learning are inseparable, a viable candidate representational scheme should be suitable for both. In line with this, evidence is growing that perception and memory-based familiarity processes once thought to be linked to anatomically clearly segregated cortical modules along the ventral pathway of the visual cortex could rely on integrated multipurpose representations within all areas [49]. In this section, we

review the two main classes of probabilistic representational schemes that are the best candidates for providing neural predictions and investigate their suitability for inference and learning.

Theoretical proposals of how probabilities can be represented in the brain fall into two main categories: schemes in which neural activities represent parameters of the probability distribution describing uncertainty in sensory variables, and schemes under which neurons represent the sensory variables themselves, such as models based on sampling. A simple example highlighting the core differences between the two approaches can be given by describing our probabilistic beliefs about the height of a chair (as in [Figure 2](#)). A parameter-based description starts with assuming that these beliefs can be described by one particular type of probability distribution, for example a Gaussian, and then specifies values for the relevant parameters of this distribution, for example the mean and (co)variance (describing our average prediction about the height of the chair and the “error bars” around it) ([Figure 3c](#)). A sampling-based description does not require that our beliefs can be described by one particular type of probability distribution. Rather, it specifies a series of possible values (samples) for the variable(s) of interest themselves, here the height of the particular chair viewed, such that if one constructed a histogram of these samples, this histogram would eventually trace out the probability distribution actually describing our beliefs ([Figure 3c](#)).

Probabilistic population codes (PPCs) are well-known examples of parameter-based descriptions and they are widely used for probabilistic modeling of inference making [\[16\]](#). Recently, neurophysiological support for PPCs in cortical areas related to decision-making was also reported [\[50\]](#). The key concept in PPCs, just as in their predecessors, kernel density estimator codes [\[51\]](#) and distributional population codes [\[15\]](#), is that neural activities encode parameters of the probability distribution that is the result of probabilistic inference ([Box 2](#)). As a consequence, a full probability distribution is represented at any moment in time and therefore changes in neural activities encode dynamically changing distributions as inferences are updated based on continuously incoming stimuli. Several recent studies explored how such representations can be used for optimal perceptual inference and decision-making in various tasks [\[52\]](#). However, a main shortcoming of PPCs is that, at present, there is no clear option to implement learning within this framework.

The alternative, sampling-based approach to representing probability distributions is based on the idea that each neuron represents an individual variable from a high-dimensional multivariate distribution of external variables, and, therefore, each possible pattern of network activities corresponds to a point in this multivariate “feature” space ([Box 2](#)). Uncertainty is naturally encoded by network dynamics that stochastically explore a series of neural activity patterns such that the corresponding features are sampled according to the particular distribution that needs to be represented. Importantly, there exist worked out examples of how learning can be implemented in this framework: almost all the classical statistical

neural networks have already been using this representational scheme implicitly [\[32,33,37\]](#). The deterministic representations used in some of the most successful statistical receptive field models [\[42,43\]](#) can also be conceived as approximate versions of the sampling-based representational approach.

Despite its appeal for learning, there have been relatively few attempts to explicitly articulate the predictions of a sampling-based representational scheme for neural activities [\[53–55\]](#). On the level of general predictions, sampling models provide a natural explanation for neural variability and co-variability [\[56\]](#), as stochastic samples vary in order to represent uncertainty. They also provide an account of bistable perception, and its neural correlates [\[57\]](#): multiple interpretations of the input correspond to multiple modes in the probability distribution over features; sequential sampling from such a distribution would produce in alternation samples from one of the peaks, but not from both at the same time [\[54,58\]](#). Although to date there is no direct electrophysiological evidence reinforcing the idea that the cortex represents distributions through samples, sampling-based representations have been recently invoked to account for several psychophysical and behavioral phenomena including stochasticity in learning, language processing and decision-making [\[59–61\]](#).

Thus, although sampling-based representations are a promising alternative to PPCs, their neurobiological ramifications are much less explored at present. In the next section, we turn to spontaneous activity in the cortex, a phenomenon so far not discussed in the context of neural representation of uncertainty, and review its potential role in developing a sound theory of sampling-based models of probabilistic neural coding.

Spontaneous activity and sampling-based representations

Modeling neural variability in evoked responses is an important first step in going beyond the modeling of receptive fields, and it is increasingly recognized as a critical benchmark for models of cortical functioning, including those positing probabilistic computations [\[16,48\]](#). Another major challenge in this direction is to accommodate spontaneous activity recorded in the awake nervous system without specific stimulation ([Box 3](#)). From a signal processing standpoint, spontaneous activity has been long considered as nuisance elicited by various aspects of stochastic neural activity [\[62\]](#), even though some proposals exist that discuss potential benefits of noise in the nervous system [\[63\]](#). However, several recent studies showed that the level of spontaneous activity is surprisingly high in some areas, and that it has a pattern highly similar to that of stimulus-evoked activity ([Box 3](#)). These findings suggest that a very large component of high spontaneous activity is probably not noise but might have a functional role in cortical computation [\[64,65\]](#).

