

Accumulation of Evidence during Sequential Decision Making: The Importance of Top–Down Factors

Floris P. de Lange,^{1,2} Ole Jensen,¹ and Stanislas Dehaene^{2,3,4,5}

¹Donders Institute for Brain, Cognition and Behaviour, Radboud University, 6525 EN Nijmegen, The Netherlands, ²Inserm, Cognitive Neuroimaging Unit, and ³Commissariat à l’Energie Atomique, Direction des Sciences du Vivant, 12BM, NeuroSpin Center, 91191 Gif-sur-Yvette, France, ⁴Université Paris-Sud 11, 91400 Orsay, France, and ⁵Collège de France, 75005 Paris, France

In the last decade, great progress has been made in characterizing the accumulation of neural information during simple unitary perceptual decisions. However, much less is known about how sequentially presented evidence is integrated over time for successful decision making. The aim of this study was to study the mechanisms of sequential decision making in humans. In a magnetoencephalography (MEG) study, we presented healthy volunteers with sequences of centrally presented arrows. Sequence length varied between one and five arrows, and the accumulated directions of the arrows informed the subject about which hand to use for a button press at the end of the sequence (e.g., LRLRR should result in a right-hand press). Mathematical modeling suggested that nonlinear accumulation was the rational strategy for performing this task in the presence of no or little noise, whereas quasilinear accumulation was optimal in the presence of substantial noise. MEG recordings showed a correlate of evidence integration over parietal and central cortex that was inversely related to the amount of accumulated evidence (i.e., when more evidence was accumulated, neural activity for new stimuli was attenuated). This modulation of activity likely reflects a top–down influence on sensory processing, effectively constraining the influence of sensory information on the decision variable over time. The results indicate that, when making decisions on the basis of sequential information, the human nervous system integrates evidence in a nonlinear manner, using the amount of previously accumulated information to constrain the accumulation of additional evidence.

Introduction

In the last decade, great progress has been made in characterizing the forward flow of neural information during simple perceptual decisions. Neurophysiological studies in monkeys have shown how parietal areas accumulate evidence during perceptual decisions. When a decision is expressed through a movement, there is an accumulation of activity related to the decision process in the cortical area that guides the motor action (Platt and Glimcher, 1999; Gold and Shadlen, 2000; Schall, 2001; Cisek and Kalaska, 2005; Gold and Shadlen, 2007). Similarly, imaging studies in humans showed that, when subjects make perceptual decisions about an object category, neuronal activity in category-specific areas gradually increases with increasing sensory evidence (Philiastides and Sajda, 2006; Kaiser et al., 2007), while the dorsolateral prefrontal cortex may integrate the output of these sensory regions to make the decision (Heekeren et al., 2004, 2006).

Although much has been learned about decisions on single percepts, much less is known about how multiple sources of evidence are combined over time to guide decision making. Consider the situation we face when deciding whether or not to cross

the street. Typically, we sequentially acquire different sources of information (“Is the traffic light green? What is the distance/speed of the cars approaching?”) and combine these sources of evidence, giving each of them a particular weight (e.g., information about the speed of the car may be more important than the traffic light). A recent study explored the neurophysiology of this type of sequential decisions (Yang and Shadlen, 2007). Monkeys learned to probabilistically associate shapes with one of two targets. During sequential presentation of the shapes, neurons in the LIP (lateral intraparietal area) showed responses consistent with these neurons combining the information from all cues over time, potentially by calculating the “log-likelihood ratio” (LLR) in favor of one response over the other, an optimal strategy for probabilistic inference.

Although the study of Yang and Shadlen shows that monkeys can combine sources of evidence over time, it remains a question how this information is integrated. The aim of this study was to study the mechanisms of sequential decision making in humans, contrasting several potential strategies that can be used for integrating evidence. We designed a task in which a sequence of arrow shapes was shown to the subject. Sequence length varied between 1 and 5. The subject’s task was to decide on the predominant direction of the arrows as soon as a go cue appeared (see Fig. 1). We tracked neural activity using magnetoencephalography (MEG) during the sequential presentation of these cues that informed the subject about the correct decision. We distinguish between three strategies that could be used to solve the task: simple addition/subtraction of each of the shapes (“naive” accu-

Received Aug. 12, 2009; revised Oct. 14, 2009; accepted Nov. 20, 2009.

This work was supported by The Netherlands Organization for Scientific Research Grant 446-07-003 (awarded to F.P.d.L.). We thank Robert Oostenveld for help with methodological issues.

Correspondence should be addressed to Floris P. de Lange, Inserm–Commissariat à l’Energie Atomique (CEA) Cognitive Neuroimaging Unit, CEA/SAC/DSV/DRM/NeuroSpin, Bât 145, Point Courrier 156, 91191 Gif-sur-Yvette, France. E-mail: florisdelange@gmail.com.

DOI:10.1523/JNEUROSCI.4080-09.2010

Copyright © 2010 the authors 0270-6474/10/300731-08\$15.00/0

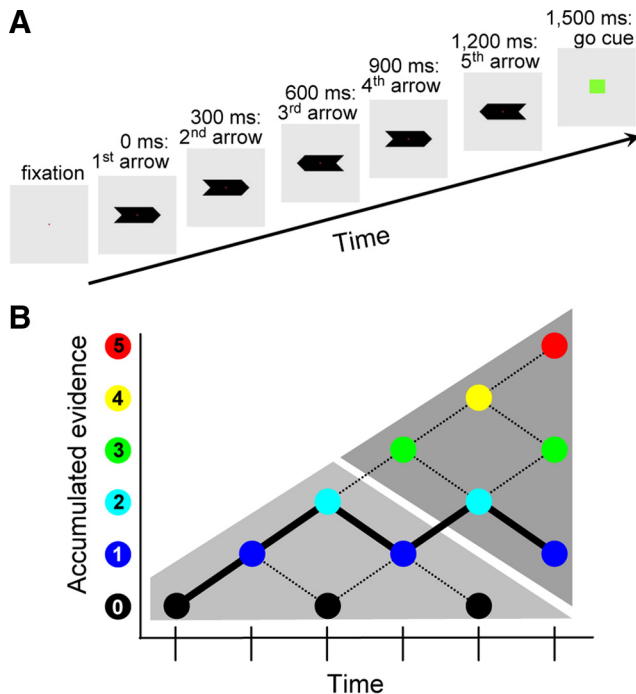


Figure 1. Task setup and diagram. **A**, Subjects were shown sequences of one to five arrows. Each arrow was shown for 100 ms followed by a 200 ms blank screen, leading to a stimulus onset asynchrony between two arrows of 300 ms. At the end of each arrow sequence, the fixation square turned green, and subjects had to decide as fast as possible whether the predominant direction of the arrow stimuli was left or right, by pressing a button with their left or right hand. **B**, The cumulative number arrows favoring one or the other direction, which is used here as a proxy for accumulated evidence (see mathematical model), is plotted as a function of time. Subjects start with no evidence for either direction (0, in black). Each incoming arrow moves the sum up or down in the diagram. The transitions of the example sequence of **A** are shown in solid lines in the diagram. States at which no decision can be made yet are shaded light gray, whereas states at which enough information has been accrued for a decision are shaded dark gray.

