



Computational Neuroscience

Quantifying the time for accurate EEG decoding of single value-based decisions



Athina Tzovara ^{a,b,*}, Ricardo Chavarriaga ^c, Marzia De Lucia ^{a,b}

^a *Electroencephalography Brain Mapping Core, Center for Biomedical Imaging (CIBM), Department of Radiology, University Hospital Center, University of Lausanne, 1011 Lausanne, Switzerland*

^b *Laboratoire de recherche en Neuroimagerie – LREN, Department of Clinical Neuroscience, Lausanne University and University Hospital, 1011 Lausanne, Switzerland*

^c *Chair in Non-Invasive Brain-Computer Interface, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

HIGHLIGHTS

- We outline a novel method for assessing time-unlocked topographic EEG activity.
- We apply this method for accurately predicting accept vs. reject decisions.
- The time when decisions have been made is modulated by task difficulty.
- It occurs ~340 ms before subjects' button press.
- This time is compatible with a diffusion process, evaluated at behavioral level.

ARTICLE INFO

Article history:

Received 18 July 2014

Received in revised form

25 September 2014

Accepted 25 September 2014

Available online 5 October 2014

Keywords:

Decision-making

Accumulation

Decoding

Single-trial

EEG

Drift diffusion model

ABSTRACT

BACKGROUND: Recent neuroimaging studies suggest that value-based decision-making may rely on mechanisms of evidence accumulation. However no studies have explicitly investigated the time when single decisions are taken based on such an accumulation process.

NEW METHOD: Here, we outline a novel electroencephalography (EEG) decoding technique which is based on accumulating the probability of appearance of prototypical voltage topographies and can be used for predicting subjects' decisions. We use this approach for studying the time-course of single decisions, during a task where subjects were asked to compare reward vs. loss points for accepting or rejecting offers.

RESULTS: We show that based on this new method, we can accurately decode decisions for the majority of the subjects. The typical time-period for accurate decoding was modulated by task difficulty on a trial-by-trial basis. Typical latencies of when decisions are made were detected at ~500 ms for 'easy' vs. ~700 ms for 'hard' decisions, well before subjects' response (~340 ms). Importantly, this decision time correlated with the drift rates of a diffusion model, evaluated independently at the behavioral level.

COMPARISON WITH EXISTING METHOD(S): We compare the performance of our algorithm with logistic regression and support vector machine and show that we obtain significant results for a higher number of subjects than with these two approaches. We also carry out analyses at the average event-related potential level, for comparison with previous studies on decision-making.

Conclusions: We present a novel approach for studying the timing of value-based decision-making, by accumulating patterns of topographic EEG activity at single-trial level.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Before making a decision, we often need to extract information from the environment and weight within possible alternative options in order to optimize a potential outcome. Perceptual decisions are known to rely on mechanisms of evidence accumulation (Gold and Shadlen, 2007; Kim and Shadlen, 1999; Shadlen and Newsome, 2001) described by drift diffusion models (DDM;

* Corresponding author at: Centre Hospitalier Universitaire Vaudois (CHUV), Mont Pâlis 16, 1011 Lausanne, Switzerland. Tel.: +41(0)21 314 6828.

E-mail address: Athina.Tzovara@uzh.ch (A. Tzovara).

(Ratcliff, 1978)). According to those models, sensory evidence, usually deriving from competing sources or masked stimuli (Heekeren et al., 2004), is integrated over time from a starting point toward one of two decision boundaries, corresponding to alternative choices. Electroencephalography (EEG) experiments have identified neurophysiological signals in humans that reflect this accumulation, either as fluctuations in oscillatory power in specific frequency bands (van Vugt et al., 2012; Wyart et al., 2012), or as the EEG activity on pre-defined groups of electrodes (O'Connell et al., 2012).

Recent studies suggest that similar mechanisms of evidence accumulation might also be involved in value-based decision making, where values corresponding to an offer are encoded, compared and accumulated over time according to DDM (Basten et al., 2010). A widespread network of brain regions is involved in this process, with the ventromedial prefrontal cortex being a value integration region (Basten et al., 2010; Philiastides et al., 2010), and parietal regions accumulating results of this integration (Basten et al., 2010; Gluth et al., 2012; Hunt et al., 2012). Because an accumulation mechanism implies a dynamic process which evolves over time, assessing the temporal dynamics of decision-making can be highly informative of the underlying neural process.