Under a sampling-based representational account, spontaneous activity could have a natural interpretation. In a probabilistic framework, if neural activities represent samples from a distribution over external variables, this distribution must be the so-called “posterior distribution”.

Box 2. Probabilistic representational schemes for inference and learning

Representing uncertainty associated with sensory stimuli requires neurons to represent the probability distribution of the environmental variables that are being inferred. One class of schemes called probabilistic population codes (PPCs) assumes that neurons correspond to parameters of this distribution (Figure 1a). A simple but highly unfeasible version of this scheme would be if different neurons encoded the elements of the mean vector and covariance matrix of a multivariate Gaussian distribution. At any given time, the activities of neurons in PPCs provide a complete description of the distribution by determining its parameters, making PPCs and other parametric representational schemes particularly suitable for real-time inference [80–82]. Given that, in general, the number of parameters required to specify a multivariate distribution scales exponentially with the number of its variables, a drawback of such schemes is that the number of neurons needed in an exact PPC representation would be exponentially large and with fewer neurons the representation becomes approximate. Characteristics of the family of representable probability distributions by this scheme are determined by the characteristics of neural tuning curves and noisiness [16] (Table I).

An alternative scheme to represent probability distributions in neural activities is based on each neuron corresponding to one of the inferred variables. For example, each neuron can encode the value of one of the variables of a multivariate Gaussian distribution. In particular, the activity of a neuron at any time can represent a sample from the distribution of that variable and a “snapshot” of the activities of many neurons therefore can represent a sample from a high-dimensional distribution (Figure 1b). Such a representation requires time to take multiple samples (i.e. a sequence of firing rate measurements) for building up an increasingly reliable estimate of the represented distribution which might be prohibitive for on-line inference, but it does not require exponentially many neurons and – given enough time – it can represent any distribution (Table I). A further advantage of collecting samples is that marginalization, an important case of computing integrals that infamously plague practical Bayesian inference, learning and decision-making, becomes a straightforward neural operation. Finally, although it is unclear how probabilistic learning can be implemented with PPCs, sampling based representations seem particularly suitable for it (see main text).