mulation), computing the LLR, and a variant of the latter that we call “mental logic.” These three strategies make different predictions about how new evidence is integrated when previous evidence has already biased the decision. We present a mathematical model that shows how these three strategies emerge naturally from an optimal Bayesian observer that combines information over time under various levels of stimulus uncertainty, and behavioral and MEG data that suggests how human subjects combine information over time.

Materials and Methods

Participants. Fourteen healthy participants (11 males, 3 females; age, mean \pm SD, 28 ± 3 years) participated in the experiment. None of the participants had a history of neurological or psychiatric disorders. All participants had normal or corrected-to-normal vision. The study was approved by the regional ethics committee (Hôpital de Bicêtre, Paris, France), and a written informed consent was obtained from the subjects according to the Declaration of Helsinki. The experiment was part of a general research program on functional neuroimaging of the human brain, which was sponsored by the Atomic Energy Commission (Denis Le Bihan).

Stimuli. The experimental stimuli were leftward and rightward pointing arrows (Fig. 1). Stimuli were black, presented on a gray background, and subtended visual angles of $2.1 \times 1.0^\circ$. Stimuli were presented using a PC running Presentation software (Neurobehavioral Systems) and projected on a screen that was ~ 70 cm away from the subject.

Experimental design. Subjects were shown sequences of one to five successive foveal arrows. Each trial started with a red fixation square, displayed for an average duration of 2000 ms (between 1750 and 2250

ms), followed by the sequence of arrows. Each arrow was presented for 100 ms, followed by a blank of 200 ms. At the end of each arrow sequence, the fixation square turned green, and the subject had to decide as fast as possible whether the predominant direction of the arrow stimuli was left or right, by pressing a button with their left or right hand (Fig. 1A). Sequences of length 1–4 were included to encourage subjects to keep actively integrating information during the whole epoch, instead of waiting until all arrows had appeared (which could be the strategy of choice when subjects know beforehand how many stimuli will be presented), but these trials were not otherwise included in the analysis. Figure 1B shows a diagram of the amount of accumulated evidence as a function of time, illustrating all the possible transitions during the task.

To encourage continuous updating of information, subjects had to respond within a 500 ms time window. Visual feedback about the subject’s performance was given on a trial-by-trial basis. Before data acquisition, participants engaged in 60 training trials to get acquainted with the task. During MEG data acquisition, subjects engaged in six task blocks, each block consisting of 120 trials. Total duration of the experiment was ~ 70 min.

Mathematical model of information processing during the task. To examine whether behavior and electrophysiological variables track optimal decision variables of the task, we constructed a mathematical model of the evolution of the LLR of one of the responses being correct (vs the other) at each time step. The model takes into account both the uncertainty about the forthcoming stimuli (given that all sequences of five arrows are equiprobable), and the additional uncertainty associated with the probability p of correctly perceiving each stimulus.

Mathematically, we describe this stimulus by a series of n numbers $\{x_i\}$, each of which is either -1 or 1 . The task requires pressing right if $\sum(x_i) > 0$, and left otherwise. Our goal is to compute $P(R|x_1 \dots x_k)$, when R is the right-hand response and x_k is the k th stimulus. The log-likelihood ratio will then simply be as follows: $\text{LLR} = \log_{10}[P(R|x_1 \dots x_k)/(1 - P(R|x_1 \dots x_k))]$. Consider first the case when $p = 1$ (arrows are always correctly perceived). Computing $P(R|x_1 \dots x_k)$, where $x_1 \dots x_k$ is a given initial sequence of k arrows ($k < n$), amounts to counting, among all the possible sequences of the remaining $n - k$ stimuli, the proportion of those that have a rightward response. Call the sum of the first k arrows s_k . After k arrows have been presented, the tree of possible stimuli is reduced to 2^{n-k} possibilities, whose remaining sum s_{n-k} is distributed according to a binomial distribution [given simply by the $(n - k)$ th line of the Pascal triangle]. Noting that $P(R|x_1 \dots x_k) = P(R|s_k) = P(s_{n-k} + s_k > 0) = P(s_{n-k} > -s_k)$, we see that the desired probability can be obtained by computing the proportion of trials, in the $(n - k)$ th line of the Pascal triangle, that exceed the midpoint minus the current sum s_k .

Note, however, that, with $p = 1$, the LLR will rise to infinity when the subject acquires certainty about his response. A more realistic model assumes that subjects have a fixed probability $p < 1$ of correctly perceiving each arrow because of noise in the sensory signal and inference. The revised theory therefore distinguishes the arrows $\{X_i\}$ objectively presented on a given trial, from the perceived arrows $\{x_i\}$. The optimal observer perceives the $\{x_i\}$ and attempts to compute $P(R|x_1 \dots x_k)$, taking into account that some of the x_i may have been misperceived. Thus, he needs to consider all the possible original stimuli $X_1 \dots X_k$ that could have led to his percept, each weighted by their probability of having been perceived $P(X_1 \dots X_k|x_1 \dots x_k)$ as follows: $P(R|x_1 \dots x_k) = \text{Sum}_{\text{all } X_1 \dots X_k} P(R|X_1 \dots X_k) P(X_1 \dots X_k|x_1 \dots x_k)$. This sum is easily computed as $P(R|x_1 \dots x_k)$ is known from the previous theory; it is $P(R|S_k)$, where $S_k = \text{sum}(X_k)$, and $P(X_1 \dots X_k|x_1 \dots x_k)$ is $p^a(1 - p)^{k-a}$, where a is the number of arrows in common between X and x .

From these equations, a MatLab program was written to compute the LLR for all possible input sequences. For simplicity, we only report here the results of simulations at three levels of perceptual certainty: absolute certainty ($p = 1$), high certainty ($p = 0.95$), and low certainty ($p = 0.6$).