Previous studies based on EEG or magnetoencephalography (MEG) have typically investigated neural activity at an across-trials level (Cohen and Ranganath, 2007; Frank et al., 2005; Talmi et al., 2012; Tzovara et al., 2012a; Yeung et al., 2005) and provided some insight on the timing of decision-making. These studies have shown that the neural correlates of outcome evaluation typically appear as modulations of EEG components, as early as ~200–300 ms after the presentation of the outcome of subjects' decisions (Cohen and Ranganath, 2007; Yeung et al., 2005). However, these findings often rely on average event-related analyses and therefore overlook variability in decision-making strategies across subjects or trials. Single-trial analyses can provide further insight into various stages of decision-making and possibly reveal mechanisms that vary on a trial-to-trial level and would remain undetected otherwise (Pernet et al., 2011).

The utility of studying single-trial activity has been highlighted by recent studies in value-based decision-making (Billeke et al., 2013; Gluth et al., 2013; Hunt et al., 2012; Philiastides et al., 2010). In particular, these studies have isolated readiness potentials in the EEG signal which reflect the emergence of value-based decisions (Gluth et al., 2013) or, in the case of MEG, have allowed tracking the information flow in brain regions linked to different stages of a decision (Hunt et al., 2012). Even though these studies clearly demonstrate the added value of single-trial analyses, they mainly rely on model-based approaches, where one needs to make strong assumptions about the neural mechanisms underlying decisions. Additionally, analyses often rely on a pre-selection of time-windows and/or electrodes of interest, which might introduce a bias in the results, or overlook phenomena occurring outside the region/time-window of interest. In the present study, we present a novel data-driven approach for studying decision-making. This approach builds upon an existing single-trial EEG method, based on modeling the distribution of voltage topographies with a mixtures of Gaussians model (GMM) (Tzovara et al., 2012b) and takes advantage of the whole electrode montage for identifying single-trial patterns of EEG activity.

We carried out an experiment where subjects made decisions by combining values (reward vs. loss points). A previous study has shown that this task follows behaviorally a DDM (Basten et al., 2010). Here, we assumed the same behavioral model and aimed at extracting neural activity preceding single decisions. We hypothesized that because of a high degree of inter-trial variability, the decision-making process would likely be revealed by single-trial instead of average analyses.

The method we introduce here is based on EEG activity starting from the offer presentation and lasting until the button press on a single-trial basis. This method can therefore deal with trials of different lengths. A simpler version of this approach has been previously used for decoding decisions in an exploration/exploitation task based on a fixed-length trial (Tzovara et al., 2012a). The present study aims also at taking those previous findings one step further by detecting the time when single decisions are made at a neural level and by coupling behavioral models' parameters with single-trial neural modulations. Information about the timing of value-based decision-making will be beneficial for a better understanding of its mechanisms in the human brain and possibly for establishing a link between perceptual and value-based decisions.

2. Materials and methods

2.1. Participants

Twelve healthy volunteers participated in this study (6 females), with mean age of 27 ± 1 years old (the values here and in the following refer to the mean \pm standard error – s.e.m.), all having normal or corrected-to-normal vision. None of the subjects reported a history of neurological disease.

2.2. Procedure and task

The experimental protocol was adapted from a previous functional magnetic resonance imaging (fMRI) study (Basten et al., 2010). It consisted of two phases: the training phase and the main experiment. Before the training phase subjects were presented with lists of 8 colors and 8 shapes, corresponding to rewards and losses respectively and were asked to memorize them. Each color was thus associated with a range of points that subjects could win (reward points) and each shape with a range of points that subjects could lose (loss points). These associations were random and unique for every subject in order to exclude the influence of low-level visual features in task efficiency. Colors corresponded to reward points and shapes to loss points for half of the participants and vice versa for the other half. Reward and loss points corresponded to benefits and costs in a previous study using a similar paradigm (Basten et al., 2010).