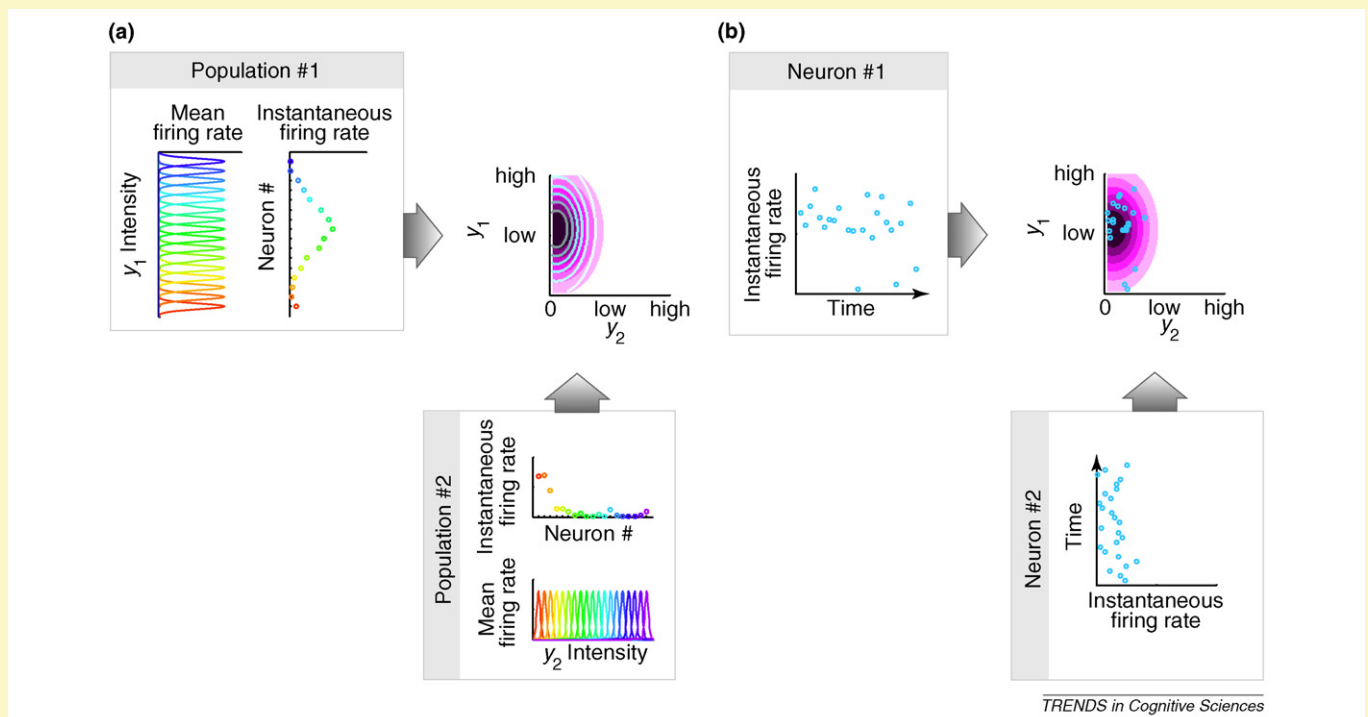


Figure 1. Two approaches to neural representations of uncertainty in the cortex. (a) Probabilistic population codes rely on a population of neurons that are tuned to the same environmental variables with different tuning curves (populations 1 and 2, colored curves). At any moment in time, the instantaneous firing rates of these neurons (populations 1 and 2, colored circles) determine a probability distribution over the represented variables (top right panel, contour lines), which is an approximation of the true distribution that needs to be represented (purple colormap). In this example, y_1 and y_2 , are independent, but in principle, there could be a single population with neurons tuned to both y_1 and y_2 . However, such multivariate representations require exponentially more neurons (see text and Table I). (b) In a sampling based representation, single neurons, rather than populations of neurons, correspond to each variable. Variability of the activity of neurons 1 and 2 through time represents uncertainty in environmental variables. Correlations between the variables can be naturally represented by co-variability of neural activities, thus allowing the representation of arbitrarily shaped distributions.

Table I. Comparing characteristics of the two main modeling approaches to probabilistic neural representations

| | PPCs | Sampling-based |
|--|---|---|
| Neurons correspond to | Parameters | Variables |
| Network dynamics required (beyond the first layer) | Deterministic | Stochastic (self-consistent) |
| Representable distributions | Must correspond to a particular parametric form | Can be arbitrary |
| Critical factor in accuracy of encoding a distribution | Number of neurons | Time allowed for sampling |
| Instantaneous representation of uncertainty | Complete, the whole distribution is represented at any time | Partial, a sequence of samples is required |
| Number of neurons needed for representing multimodal distributions | Scales exponentially with the number of dimensions | Scales linearly with the number of dimensions |
| Implementation of learning | Unknown | Well-suited |

Box 3. Spontaneous activity in the cortex

Spontaneous activity in the cortex is defined as ongoing neural activity in the absence of sensory stimulation [83]. This definition is the clearest in the case of primary sensory cortices where neural activity has traditionally been linked very closely to sensory input. Despite some early observations that it can influence behavior, cortical spontaneous activity has been considered stochastic noise [84]. The discovery of retinal and later cortical waves [85] of neural activity in the maturing nervous system has changed this view in developmental neuroscience, igniting an ongoing debate about the possible functional role of such spontaneous activity during development [86].

Several recent results based on the activities of neural populations initiated a similar shift in view about the role of spontaneous activity in the cortex during real-time perceptual processes [65]. Imaging and multi-electrode studies showed that spontaneous activity has large scale spatiotemporal structure over millimeters of the cortical surface, that the mean amplitude of this activity is comparable to that of evoked activity and it links distant cortical areas together [64,87,88] (Figure 1). Given the high energy cost of cortical spike activity [89], these findings argue against the idea of spontaneous activity being mere noise. Further investigations found that spontaneous activity shows repetitive patterns [90,91], it reflects the structure of the

underlying neural circuitry [67], which might represent visual attributes [66], that the second order correlational structure of spontaneous and evoked activity is very similar and it changes systematically with age [64]. Thus, cell responses even in primary sensory cortices are determined by the combination of spontaneous and bottom-up, external stimulus-driven activity.

The link between spontaneous and evoked activity is further promoted by findings that after repetitive presentation of a sensory stimulus, spontaneous activity exhibits patterns of activity reminiscent to those seen during evoked activity [92]. This suggests that spontaneous activity might be altered on various time scales leading to perceptual adaptation and learning. These results led to an increasing consensus that spontaneous activity might have a functional role in perceptual processes that is related to internal states of cell assemblies in the brain, expressed via top-down effects that embody expectations, predictions and attentional processes [93] and manifested in modulating functional connectivity of the network [94]. Although there have been theoretical proposals of how bottom-up and top-down signals could jointly define perceptual processes [55,95], the rigorous functional integration of spontaneous activity in such a framework has emerged only recently [53].