MEG measurements. Ongoing brain activity was recorded using a whole-head MEG with 275 axial gradiometers (VSM/CTF Systems). Data were collected at 12 kHz and downsampled at 600 Hz. To prevent aliasing, an eighth-order elliptic infinite impulse response filter was used with a cutoff at one-fourth of the sampling frequency, a 0.1 dB pass band ripple, and 120 dB attenuation at the Nyquist frequency as an antialiasing

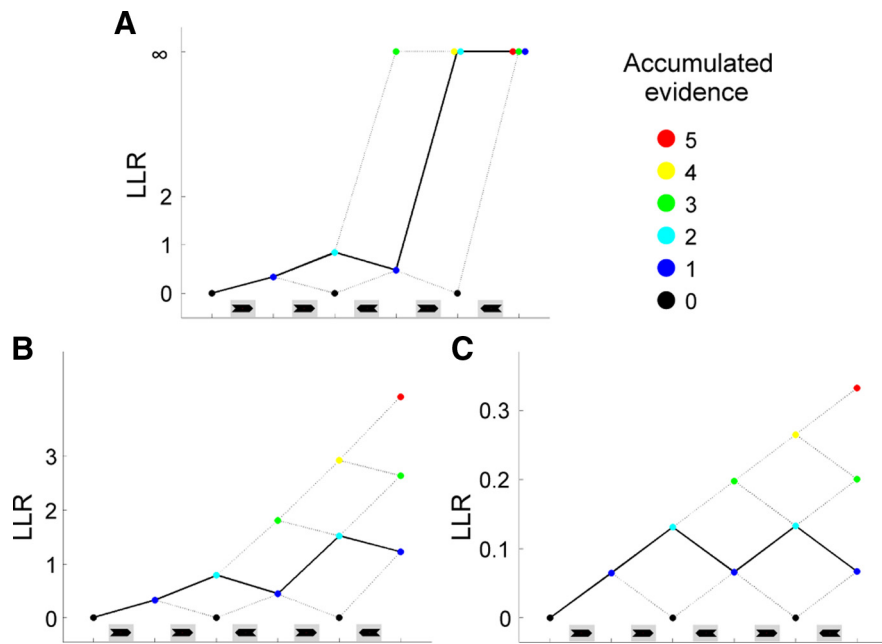


Figure 2. Predictions from mathematical model. The LLR is plotted as a function of the amount of accumulated evidence and three scenarios of perceptual uncertainty: complete perceptual certainty ($p = 1$) (A), high amount of perceptual certainty ($p = 0.95$) (B), and low amount of perceptual certainty ($p = 0.6$) (C).

filter. Head localization was monitored continuously during the experiment using coils that were placed at the cardinal points of the head (nasion, left and right ear canal). The magnetic fields produced by these coils were used to measure the position of the subject's head with respect to the MEG sensor array. In addition to the MEG, the electro-oculogram (EOG) was recorded from the supraorbital and infraorbital ridge of the left eye for the subsequent artifact rejection. Also, the electromyogram (EMG) and electrocardiogram (ECG) were recorded using 10-mm-diameter Ag–AgCl surface electrodes. EMG electrodes were placed on the left and right forearm, in a “belly–tendon” arrangement, after standard skin preparation.

Data analysis. All data analysis was performed using the FieldTrip toolbox developed at Donders Institute for Brain, Cognition and Behaviour (<http://www.ru.nl/fcdonders/fieldtrip>) using MatLab 7 (Mathworks). Data were checked for artifacts using a semiautomatic routine that helped detecting and rejecting eye blinks, muscle artifacts, and jumps in the MEG signal caused by the superconducting quantum interference device (SQUID) electronics. Subsequently, independent component analysis (Bell and Sejnowski, 1995) was used to remove any heart artifacts and eye movements not rejected by the semiautomatic routine (Jung et al., 2000). First, 275 independent components (ICs) were estimated from the full dataset. Then, ICs that had both a strong correlation with the ECG or EOG channels and a topography that was consistent with the corresponding ocular or heartbeat artifact were removed from the data set. Finally, we low-pass filtered the data using a two-pass Butterworth filter with a filter order of 6, and a frequency cutoff of 40 Hz.

For the subsequent analysis, we only considered correct trials with the longest sequence length (i.e., consisting of five arrows). We calculated an estimate of the planar gradient for the data analysis on the sensor level (Bastiaansen and Knösche, 2000). The horizontal and vertical components of the planar gradients were calculated for each sensor using the signals from the neighboring sensors, thus approximating the signal measured by MEG systems with planar gradiometers. This approach has been successfully applied in previous MEG studies (Bauer et al., 2006; Osipova et al., 2006; de Lange et al., 2008). The planar field gradient simplifies the interpretation of the sensor-level data because the maximal signal typically is located above the source (Hamalainen et al., 1993).

We used a multiple regression analysis to isolate the unique contributions of the following experimental factors: (1) the number of accumulated arrows at stimulus onset [hereafter called “evidence” for simplicity,

given its theoretical correlation with LLR; range (0–5)], (2) whether the stimulus changed direction with respect to the preceding stimulus or not [“change”; 0 (no change) or 1 (change)], and (3) the temporal position in the sequence at which the stimulus was presented [“time”; range (1–5)]. Although evidence and change are orthogonal with respect to each other, there is a degree of correlation between evidence and time ($r = 0.42$), necessitating a multiple regression approach to discern the unique contributions of each of these factors from their commonly explained variance.

We established the significance of the differences in field strength for each experimental factor at the cluster level, using a nonparametric cluster randomization test (Nichols and Holmes, 2002; Maris and Oostenveld, 2007). This test effectively controls the type I error rate in situations involving multiple comparisons (such as 275 sensors) by clustering neighboring sensor pairs that exhibit the same effect. The randomization method first identified sensors whose t statistics exceeded a critical value when comparing two conditions sensor by sensor ($p < 0.01$, two-sided). In the second step, to correct for multiple comparisons, contiguous sensors (separated by < 5 cm) that exceeded the critical value (as defined in the first step) were considered a cluster. The cluster-

level test statistic was defined from the sum of the t values of the sensors in a given cluster. The cluster with the maximum sum was used in the test statistics. The type I error rate for the complete set of 275 sensors was controlled by evaluating the cluster level test statistic under the randomization null distribution of the maximum cluster-level test statistic. This was obtained by randomizing the data between the two conditions across multiple subjects calculating t statistics for the new set of clusters. A reference distribution of cluster-level t statistics was created from 2000 randomizations. The p value was estimated according to the proportion of the randomization null distribution exceeding the observed maximum cluster-level test statistic (the so-called Monte Carlo p value). The anatomical label of each reported cluster refers to the localization scheme used for labeling the SQUID sensors (VSM/CTF Systems).