When subjects were confident that they had memorized these associations, they underwent two training sessions for learning rewards and losses, consisting of three blocks of 57 trials each (171 trials in total). In every trial they were presented with a pair of different colors (i.e. a 'wavy' vertical bar of different colors) or different shapes and were asked to indicate with a button press which one was more advantageous, corresponding to a higher reward or a smaller loss. They had 2 s to respond; after that they received visual feedback for their choice and saw the actual values of the two stimuli. The feedback was displayed on the screen for 1 s and was followed by a fixation cross, lasting 1.00 ± 0.12 s. The presented stimuli had values randomly and equiprobably drawn from the interval that subjects had initially seen, spanning from 1 up to 80 points, with a step of 10. For example, the stimulus corresponding to the lowest reward interval could give 1 up to 10 points and the one corresponding to the greatest loss –80 up to –71 loss points.

After completing the two training sessions, subjects were faced with the main experiment which was a gambling task (Fig. 1a). Subjects were presented with offers, consisting of combinations of the previously learned stimuli (i.e. shapes filled with color) and their task was to decide whether they would accept or reject them, through a button press with their index or middle finger. This decision was based on whether the overall value of the offer, when subtracting the reward from the loss points would lead to a win or

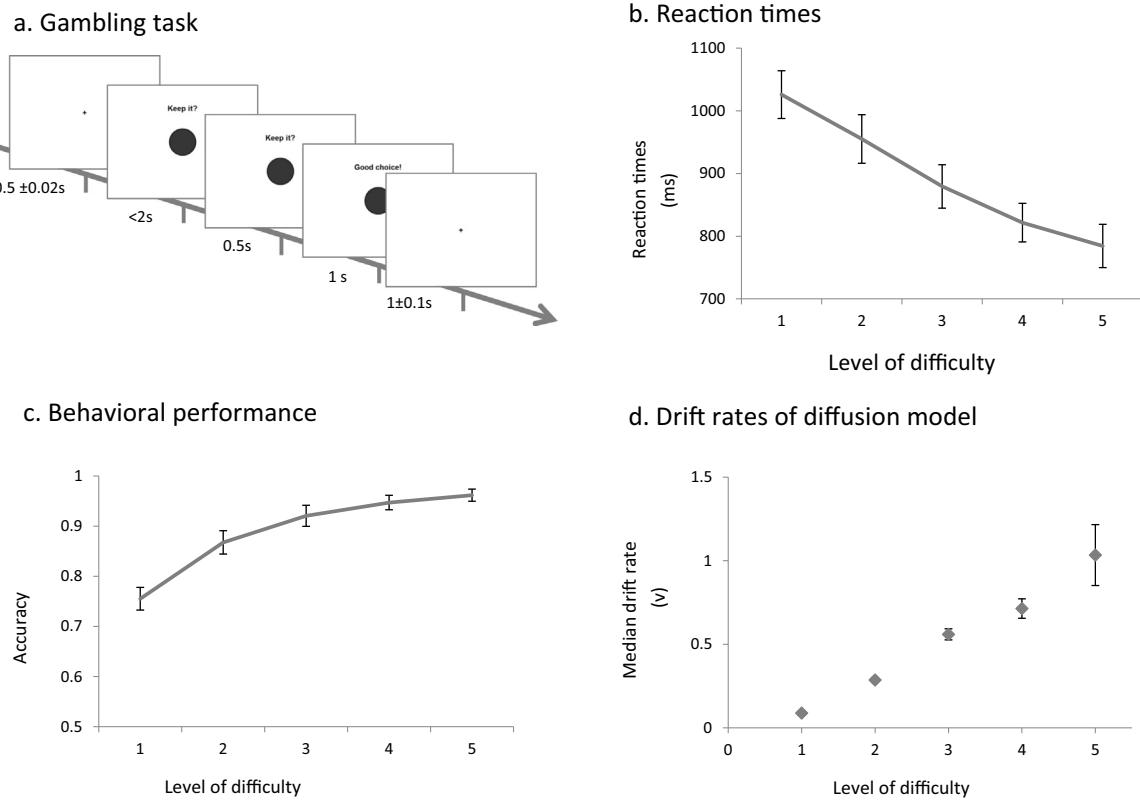


Fig. 1. Description of the experiment and behavioral results. (a) Each trial started with a fixation cross (0.50 ± 0.02 s) which was followed by the offer presentation, consisting of a combination of the previously learned colors and shapes. Subjects had 2 s maximum to respond through a button press whether they would accept or reject the offer, based on whether it consisted of more reward or loss points. The offer remained on the screen for 0.5 s after the button press and was then followed by a feedback ("Good choice"/"Bad choice"). (b) reaction times across difficulty levels (1–5). Level of difficulty is measured by the absolute difference in the reward vs. loss points that are contained within each offer. Level 1 corresponds to the highest difficulty while level 5 to the easiest decisions. Subjects were faster in responding as difficulty decreased. (c) Similarly, the ratio of correct responses consistently increased as the level of difficulty decreased. (d) Median drift rates, across subjects, as they were estimated by the drift diffusion model. Drift rates increase as the level of difficulty decreases, leading to a faster accumulation process and therefore decisions.

loss. For each trial subjects saw a fixation cross for 0.5 ± 0.02 s which was followed by the display of the offer (Fig. 1a). Subjects had 2 s to make a decision, starting immediately after the offer display. Half a second after they made a decision by pressing a button, they saw a message informing them if their choice was advantageous (Fig. 1a) for 1 s, which was followed by a fixation cross that remained on the screen for 1.0 ± 0.1 s. In order to simulate a realistic situation and to make the experiment more motivating, subjects were told that they had a total of 500 points to begin with and that their decisions would influence this amount: if they correctly accepted an offer corresponding to a total positive gain, this gain would be added to their overall points, however, if they erroneously accepted an offer corresponding to a loss, this would be subtracted from their total points.