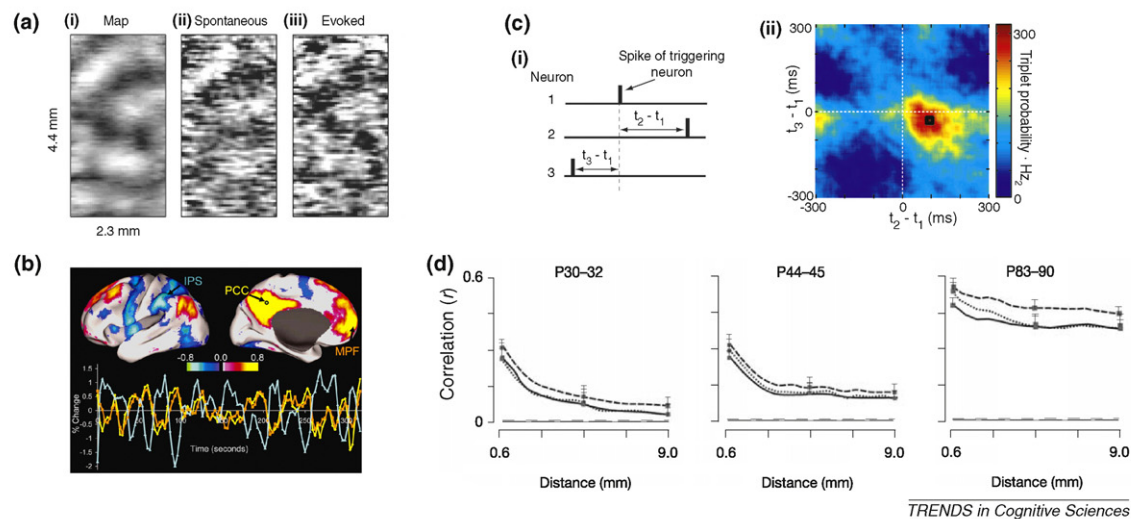


Figure 1. Characteristics of cortical spontaneous activity. (a) There is a significant correlation between the orientation map of the primary visual cortex of anesthetized cat (left panel), optical image patterns of spontaneous (middle panel) and visually evoked activities (right panel) (adapted with permission from [66]). (b) Correlational analysis of BOLD signals during resting state reveals networks of distant areas in the human cortex with coherent spontaneous fluctuations. There are large scale positive intrinsic correlations between the seed region PCC (yellow) and MPF (orange) and negative correlations between PCC and IPS (blue) (adapted with permission from [98]). (c) Reliably repeating spike triplets can be detected in the spontaneous firing of the rat somatosensory cortex by multielectrode recording (adapted with permission from [91]). (d) Spatial correlations in the developing awake ferret visual cortex of multielectrode recordings show a systematic pattern of emerging strong correlations across several millimeters of the cortical surface and very similar correlational patterns for dark spontaneous (solid line) and visually driven conditions (dotted and dashed lines for random noise patterns and natural movies, respectively) (adapted with permission from [64]).

The posterior distribution is inferred by combining information from two sources: the sensory input, and the prior distribution describing *a priori* beliefs about the sensory environment (Figure 3b). Intuitively, in the absence of sensory stimulation, this distribution will collapse to the prior distribution, and spontaneous activity will represent this prior (Figure 4).

This proposal linking spontaneous activity to the prior distribution has implications that can address many of the issues developed in this review. It provides an account of spontaneous activity that is consistent with one of its main features: its remarkable similarity to evoked activity [64,66,67]. A general feature of statistical models that are appropriately describing their inputs is that the prior distribution and the average posterior distribution closely match each other [68]. Thus, if evoked and spontaneous

activities represent samples from the posterior and prior distributions, respectively, under an appropriate model of the environment, they are expected to be similar [53]. In addition, spontaneous activity itself, as prior expectation, should be sufficient to evoke firing in some cells without sensory input, as was observed experimentally [67].

Statistical neural networks also suggest that sampling from the prior can be more than just a byproduct of probabilistic inference: it can be computationally advantageous for the functioning of the network. In the absence of stimulation, during awake spontaneous activity, sampling from the prior can help with driving the network close to states that are probable to be valid inferences once input arrives, thus potentially shortening the reaction time of the system [69]. This “priming” effect could present an alternative account of why human subjects are able to sort