First, we assessed the effects of previous evidence and time on the evoked fields just after stimulus onset [interval (0–50 ms)], where the current arrow did not have time to impact on brain activity. This “baseline” interval was physiologically motivated, for it takes ~ 50 ms for a visual stimulus to reach the cortex (Bullier and Nowak, 1995; Nowak et al., 1995). To investigate how each new arrow influenced subsequent stimulus processing, we assessed the effects of evidence, change, and time on evoked fields in two time windows: early (150–225 ms) and late (225–300 ms) after stimulus onset. Activity changes, over and above the “baseline” activity level at stimulus onset, were obtained by subtracting the average activity of the baseline interval (0–50 ms).

Results

Mathematical model of information processing during the task

Figure 2 shows how the theoretical LLR evolves as a function of time (the abscissa) and of the evidence provided by the different types of sequential stimuli [the color of the curves: from black, no evidence (equal number of left and right arrows), to red, five arrows pointing in the same direction]. Interestingly, by changing the value of a single parameter, the uncertainty in the stimulus, we obtain three qualitatively different types of behavior. Under conditions of no noise in the system ($p = 1$), the LLR “jumps” to an infinite amount of certainty about the decision to make as soon as enough arrows have been observed that have the

same direction (Fig. 2A). After this point, new information no longer changes the decision certainty. This optimal strategy is what we referred to as “mental logic” in Introduction. It exploits all the information available, including that about the maximal sequence length. For instance, it takes into account that, given a maximum of five arrows, any set of three arrows pointing in the same direction suffices to obtain certainty, regardless of subsequent stimuli.

When a small amount of noise is added to the system ($p = 0.95$), the diagram shows a nonlinear increase in the amount of evidence at the moment that enough stimuli have been observed with the same direction (corresponding to achieving “near-certainty”), but new stimuli can still influence the LLR (Fig. 2B). Finally, when there is a significant amount of noise in the system ($p = 0.6$), the transition diagram becomes essentially tree-like, as the LLR is essentially linearly increasing with the sum of the arrows. This finding indicates that a close approximation to the true LLR can be obtained by simply adding or subtracting a fixed quantity from a running total for each incoming arrow, and justifies our use of this sum as a proxy for evidence in the following analyses (Fig. 2C; compare with Fig. 1B).

In summary, the theoretical analysis suggests that when perceiving the arrows optimally and performing as perfect “noiseless” logicians ($p = 1$), subjects should cease updating their internal state for arrows arriving late in the sequence, whenever the earlier arrows convey enough certainty. Mathematical analysis, however, shows that a much simpler strategy of simply adding all incoming arrows can be optimal too, particularly under conditions of faulty perception.

Behavioral results

Reaction times (RTs) decreased monotonically with longer sequence length (Fig. 3A) ($T_{(1,13)} = 10.8$; $p < 0.001$), likely because of the increasing probability of an impending go cue as a function of sequence length. Moreover, subjects responded faster when they had accumulated more evidence (Fig. 3A) ($T_{(1,13)} = 4.13$; $p = 0.0012$). Error rates were generally low (mean \pm SD, $6.0 \pm 2.3\%$), indicating that the subjects were well able to perform the task. As can be seen in Figure 3B, the probability of responding left or right was modulated by the amount and direction of the evidence accumulated ($T_{(1,13)} = 72.3$; $p < 0.001$). There was also an effect of arrow position on reaction time, suggesting that later information was weighed more heavily than earlier information (Fig. 3C). When the evidence conveyed by the fifth (last) stimulus was congruent with the final decision, this led to a 51 ms RT gain ($T_{(1,13)} = 8.82$; $p < 0.001$). Congruence of the fourth stimulus also led to a RT gain, but of smaller magnitude (17 ms; $T_{(1,13)} = 3.22$; $p = 0.0067$), whereas there was no significant influence of congruence of the first three arrows (all $p > 0.5$). *Post hoc* tests show that the congruence effect of the last arrow was observed when previous evidence was already sufficient for forming a decision (evidence = 2: $T_{(1,13)} = 9.86$, $p <$

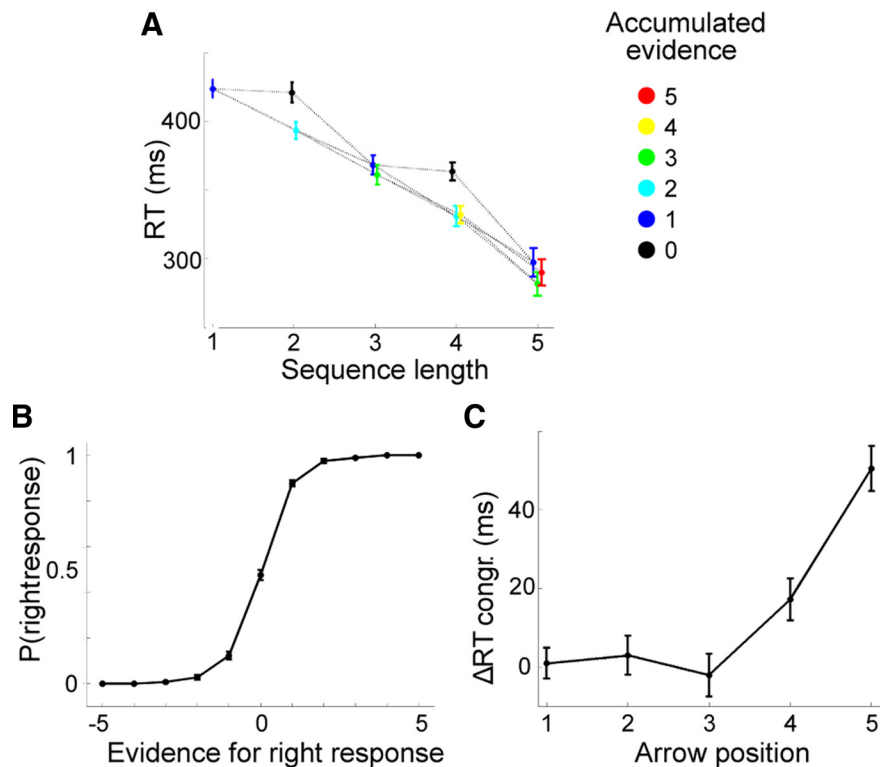


Figure 3. Behavioral results suggesting that subjects accumulate incoming evidence. **A**, RTs as a function of accumulated evidence and sequence length. **B**, Subject’s responses as a function of accumulated evidence for a rightward response. **C**, Influence of the position of the arrow in the sequence on RT. The later arrows (fourth and fifth in the sequence) showed a larger contribution to subject’s RT than arrows presented early in the sequence. Error bars indicate SEM.