The experimental protocol used in this study allows varying the levels of decision difficulty, depending on the difference of reward minus loss points for any given offer. A previous fMRI study using this experiment (Basten et al., 2010) showed that behaviorally, the greatest the distance between the absolute values of reward and loss for a given offer, the easiest it is for subjects to make a decision, as measured by reduced reaction times and increased accuracy. In this sense, we used the difference between reward and loss points' range as an index of decision difficulty, which ranged from 1 for hard decisions to 5 for the easiest ones.

In total, we presented subjects with 8 blocks, including all possible combinations of stimuli with few exceptions. We did not present any of those combinations where the reward and loss points would be equal, and we did not present most of the combinations that contained extreme values (i.e. offers containing stimuli

that correspond to the maximum or minimum of reward or loss points). The reason for excluding these extreme values was that subjects would likely not need to compute the values of both stimuli but would instead base their decisions on the extreme value alone. Each block contained every color/shape combination twice, with randomized order of presentation. This resulted in 76 trials per block, or 608 trials in total per subject. The whole experiment, including the training session and breaks between blocks, lasted approximately 1.5 h per participant.

2.3. Behavioral models

We modeled the subjects' behavior according to a DDM (Ratcliff, 1978). The suitability of this model to explain this type of decisions has been shown in a previous study (Basten et al., 2010). Here, we assumed the same type of behavioral model and computed its parameters in a similar way as in (Basten et al., 2010). Specifically, this previous study has shown that drift rates vary according to the level of difficulty, while the rest of the DDM's parameters did not differ (distance between boundaries, the starting point and the total non-decision time). Here, we employed a similar approach and computed separate DDMs for each subject, with reaction times and type of responses being the dependent variables of the models. For computing the DDMs we used the DMA toolbox (VandeKerckhove and Tuerlinckx, 2007).

We also computed additional behavioral models, for which we kept all the parameters fixed across levels of difficulty and we compared them to the models where the drift rates varied across difficulty levels. This comparison was based on the Bayesian

Information Criterion (BIC; (VandeKerckhove and Tuerlinckx, 2007). For every subject, the model where the drift rates vary across difficulty levels provided lower values of BIC than the model where parameters were considered constant across difficulty. Indeed, in these two cases we obtained a range of BIC values of 969–1653 for the models with fixed parameters, while the more general one provided a range of values of 766–1507. In the main results we thus always consider the model with varying drift rates across different levels of difficulty.