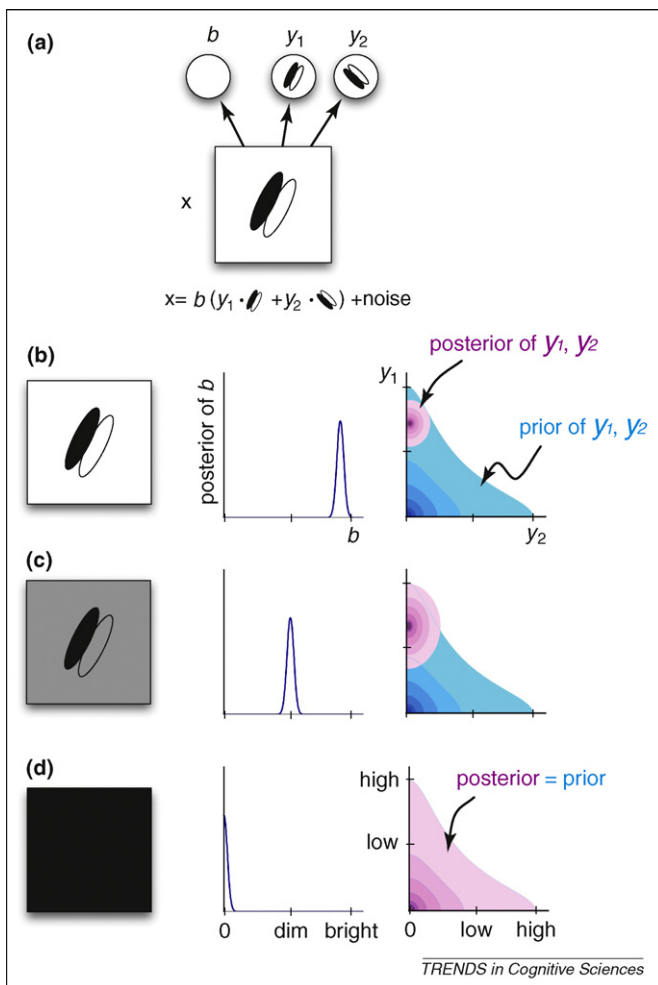


Figure 4. Relating spontaneous activity in darkness to sampling from the prior, based on the encoding of brightness in the primary visual cortex. **(a)** A statistically more efficient toy model of the early visual system [47,99] (Figure 3b). An additional feature variable, b , has a multiplicative effect on other features, effectively corresponding to the overall luminance. Explaining away this information removes redundant correlations thus improving statistical efficiency. **(b–c)** Probabilistic inference in such a model results in a luminance-invariant behavior of the other features, as observed neurally [100] as well as perceptually [101]: when the same image is presented at different global luminance levels (left), this difference is captured by the posterior distribution of the “brightness” variable, b (center), whereas the posterior for other features, such as y_1 and y_2 , remains relatively unaffected (right). **(d)** In the limit of total darkness (left), the same luminance-invariant mechanism results in the posterior over y_1 and y_2 collapsing to the prior (right). In this case, the inferred brightness, b , is zero (center) and as b explains all of the image content, there is no constraint left for the other feature variables, y_1 and y_2 (the identity in a becomes $0 = 0 (y_1 \cdot w_1 + y_2 \cdot w_2)$, which is fulfilled for every value of y_1 and y_2).

images into natural/non-natural categories in a matter of ~ 150 ms [70], which is traditionally taken as evidence for the dominance of feed-forward processing in the visual system [71]. Finally, during off-line periods, such as sleeping, sampling from the prior could have a role in tuning synaptic weights thus contributing to the refinement of the internal model of the sensory environment as suggested by statistical neural network models [32,33].

Importantly, the proposal that spontaneous activity represents samples from the prior also provides a way to test a direct link between statistically optimal inference and learning. A match between the prior and the average posterior distribution in a statistical model is expected to develop gradually as learning proceeds [68], and this gradual match could be tracked experimentally by comparing

Box 4. Questions for future research

- Exact probabilistic computation in the brain is not feasible. What are the approximations that are implemented in the brain and to what extent can an approximate computation scheme still claim that it is probabilistic and optimal?
- Probabilistic learning is presently described at the neural level as a simple form of parameter learning (so-called maximum likelihood learning) at best. However, there is ample behavioral evidence for more sophisticated forms of probabilistic learning, such as model selection. These forms of learning require a representation of uncertainty about parameters, or models, not just about hidden variables. How do neural circuits represent parameter uncertainty and implement model selection?
- Highly structured neural activity in the absence of external stimulation has been observed both in the neocortex and in the hippocampus, under the headings “spontaneous activity” and “replay”, respectively. Despite the many similarities these processes show there has been little attempt to study them in a unified framework. Are the two phenomena related, is there a common function they serve?
- Can a convergence between spontaneous and evoked activities be predicted from premises that are incompatible with spontaneous activity representing samples from the prior, for example with simple correlational learning schemes?
- Can some recursive implementation of probabilistic learning link learning of low-level attributes, such as orientations, with high-level concept learning, that is, can it bridge the subsymbolic and symbolic levels of computation?
- What is the internal model according to which the brain is adapting its representation? All the probabilistic approaches have preset prior constraints that determine how inference and learning will work. Where do these constraints come from? Can they be mapped to biological quantities?

spontaneous and evoked population activities at successive developmental stages. Such testable predictions can confirm if sampling-based representations are present in the cortex and verify the proposed link between spontaneous activity and sampling-based coding.

Concluding remarks and future challenges

In this review, we have argued that in order to develop a unified framework that can link behavior to neural processes of both inference and learning, a key issue to resolve is the nature of neural representations of uncertainty in the cortex. We compared potential candidate neural codes that could link behavior to neural implementations in a probabilistic way by implementing computations with and learning of probability distributions of environmental features. Although explored to different extents, these coding frameworks are all promising candidates, yet each of them has shortcomings that need to be addressed in future research (Box 4). Research on PPCs needs to make viable proposals on how learning could be implemented with such representations, whereas the main challenge for sampling-based methods is to demonstrate that this scheme could work for non-trivial, dynamical cases in real time.

Most importantly, a tighter connection between abstract computational models and neurophysiological recordings in behaving animals is needed. For PPCs, such interactions between theoretical and empirical investigations have just begun [50]; for sampling-based methods it is still almost non-existent beyond the description of receptive fields. Present day data collection methods, such

as modern imaging techniques and multi-electrode recording systems, are increasingly available to provide the necessary experimental data for evaluating the issues raised in this review. Given the complex and non-local nature of computation in probabilistic frameworks, the main theoretical challenge remains to map abstract probabilistic models to neural activity in awake behaving animals to further our understanding of cortical representations of inference and learning.