0.001; evidence = 4: $T_{(1,13)} = 2.32$, $p = 0.037$). This behavior does not correspond to purely logical decision making, since new information continued to influence behavioral performance after the decision boundary has been crossed. The data further showed a significantly larger influence of arrow congruence when previous evidence was lower (evidence = 2 vs evidence = 4: $T_{(1,13)} = 3.40$, $p = 0.0048$). Therefore, the results are also not consistent with a simple summation model, since new information was weighted differently for different levels of previous evidence. The intermediate Bayesian model did capture this behavior rather well. This model showed no differential change in LLR for different levels of evidence at either very low or very high levels of sensory noise, but a maximal difference in information conveyed by the incongruent stimulus between different levels of evidence at an intermediate reliability of $p = 0.84$ (see supplemental Fig. S1, available at www.jneurosci.org as supplemental material).

Parietal and central activity are inversely related to accumulated evidence

We first investigated whether the amount of previously accumulated evidence (evidence) influences neural activity levels at the onset of processing new stimuli. For this, we compared activity at stimulus onset (0–50 ms after stimulus presentation) as a function of the amount of previously accumulated evidence. Figure 4A shows sensors that showed a significant modulation of neural activity at stimulus onset as a function of evidence. As can be seen from Figure 4B, an inverse relationship is observed: when more evidence has been accumulated, neural activity before the upcoming stimulus is lower. This effect is significant over central (cluster size = 18; cluster $T =$

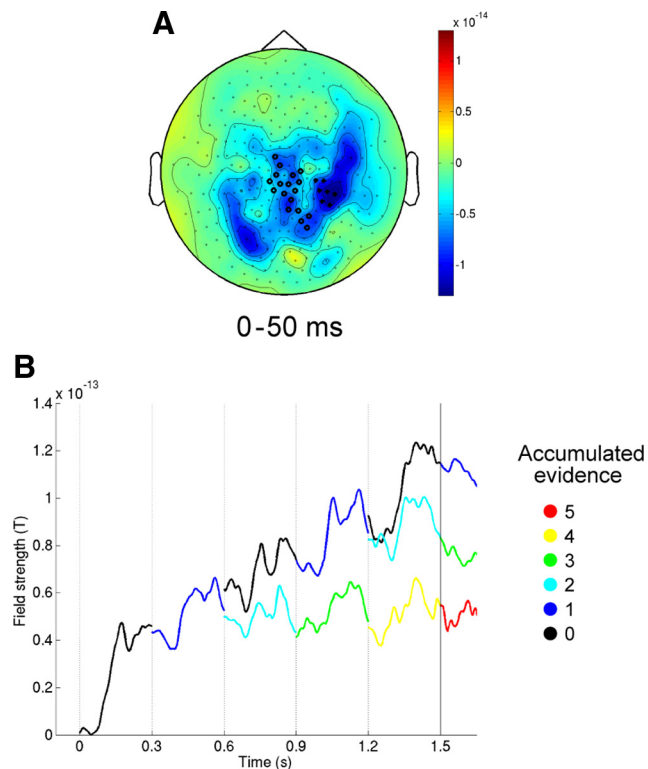


Figure 4. Amount of previous evidence determines neural activity at stimulus onset. **A**, Grand average of the topography of the influence of previously accumulated evidence (evidence) on brain activity at the onset of the incoming stimulus (0–50 ms). The markers indicate clusters of significant regression weights ($p < 0.05$, corrected for multiple comparisons). The right parietal cluster is denoted by “+,” and the central cluster is denoted by “o.” **B**, Time course of sensors showing a significant influence of evidence on activity at stimulus onset. For illustration purposes, as well as to be able to dissociate between the effects of evidence and change (see below), we only plot trials in which there was no stimulus change. It can be seen that both at the onset of the third stimulus (600 ms), the fourth stimulus (900 ms) and the fifth stimulus (1200 ms) there is overall lower activity when more evidence has been accumulated.

62.6; $p_{\text{corrected}} = 0.0005$) and right parietal sensors (cluster size = 7; cluster $T = 27.6$; $p_{\text{corrected}} = 0.0075$).

We then investigated whether evidence also influenced the transient activity induced by the upcoming stimulus (over and above the baseline shift). For this, we subtracted from each trace the baseline at the onset of the stimulus (0–50 ms after stimulus presentation). As can be seen from Figure 5, the increase in activity induced by each stimulus was lower when more evidence has been accumulated. This effect was significant in central sensors, both in the interval 150–225 ms (cluster size = 5; cluster $T = 16.9$; $p_{\text{corrected}} = 0.034$) and 225–300 ms (cluster size = 6; cluster $T = 24.4$; $p_{\text{corrected}} = 0.017$).

Similarly to the behavioral analysis, we also compared the effect of congruency of the last arrow between different levels of evidence (evidence = 2 vs evidence = 4). When the last arrow was congruent with previous evidence, this resulted in a trend of larger frontal activity when evidence was lower (evidence = 2 vs evidence = 4; cluster size = 3; cluster $T = 10.7$; $p_{\text{corrected}} = 0.079$), in the interval 225–300 ms after arrow onset. This frontal cluster had a similar topography as the sensors showing a main effect of time (supplemental Fig. S2A, available at www.jneurosci.org as supplemental material) (see also Fig. 7A) (see below). Incongruency of the last arrow resulted in a significantly larger right parietal activity when evidence was higher in the interval 225–300 ms after arrow onset (evidence = 4 vs evidence = 2; cluster size = 9;