2.4. EEG acquisition and preprocessing

Continuous 128-channel EEG was acquired through a Biosemi Active II portable system with a sampling rate of 1024 Hz. EEG was referenced to the common mode sense (CMS; active electrode) and grounded to the driven right leg (DRL; passive electrode), which functions as a feedback loop driving the average potential across the electrode montage to the amplifier zero. EEG was recorded during the training phase and during the main experiment, however, in the context of this study, we only processed data acquired during the main experiment.

Data preprocessing was done with Cartool (Brunet et al., 2011). EEG data were bandpass filtered offline (0.18–40 Hz) and eye movements, muscle or other artifacts were rejected in a semi-automatic way, by applying a criterion of $\pm 60 \mu\text{V}$. Peristimulus epochs of EEG, spanning from -100 ms before the offer presentation and lasting up to the subjects' button press for each trial, were extracted. Data from artifacted electrodes because of a bad contact were identified visually for each participant (6 ± 1 electrodes per subject) and were interpolated (Perrin et al., 1987). For the single-trial analysis data were down-sampled to 64 electrodes to reduce the computational load and were normalized by their instant EEG power, i.e. the mean of the squared values across electrodes:

$$\mathbf{m} = \frac{\mathbf{m}_0}{\sqrt{\frac{1}{N} \sum_{i=1}^N m_i^2}} \quad (1)$$

Here \mathbf{m} refers to a normalized single-trial voltage topography and consists of a 64-dimensional vector with voltage values across all the N electrodes in the montage: $\mathbf{m} = [\mathbf{m}_1, \dots, \mathbf{m}_N]$ (where $N=64$). \mathbf{m}_0 refers to the original voltage topography.

Only trials where subjects made a 'correct' decision were kept for further analysis. These corresponded either to an accepted offer which would give more reward than loss points, or to a rejected offer which would lead to a loss (i.e. more loss than reward points). The total number of trials that were considered for accept and reject decisions was 227 ± 6 and 221 ± 6 trials per subject, respectively, corresponding to $75 \pm 2\%$ and $73 \pm 2\%$ of all the recorded single-trials. The total number of accept and reject trials did not differ significantly for the two conditions (t -test, $t(11)=1.06$; $p=0.31$).

2.5. EEG analysis

2.5.1. Multivariate EEG decoding based on accumulation of topographic activity

The single-trial topographic EEG analysis was carried out at the single-subject level and was based on a multivariate EEG decoding approach, whose suitability for decoding subjects' decisions has been shown in a previous study (Tzovara et al., 2012a). In this previous study, we evaluated the time of decisions at the neural level by taking advantage of the EEG activity locked to the reward presentation. In the present study, we aimed at detecting such a time at the level of each single decision/trial.

The analysis presented here consists of three main steps: (a) cross-validation: extraction of template maps, (b) validation: decoding decisions based on the extracted maps, consisting of an

evaluation of the generalization of these maps in new trials and (c) detection of critical time-points, which aims at quantifying the timing of single decisions. In the following we present in detail every one of these steps.

2.5.1.1. Cross-validation: extraction of template maps. The employed approach consists of modeling single-trial voltage topographies (i.e. the configuration of electric activity across all electrodes at any given instant) based on GMM models (see similar applications of the GMM for single-trial EEG analysis in (Bernasconi et al., 2011; Cossy et al., 2014; De Lucia et al., 2012)). The probability distribution of a GMM model with Q Gaussians in total was computed in an N -dimensional space ($N=\text{total number of electrodes}$) as:

$$P(\mathbf{m}|\mu, \sigma) = \sum_{k=1}^Q p_k P_k(\mathbf{m}|\mu_k, \sigma_k) \quad (2)$$

P_k denotes the k -th Gaussian distribution with mean μ_k and covariance σ_k . p_k is the prior probability of the k -th Gaussian. In the following, we replace " μ_k, σ_k " with c_k for reasons of simplicity. The parameters of the GMM models were estimated based on an expectation–maximization algorithm (Bishop, 1995; Dempster et al., 1977), which iteratively computes the μ_k, σ_k and p_k parameters by minimizing an error function:

$$E = -\ln L = -\sum_i \ln \left\{ \sum_{k=1}^Q P(\mathbf{m}|c_k)p_k \right\} \quad (3)$$

The index i spans over all the instantaneous single-trial voltage topographies of the training dataset. This step does not take into account the temporal order of appearance of voltage topographies (Tzovara et al., 2012b) and thus extracts few that are representative of the whole dataset (i.e. the means of the Gaussians, or template maps). The extraction of template maps was done for every experimental condition separately (i.e. accept vs. reject decisions).