Acknowledgements

This work was supported by the Swartz Foundation (J.F., G.O., P.B.), by the Swiss National Science Foundation (P.B.) and the Wellcome Trust (M.L.). We thank Peter Dayan, Maneesh Sahani and Jeff Beck for useful discussions.

References

- 1 Barlow, H.B. (1961) Possible principles underlying the transformations of sensory messages. In *Sensory Communication* (Rosenblith, W., ed.), pp. 217–234, MIT Press
- 2 Geisler, W.S. and Diehl, R.L. (2002) Bayesian natural selection and the evolution of perceptual systems. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 357, 419–448
- 3 Smith, J.D. (2009) The study of animal metacognition. *Trends Cogn. Sci.* 13, 389–396
- 4 Pouget, A. et al. (2004) Inference and computation with population codes. *Annu. Rev. Neurosci.* 26, 381–410
- 5 Helmholtz, H.V. (1925) *Treatise on Physiological Optics*, Optical Society of America
- 6 Green, D.M. and Swets, J.A. (1966) *Signal Detection Theory and Psychophysics*, John Wiley and Sons
- 7 Cox, R.T. (1946) Probability, frequency and reasonable expectation. *Am. J. Phys.* 14, 1–13
- 8 Bernoulli, J. (1713) *Ars Conjectandi*, Thurnisiorum
- 9 Atkins, J.E. et al. (2001) Experience-dependent visual cue integration based on consistencies between visual and haptic percepts. *Vision Res.* 41, 449–461
- 10 Ernst, M.O. and Banks, M.S. (2002) Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* 415, 429–433
- 11 Kording, K.P. and Wolpert, D.M. (2004) Bayesian integration in sensorimotor learning. *Nature* 427, 244–247
- 12 Weiss, Y. et al. (2002) Motion illusions as optimal percepts. *Nat. Neurosci.* 5, 598–604
- 13 Trommershäuser, J. et al. (2008) Decision making, movement planning and statistical decision theory. *Trends Cogn. Sci.* 12, 291–297
- 14 Kording, K. (2007) Decision theory: What “should” the nervous system do? *Science* 318, 606–610
- 15 Zemel, R. et al. (1998) Probabilistic interpretation of population codes. *Neural Comput.* 10, 403–430
- 16 Ma, W.J. et al. (2006) Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438
- 17 Beck, J.M. et al. (2008) Probabilistic population codes for Bayesian decision making. *Neuron* 60, 1142–1152
- 18 Jacobs, R.A. et al. (1991) Adaptive mixtures of local experts. *Neural Comput.* 3, 79–87
- 19 Geisler, W.S. et al. (2001) Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res.* 41, 711–724
- 20 Griffiths, T.L. and Tenenbaum, J.B. (2006) Optimal predictions in everyday cognition. *Psychol. Sci.* 17, 767–773
- 21 Chater, N. et al. (2006) Probabilistic models of cognition: conceptual foundations. *Trends Cogn. Sci.* 10, 287–291
- 22 Courville, A.C. et al. (2004) Model uncertainty in classical conditioning. In *Advances in Neural Information Processing Systems 16* (Thrun, S. et al., eds), pp. 977–984, MIT Press
- 23 Blaisdell, A.P. et al. (2006) Causal reasoning in rats. *Science* 311, 1020–1022
- 24 Gopnik, A. et al. (2004) A theory of causal learning in children: causal maps and Bayes nets. *Psychol. Rev.* 111, 3–32
- 25 Kemp, C. and Tenenbaum, J.B. (2008) The discovery of structural form. *Proc. Natl. Acad. Sci. U. S. A.* 105, 10687–10692
- 26 Orbán, G. et al. (2008) Bayesian learning of visual chunks by human observers. *Proc. Natl. Acad. Sci. U. S. A.* 105, 2745–2750
- 27 Buonomano, D.V. and Merzenich, M.M. (1998) Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* 21, 149–186
- 28 Gilbert, C.D. et al. (2009) Perceptual learning and adult cortical plasticity. *J. Physiol. Lond.* 587, 2743–2751