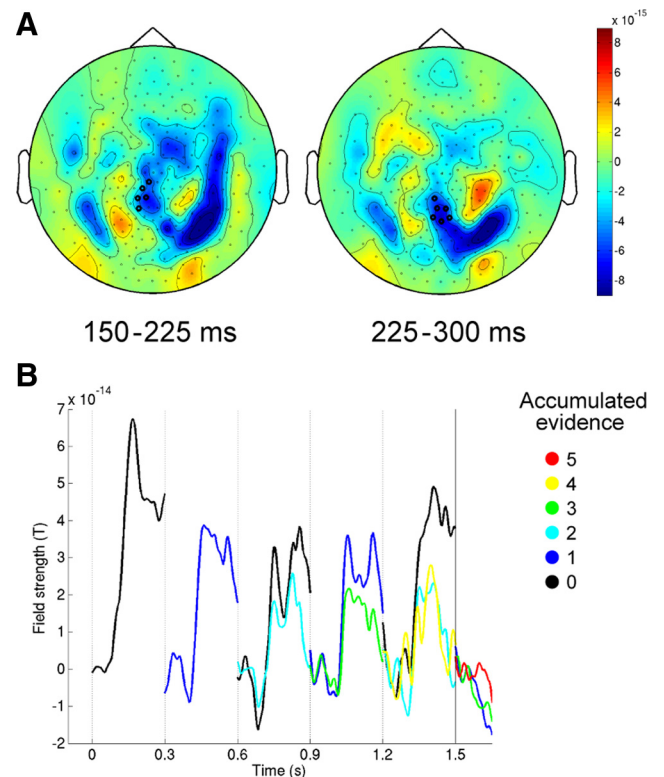


Figure 5. Amount of accumulated evidence changes processing of new stimuli. **A**, Grand average of the topography of the influence of accumulated evidence (evidence) on brain activity changes induced by the new stimulus, at two time intervals (150–225 and 225–300 ms after stimulus onset). The dots indicate clusters of significant regression weights ($p < 0.05$, corrected for multiple comparisons). **B**, Time course of the sensors showing a significant influence of evidence on activity changes. More evidence leads to attenuated activity changes for the third stimulus (750–900 ms), the fourth stimulus (1050–1200 ms), and the fifth stimulus (1350–1500 ms). Other conventions are as in Figure 4.

cluster $T = 30.9$; $p_{\text{corrected}} = 0.0055$). This parietal cluster had a similar topography as the sensors showing the main effect of evidence (supplemental Fig. S2B, available at www.jneurosci.org as supplemental material) (Fig. 5A). This pattern is in good correspondence with the behavioral results, insofar as it shows that (1) new information continued to influence accumulation-related neural activity after the decision boundary had been crossed and (2) there was a differential influence of information under different levels of evidence.

Occipital and fronto-central regions react to a change in stimulus direction

Next, we looked at neural activity changes as a function of whether the stimulus conveyed a change in the direction of evidence (change) with respect to the preceding stimulus—a behaviorally relevant factor that slowed response times. As can be seen from Figure 6, a change in the direction of evidence resulted in a large amplification of neural activity. This amplification was visible after 150–225 ms in bilateral occipital sensors (cluster size = 63; cluster $T = 236.8$; $p_{\text{corrected}} < 0.0001$) and was then propagated to central and frontal sensors, where it led to a significant difference after 225–300 ms after stimulus onset (cluster size = 140; cluster $T = 587.4$; $p_{\text{corrected}} < 0.0001$). It is noteworthy that the larger response for change happens at all levels of evidence accumulated. Indeed, well after the decision boundary has been crossed, a change in evidence still leads to a sizeable increase in neural activity, in both occipital and central-frontal sensors.

Frontal signals are enhanced and posterior signals are attenuated as a function of time

Finally, we looked at neural activity as a function of time (time). In general, virtually all sensors showed a significant ramping up of activity at stimulus onset as a function of time (cluster size = 273; cluster $T = 1602$; $p_{\text{corrected}} < 0.0001$). Looking at the influence of time on changes in activity induced by the stimulus (over and above baseline shifts), we found that occipital signals were suppressed over time, whereas frontal signals became stronger (Fig. 7). This led to a significant negative regression over posterior sensors 150–225 ms after stimulus onset (cluster size = 10; cluster $T = 36.1$; $p_{\text{corrected}} = 0.0045$) and a significant positive regression over frontal sensors both 150–225 ms after stimulus onset (cluster size = 13; cluster $T = 47.5$; $p_{\text{corrected}} = 0.0005$) and 225–300 ms after stimulus onset (cluster size = 6; cluster $T = 22.7$; $p_{\text{corrected}} = 0.024$).

Discussion

In this study, we have investigated how human observers accumulate evidence for one of two choices as the decision process is gradually constrained by serially incoming sensory information. We contrasted three possibilities for evidence accumulation: mental logic, Bayesian integration, and addition/subtraction. Using a simple mathematical model, we found that these three types of behavior could arise from an optimal decision process under different levels of noise. Behaviorally, response times and error rates supported a nonlinear but continuous accumulation process, extending beyond the point at which logical reasoning sufficed to reach a decision. MEG recordings showed a correlate of evidence accumulation over parietal and central cortex that was inversely related to the amount of accumulated evidence. This suggests these areas upregulate their response under situations of increasing uncertainty. We also observed larger activity over occipital and fronto-central cortex when the stimulus conflicted with the directly preceding stimulus—again, regardless of whether that change was logically irrelevant or not. This activity increase was regardless of the amount of evidence accumulated. Finally, as the decision process unfolded over time, frontal cortex increased its activity while occipito-parietal activity was suppressed. Together, the results highlight top-down and bottom-up components of evidence accumulation during sequential decision making, which we will discuss below.

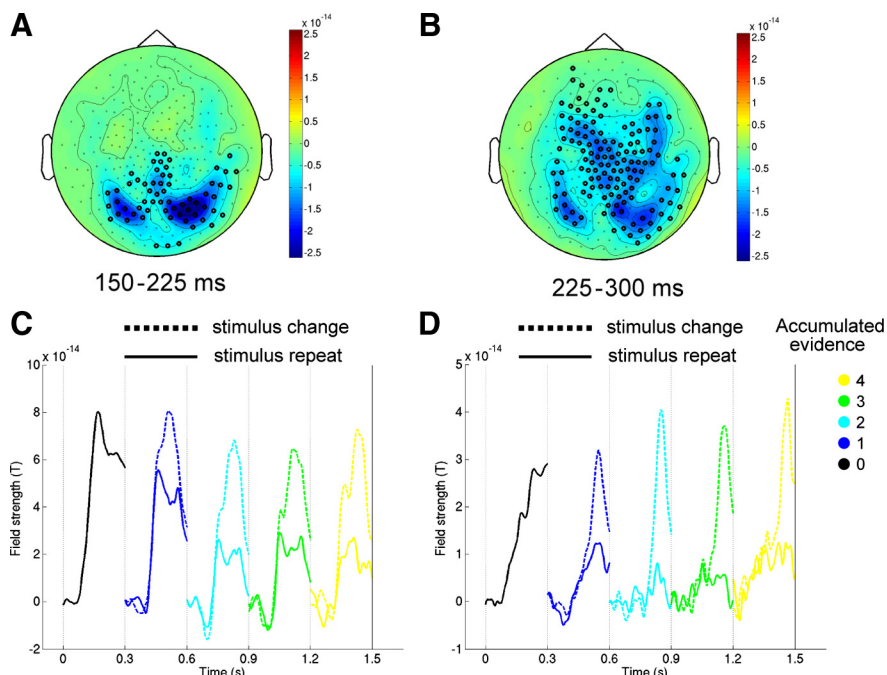


Figure 6. Change in evidence leads to increased activity. **A, B**, Grand average of the topography of the influence of a change in arrow direction (change) on brain activity change, at two time intervals [150–225 ms (**A**) and 225–300 ms (**B**) after stimulus onset]. The dots indicate clusters of significant regression weights ($p < 0.05$, corrected for multiple comparisons). **C, D**, Time course of posterior sensors showing a significant influence of change after 150–225 ms (**C**), and of frontal sensors showing a significant influence of change after 225–300 ms (**D**). For illustration purposes, as well as to be able to dissociate between evidence and change (see below), we present one level of evidence for each time window.