Model estimation was based on one part of the available trials (training dataset, on average 173 ± 4 trials per subject and type of decision) and was then tested on another part (test dataset, on average 19 ± 4 trials per subject and decision), in a 10-fold cross validation procedure. Final results were confirmed on a validation dataset (26 ± 3 trials per subject and condition), which consisted of trials that were never used at any point in the cross-validation, for training or testing the models (see also 2.5.1.2 Validation: decoding decisions based on the extracted maps). The total number of trials extracted within each dataset did not differ significantly across subjects for accept vs. reject decisions for any level of difficulty (paired t -tests, $|t| \leq 1.49$, $p \geq 0.17$ for training/test datasets and $|t| \leq 1.25$, $p \geq 0.24$ for the validation dataset).

Since it is not possible to estimate beforehand the optimal number of Gaussians in the GMM models, we trained a range of models with a total number of Gaussians from 3 up to 11 for each condition. This range of parameters was chosen in an empirical way, based on the range that provided satisfying results on previous studies using a similar approach (Tzovara et al., 2012a). On a second step we selected the optimal pair of models by maximizing decoding performance across the 10 test datasets, in a cross validation procedure.

We made the hypothesis that the neural generators that are responsible for each decision (accept or reject) are each reflected in one voltage topography. Therefore, by maximizing discrimination between accept vs. reject decisions in the training dataset, we extracted two template maps, one per type of decision.

Based on the computed GMM models, we quantified the degree to which the recorded voltage topographies are represented by the template maps with posterior probabilities:

$$P(c_k|\mathbf{m}) = \frac{P(\mathbf{m}|c_k) \cdot p_k}{p(\mathbf{m})} \quad (4)$$

where \mathbf{m} refers to a recorded topography and contains an array of voltage values recorded at individual electrodes, c_k to one Gaussian within the mixture (one template map), $p(\mathbf{m})$ to the unconditional density function, which is the probability of observing \mathbf{m} irrespective of the Gaussian in the mixture, and p_k to the prior probability of the k -th Gaussian.

The Bayes Factor (BF; Raftery, 1995) was used to quantify the relative probability for a given topography to be represented by the model computed from accept vs. the model computed from reject trials. So, at each latency h within the t -th trial, the BF was computed for any given set of topographies c_{Accept_k} and $c_{Reject_{k'}}$, as:

$$BF_{t,h} = \frac{P(c_{Accept_k}|\mathbf{m}_{t,h})}{P(c_{Reject_{k'}}|\mathbf{m}_{t,h})} \quad (5)$$

where $\mathbf{m}_{t,h}$ is the observed topography of the t -th trial, at latency h . c_{Accept_k} and $c_{Reject_{k'}}$ refer to the k -th template map within the GMM generated for accept and reject decisions, respectively. k and k' range from 1 up to the total number of Gaussians in each of the two different models. The indexes k and k' are not denoted in the BF for reasons of simplicity, as only one pair of mixtures is eventually selected (see below).

The Bayes Factor was computed in the same way for topographies belonging to both conditions and represented a measure of the confidence with which we can assign a given topography ($\mathbf{m}_{t,h}$) to the template maps corresponding to each decision (c_{Accept_k} or $c_{Reject_{k'}}$). In general, a BF greater than 1 suggests that the specific topography is better represented by the template map for accept than for reject decisions and vice versa. However, a BF of less than 20 cannot be easily interpreted as a strong statistical evidence of the predominance of one template map over the other (Raftery, 1995; Wetzels et al., 2011). For this reason, we never considered the absolute value as an evidence for a difference between accept and reject decisions. The BF was used for building a discriminant function as described below.