- 29 Deneve, S. (2008) Bayesian spiking neurons II: learning. *Neural Comput.* 20, 118–145
- 30 Lengyel, M. and Dayan, P. (2007) Uncertainty, phase and oscillatory hippocampal recall. In *Advances in Neural Information Processing Systems 19* (Scholkopf, B. et al., eds), pp. 833–840, MIT Press
- 31 Rumelhart, D.E. et al., eds (1986) *Parallel Distributed Processing – Explorations in the Microstructure of Cognition*, MIT Press
- 32 Hinton, G.E. et al. (1995) The wake–sleep algorithm for unsupervised neural networks. *Science* 268, 1158–1161
- 33 Hinton, G.E. and Sejnowski, T.J. (1986) Learning and relearning in Boltzmann machines. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (Rumelhart, D.E. and McClelland, J.L., eds), MIT Press
- 34 Neal, R.M. (1996) *Bayesian Learning for Neural Networks*, Springer-Verlag
- 35 Hinton, G.E. (2007) Learning multiple layers of representation. *Trends Cogn. Sci.* 11, 428–434
- 36 Dayan, P. (1999) Recurrent sampling models for the Helmholtz machine. *Neural Comput.* 11, 653–677
- 37 Dayan, P. and Hinton, G.E. (1996) Varieties of Helmholtz machine. *Neural Netw.* 9, 1385–1403
- 38 Hinton, G.E. and Ghahramani, Z. (1997) Generative models for discovering sparse distributed representations. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 352, 1177–1190
- 39 Lucke, J. and Sahani, M. (2008) Maximal causes for non-linear component extraction. *J. Mach. Learn. Res.* 9, 1227–1267
- 40 Bell, A.J. and Sejnowski, T.J. (1997) The “independent components” of natural scenes are edge filters. *Vision Res.* 37, 3327–3338
- 41 Berkes, P. and Wiskott, L. (2005) Slow feature analysis yields a rich repertoire of complex cell properties. *J. Vis.* 5, 579–602
- 42 Karklin, Y. and Lewicki, M.S. (2009) Emergence of complex cell properties by learning to generalize in natural scenes. *Nature* 457, 83–85
- 43 Olshausen, B.A. and Field, D.J. (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607–609
- 44 Rao, R.P.N. and Ballard, D.H. (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87
- 45 Roweis, S. and Ghahramani, Z. (1999) A unifying review of linear Gaussian models. *Neural Comput.* 11, 305–345
- 46 Turner, R. and Sahani, M. (2007) A maximum-likelihood interpretation for slow feature analysis. *Neural Comput.* 19, 1022–1038
- 47 Schwartz, O. and Simoncelli, E.P. (2001) Natural signal statistics and sensory gain control. *Nat. Neurosci.* 4, 819–825
- 48 Olshausen, B.A. and Field, D.J. (2005) How close are we to understanding V1? *Neural Comput.* 17, 1665–1699
- 49 Lopez-Aranda, M.F. et al. (2009) Role of layer 6 of V2 visual cortex in object-recognition memory. *Science* 325, 87–89
- 50 Yang, T. and Shadlen, M.N. (2007) Probabilistic reasoning by neurons. *Nature* 447, 1075–1080
- 51 Anderson, C.H. and Van Essen, D.C. (1994) Neurobiological computational systems. In *Computational Intelligence Imitating Life* (Zurada, J.M. et al., eds), pp. 213–222, IEEE Press
- 52 Ma, W.J. et al. (2008) Spiking networks for Bayesian inference and choice. *Curr. Opin. Neurobiol.* 18, 217–222
- 53 Berkes, P. et al. (2009) Matching spontaneous and evoked activity in V1: a hallmark of probabilistic inference. In *Frontiers in Systems Neuroscience. Conference Abstract: Computational and systems neuroscience DOI:10.3389/conf.neuro.06.2009.03.314*
- 54 Hoyer, P.O. and Hyvarinen, A. (2003) Interpreting neural response variability as Monte Carlo sampling of the posterior. In *Advances in Neural Information Processing Systems 15* (Becker, S. et al., eds), pp. 277–284, MIT Press