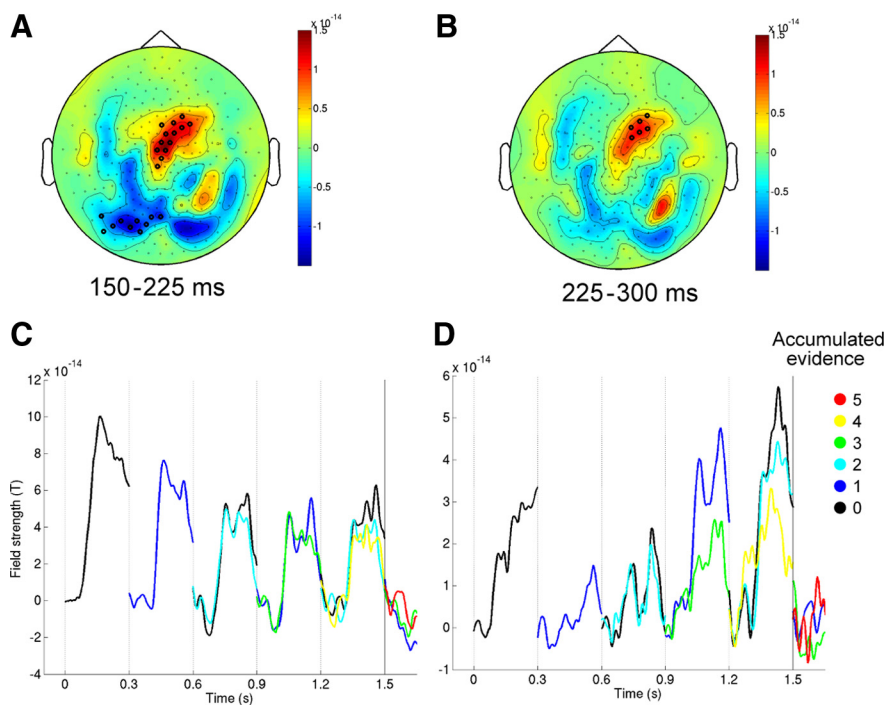


Figure 7. Occipital activity attenuates and frontal activity increases over time. **A, B**, Grand average of the topography of the influence of time on changes in brain activity induced by the stimulus, at two time intervals [150–225 ms (**A**) and 225–300 ms (**B**) after stimulus onset]. The dots indicate clusters of significant regression weights ($p < 0.05$, corrected for multiple comparisons). **C**, Time course of occipital sensors showing a significant negative relationship with time after 150–225 ms. **D**, Time course of frontal sensors showing a significant positive relationship with time both after 150–225 ms and after 225–300 ms. Other conventions are as in Figure 4.

Parietal and central signals are inversely related to nonlinear evidence accumulation

There was a gradual buildup of activity with each incoming stimulus, consistent with the concept of accumulation of evidence. However, both the amount of activity (Fig. 4) and change in activity for each new stimulus (Fig. 5) were inversely related to the amount of previous evidence accumulated. In other words, when a larger amount of evidence had already been accumulated, new sensory evidence had a lower impact on brain activity in parietal and central areas. This indirect signature of evidence accumulation is quite distinct from the bottom–up sensory evidence accumulation that shows a positive relationship between activation strength and amount of evidence accumulated and has been previously observed using single-cell recordings in macaques (Platt and Glimcher, 1999; Shadlen and Newsome, 2001; Yang and Shadlen, 2007). It is also opposite to what would be expected from the mathematical model, which shows larger changes in LLR for arrows that are presented later in the sequence. We suggest that this new observation of a negative modulation of brain activity with evidence likely reflects a top–down influence on sensory processing, effectively constraining the influence of sensory information on the decision variable. It is well known that internal state variables like expectation, bias, and reward can all help to reduce the computational load imposed by the information in the outside world (Gilbert and Sigman, 2007; Nienborg and Cumming, 2009; Rushworth et al., 2009). These internal state variables can subsequently impact visual and attentional processing (Murray et al., 2002; Summerfield and Koechlin, 2008). As such, our data are consistent with the hypothesis that subjects formed expectations about the outcome of each trial based on the basis of initial evidence and used this information to modulate the amount of resources dedicated to the stimuli that were subsequently presented. Although the current set of data was not optimally suited to quantitatively link the mathematical model to the observed MEG responses, this limitation could be potentially overcome by using a larger set of shapes that convey probabilistic information about trial outcome (cf. Yang and Shadlen, 2007). It would be interesting to see whether and how previous expectations change the monotonic relationship between LLR and neural activity that has been observed in these settings.

Changing the evidence leads to amplification of occipital and fronto-central cortex

When the direction of a stimulus was opposite to the direction of the preceding stimulus, this resulted in a large increase in activity that started in occipital cortex and spread to central and frontal regions. Interestingly, this activity difference was equally present for all levels of accumulated evidence. Although it is possible that the larger response to a different stimulus can be partly explained as a sensory phenomenon [e.g., repetition suppression (Grill-Spector et al., 2006)], the large spread of activity to central and frontal areas suggests that it is not merely a sensory effect. It is possible that the larger activity for a change in stimulus direction is caused by a form of task switching, either because of the change from addition of evidence to subtraction of evidence, or equivalently, to the change from engaging one accumulator to the other. Future studies that manipulate changes in accumulation independently from stimulus identity are required to dissociate between these possibilities.