- 55 Lee, T.S. and Mumford, D. (2003) Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. (A)* 20, 1434–1448
- 56 Zohary, E. *et al.* (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370, 140–143
- 57 Leopold, D.A. and Logothetis, N.K. (1996) Activity changes in early visual cortex reflect monkeys' percepts during binocular rivalry. *Nature* 379, 549–553
- 58 Sundareswara, R. and Schrater, P.R. (2008) Perceptual multistability predicted by search model for Bayesian decisions. *J. Vis.* 8, 12.1–12.19
- 59 Daw, N. and Courville, A. (2008) The rat as particle filter. In *Advances in Neural Information Processing Systems 20* (Platt, J. *et al.*, eds), pp. 369–376, MIT Press
- 60 Levy, R.P., Reali, F. and Griffiths, T.L. (2009) Modeling the effects of memory on human online sentence processing with particle filters. In *Advances in Neural Information Processing Systems 21* (Koller, P. *et al.*, eds), pp. 937–944, MIT Press
- 61 Vul, E. and Pashler, H. (2008) Measuring the crowd within – probabilistic representations within individuals. *Psychol. Sci.* 19, 645–647
- 62 Shadlen, M.N. and Newsome, W.T. (1994) Noise, neural codes and cortical organization. *Curr. Opin. Neurobiol.* 4, 569–579
- 63 Anderson, J.S. *et al.* (2000) The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science* 290, 1968–1972
- 64 Fiser, J. *et al.* (2004) Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 431, 573–578
- 65 Ringach, D. (2009) Spontaneous and driven cortical activity: implications for computation. *Curr. Opin. Neurobiol.* 19, 439–444
- 66 Kenet, T. *et al.* (2003) Spontaneously emerging cortical representations of visual attributes. *Nature* 425, 954–956
- 67 Tsodyks, M. *et al.* (1999) Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 286, 1943–1946
- 68 Dayan, P. and Abbott, L.F. (2001) *Theoretical Neuroscience*, MIT Press
- 69 Neal, R.M. (2001) Annealed importance sampling. *Stat. Comput.* 11, 125–139
- 70 Thorpe, S. *et al.* (1996) Speed of processing in the human visual system. *Nature* 381, 520–522
- 71 Serre, T. *et al.* (2007) A feedforward architecture accounts for rapid categorization. *Proc. Natl. Acad. Sci. U. S. A.* 104, 6424–6429
- 72 Saffran, J.R. *et al.* (1996) Statistical learning by 8-month-old infants. *Science* 274, 1926–1928
- 73 Seitz, A.R. *et al.* (2007) Simultaneous and independent acquisition of multisensory and unisensory associations. *Perception* 36, 1445–1453
- 74 Pena, M. *et al.* (2002) Signal-driven computations in speech processing. *Science* 298, 604–607
- 75 Turk-Browne, N.B. *et al.* (2009) Neural evidence of statistical learning: efficient detection of visual regularities without awareness. *J. Cogn. Neurosci.* 21, 1934–1945
- 76 Fiser, J. and Aslin, R.N. (2001) Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychol. Sci.* 12, 499–504
- 77 Fiser, J. and Aslin, R.N. (2002) Statistical learning of new visual feature combinations by infants. *Proc. Natl. Acad. Sci. U. S. A.* 99, 15822–15826
- 78 Kirkham, N.Z. *et al.* (2002) Visual statistical learning in infancy: evidence for a domain general learning mechanism. *Cognition* 83, B35–B42
- 79 Fiser, J. and Aslin, R.N. (2005) Encoding multielement scenes: statistical learning of visual feature hierarchies. *J. Exp. Psychol. Gen.* 134, 521–537
- 80 Beck, J.M. and Pouget, A. (2007) Exact inferences in a neural implementation of a hidden Markov model. *Neural Comput.* 19, 1344–1361
- 81 Deneve, S. (2008) Bayesian spiking neurons I: inference. *Neural Comput.* 20, 91–117
- 82 Huys, Q.J.M. *et al.* (2007) Fast population coding. *Neural Comput.* 19, 404–441
- 83 Creutzfeldt, O.D. *et al.* (1966) Relations between EEG phenomena and potentials of single cortical cells: II. Spontaneous and convulsoid activity. *Electroencephalogr. Clin. Neurophysiol.* 20, 19–37
- 84 Tolhurst, D.J. *et al.* (1983) The statistical reliability of signals in single neurons in cat and monkey visual-cortex. *Vision Res.* 23, 775–785
- 85 Wu, J.Y. *et al.* (2008) Propagating waves of activity in the neocortex: what they are, what they do. *Neuroscientist* 14, 487–502
- 86 Katz, L.C. and Shatz, C.J. (1996) Synaptic activity and the construction of cortical circuits. *Science* 274, 1133–1138
- 87 Arieli, A. *et al.* (1995) Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *J. Neurophysiol.* 73, 2072–2093
- 88 Fox, M.D. and Raichle, M.E. (2007) Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat. Rev. Neurosci.* 8, 700–711
- 89 Attwell, D. and Laughlin, S.B. (2001) An energy budget for signaling in the grey matter of the brain. *J. Cereb. Blood Flow Metab.* 21, 1133–1145
- 90 Ikegaya, Y. *et al.* (2004) Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 304, 559–564
- 91 Luczak, A. *et al.* (2007) Sequential structure of neocortical spontaneous activity in vivo. *Proc. Natl. Acad. Sci. U. S. A.* 104, 347–352
- 92 Han, F. *et al.* (2008) Reverberation of recent visual experience in spontaneous cortical waves. *Neuron* 60, 321–327
- 93 Gilbert, C.D. and Sigman, M. (2007) Brain states: top-down influences in sensory processing. *Neuron* 54, 677–696
- 94 Nauhaus, I. *et al.* (2009) Stimulus contrast modulates functional connectivity in visual cortex. *Nat. Neurosci.* 12, 70–76
- 95 Yuille, A. and Kersten, D. (2006) Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* 10, 301–308
- 96 Ernst, M.O. and Bulthoff, H.H. (2004) Merging the senses into a robust percept. *Trends Cogn. Sci.* 8, 162–169
- 97 Kording, K.P. and Wolpert, D.M. (2006) Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.* 10, 319–326
- 98 Fox, M.D. *et al.* (2005) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc. Natl. Acad. Sci. U. S. A.* 102, 9673–9678
- 99 Berkes, P. *et al.* (2009) A structured model of video reproduces primary visual cortical organisation. *PLoS Comput. Biol.* 5, e1000495
- 100 Rossi, A.F. *et al.* (1996) The representation of brightness in primary visual cortex. *Science* 273, 1104–1107
- 101 Adelson, E.H. (1993) Perceptual organization and the judgment of brightness. *Science* 262, 2042–2044