Time matters during sequential decision making

There were opposite effects of the passage of time on occipito-parietal and frontal sensors (Fig. 7). Although occipito-parietal

responses to stimuli were attenuated over time, frontal signals showed increasing responses to new stimuli, suggesting that information that was acquired at a later moment in time was weighted more heavily in these regions. There was an interesting behavioral counterpart to this phenomenon: congruency of the first three stimuli with the final decision had no effect on reaction times, whereas there was a moderate effect (17 ms) for the fourth stimulus and a large effect (51 ms) for the fifth stimulus (Fig. 3C). The sensory attenuation in posterior cortex and enhancement of medial frontal cortex may therefore be the neural counterpart and explanation for this behaviorally observed “recency” effect.

Mental logic or accumulation?

The finding that a change in evidence led to a large increase in activity, in both sensory and decision-related areas, and regardless of the amount of evidence accumulated (Fig. 6) is clearly not in line with a “mental logic” model, in which conflicting evidence should have no influence after the decision threshold has been crossed. Nevertheless, it is also obvious from the data that the amount of previously accumulated evidence had a large impact on processing of new stimuli (Fig. 5). This argues against a simple addition model: if new evidence is simply added or subtracted, the accumulation process should not interact with its history. Our data indicate that the impact of congruent or incongruent information on brain activity and behavior depends on the amount of already accrued information. The fact that new stimuli lead to a larger response under conditions of increasing uncertainty suggests that the brain weighs the attentional resources according to the certainty it has about the outcome. Together, our results offer no support for the mental logic model, and generally support the view that incoming evidence is continuously accumulated, but with two violations from a purely additive model: new information is differentially weighted depending on the previously accumulated evidence, whereas recent conflicting information can affect the decision process even after the decision boundary should have been crossed.

Our results could be interpreted as suggesting that human subjects are inherently irrational and let their decisions be influenced by irrelevant variables. However, our paradigm, with one arrow every 300 ms, also leaves open the possibility that subjects were simply unable to deploy their logical reasoning abilities at this high speed, and therefore resorted to a faster but less precise accumulation system. We deem it likely that, if the current task was slowed, additional logical inference processes would become deployed, as they were indeed observed in a related but slower sequential logic task studied with functional magnetic resonance imaging (Landmann et al., 2007).

In conclusion, we demonstrated that sequential decision making in humans proceeds by accumulation of serially incoming evidence, as originally reported in monkeys (Yang and Shadlen, 2007). However, our indirect noninvasive MEG recordings were unable to pick up the increasing neural signals that directly reflect this evidence accumulation process. Instead, our results showed a new and unexpected effect: the negative impact of previous evidence accumulation on brain activation level and on incoming sensory-evoked activation. These processes complement the well described forward flow of neural information during decision making and indicate that, when making decisions on the basis of sequential information, the human nervous system uses the amount of previously accumulated information to constrain the accumulation of additional evidence.

References

- Bastiaansen MC, Knösche TR (2000) Tangential derivative mapping of axial MEG applied to event-related desynchronization research. *Clin Neurophysiol* 111:1300–1305.
- Bauer M, Oostenveld R, Peeters M, Fries P (2006) Tactile spatial attention enhances gamma-band activity in somatosensory cortex and reduces low-frequency activity in parieto-occipital areas. *J Neurosci* 26:490–501.
- Bell AJ, Sejnowski TJ (1995) An information-maximization approach to blind separation and blind deconvolution. *Neural Comput* 7:1129–1159.
- Bullier J, Nowak LG (1995) Parallel versus serial processing: new vistas on the distributed organization of the visual system. *Curr Opin Neurobiol* 5:497–503.
- Cisek P, Kalaska JF (2005) Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45:801–814.
- de Lange FP, Jensen O, Bauer M, Toni I (2008) Interactions between posterior gamma and frontal alpha/beta oscillations during imagined actions. *Front Hum Neurosci* 2:7.
- Gilbert CD, Sigman M (2007) Brain states: top-down influences in sensory processing. *Neuron* 54:677–696.
- Gold JL, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–394.
- Gold JL, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Grill-Spector K, Henson R, Martin A (2006) Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci* 10:14–23.
- Hamalainen M, Hari R, Ilmoniemi R, Knuutila J, Lounasmaa O (1993) Magnetoencephalography. Theory, instrumentation and applications to the noninvasive study of brain function. *Rev Mod Phys* 65:413–497.
- Heekeren HR, Marrett S, Bandettini PA, Ungerleider LG (2004) A general mechanism for perceptual decision-making in the human brain. *Nature* 431:859–862.
- Heekeren HR, Marrett S, Ruff DA, Bandettini PA, Ungerleider LG (2006) Involvement of human left dorsolateral prefrontal cortex in perceptual decision making is independent of response modality. *Proc Natl Acad Sci U S A* 103:10023–10028.
- Jung TP, Makeig S, Westerfield M, Townsend J, Courchesne E, Sejnowski TJ (2000) Removal of eye activity artifacts from visual event-related potentials in normal and clinical subjects. *Clin Neurophysiol* 111:1745–1758.
- Kaiser J, Lennert T, Lutzenberger W (2007) Dynamics of oscillatory activity during auditory decision making. *Cereb Cortex* 17:2258–2267.
- Landmann C, Dehaene S, Pappata S, Jobert A, Bottlaender M, Roumenov D, Le Bihan D (2007) Dynamics of prefrontal and cingulate activity during a reward-based logical deduction task. *Cereb Cortex* 17:749–759.
- Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG- and MEG-data. *J Neurosci Methods* 164:177–190.
- Murray SO, Kersten D, Olshausen BA, Schrater P, Woods DL (2002) Shape perception reduces activity in human primary visual cortex. *Proc Natl Acad Sci U S A* 99:15164–15169.
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp* 15:1–25.
- Nienborg H, Cumming BG (2009) Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459:89–92.
- Nowak LG, Munk MH, Girard P, Bullier J (1995) Visual latencies in areas V1 and V2 of the macaque monkey. *Vis Neurosci* 12:371–384.
- Ospova D, Takashima A, Oostenveld R, Fernández G, Maris E, Jensen O (2006) Theta and gamma oscillations predict encoding and retrieval of declarative memory. *J Neurosci* 26:7523–7531.
- Philiastides MG, Sajda P (2006) Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cereb Cortex* 16:509–518.
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238.
- Rushworth MF, Mars RB, Summerfield C (2009) General mechanisms for making decisions? *Curr Opin Neurobiol* 19:75–83.
- Schall JD (2001) Neural basis of deciding, choosing and acting. *Nat Rev Neurosci* 2:33–42.
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916–1936.
- Summerfield C, Koehlin E (2008) A neural representation of prior information during perceptual inference. *Neuron* 59:336–347.
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080.