2.5.1.2. Validation: decoding decisions based on the extracted maps. For decoding accept vs. reject decisions we considered as discriminant function the accumulation of the $BF_{t,h}$ across the whole trial t :

$$DF_t = \frac{1}{L_t} \sum_{i=1}^{L_t} (BF_{t,i}) \quad (6)$$

where L_t is the t -th trial's length in time-points. This discriminant function, obtained by averaging the BF over time, corresponds to a majority vote rule for assigning each trial to the accept or reject condition. This rule is known to be more robust to errors and to outperform other classifiers, such as a product rule (Kittler et al., 1998; Tax et al., 2000).

For every latency and trial of the training dataset we computed all possible combinations of the Bayes Factor (of k and k') and selected the most discriminant pair of template maps. These template maps were extracted from the training datasets and were then used for decoding accept vs. reject decisions in the test and validation datasets.

Decoding performance was measured as the Area Under the receiver operator characteristic Curve (AUC Green and Swets, 1966), which ranges from 0 to 1 for a perfect classification. Chance levels were quantified by randomly shuffling the true labels of the training datasets and recomputing the GMM models 100 times. We

then computed the decoding performance in the validation dataset based on these random models and estimated a distribution of chance levels. We compared the decoding performance obtained on the basis of the real GMM models with this distribution of chance values with a signed Wilcoxon ranking test ($p < 0.001$).

2.5.1.3. Detection of critical time-points. The decoding scheme described above allowed predicting subjects' decisions based on single-trial EEG activity, recorded at any time-point within a given trial. For tracking the timing of this single-trial activity, in a data-driven way, we accumulated the BF over time:

$$ABF_{t,h} = \frac{1}{h} \sum_{i=1}^h (BF_{t,i}) \quad (7)$$

where h refers to any time-point within the t -th trial, thus h ranging from 1 up to L_t , where L_t is equal to the t -th trial's length. Please note that for $h=L_t$ Eq. (7) is equivalent to Eq. (6) i.e. the average BF over the whole trial's length.

We observed in the trials of the training dataset that the accumulated BF typically increased or decreased until it reached a 'plateau', after which its values remained relatively stable (Fig. 2e for few exemplar trials). In order to detect quantitatively those plateaus we considered the first derivative of the accumulated BF. Formally, a plateau would correspond to a 'small' change in this derivative. For each trial of the training dataset we computed the distribution of the values of the derivatives (computed as the difference of the BF between contiguous time points) irrespective of time. We then considered the lowest values corresponding to the first 5% of this distribution. By keeping the median of these lowest 5% across all trials we extracted one threshold for each subject which would correspond to a plateau in the ABF.

We then used this threshold on the test and validation trials to examine whether they would also exhibit a plateau. This generalization step is crucial in determining the suitability of the method we used here for plateau detection. In the case of the test trials, we searched for periods of stability, starting from the beginning of each trial and moving on. We refer to the onset of these plateaus as critical time-points (t^*) and we hypothesized that they correspond to the time-points where enough evidence has already been accumulated for accurately decoding subjects' decisions. We tested this hypothesis by decoding accept vs. reject decisions, by taking into account activity from the beginning of the trial up to the critical time-points. We expected that if a decision has already been reached at the t^* , decoding performance should be (a) above chance levels and (b) not significantly different from what was obtained by considering the accumulation of the BF during the whole trial. Both of these conditions were tested with Wilcoxon signed rank tests ($p \leq 0.001$).

2.5.2. Comparison with average ERP analysis

For reasons of comparison with what can be obtained from a 'classic' average event-related potential (ERP) analysis, we analyzed average ERPs at the level of the group, in a two-factorial design (decision \times difficulty). For this part of the analysis, the extracted trials (corresponding to accept and reject decisions) were further sub-divided in two categories each, according to the level of difficulty (easy vs. hard). We considered 'easy' trials those where the absolute difference between the range of rewards and losses was of 4 or 5 and 'hard' those with a difference of 1 or 2. The length of each trial was equal to the single value of reaction time. For this analysis we kept only trials with a reaction time longer or equal than 500 ms and considered activity only up to 500 ms for all of them (on average 60 ± 7 trials per condition and subject, corresponding to $27 \pm 3\%$ of all the available single-trials).

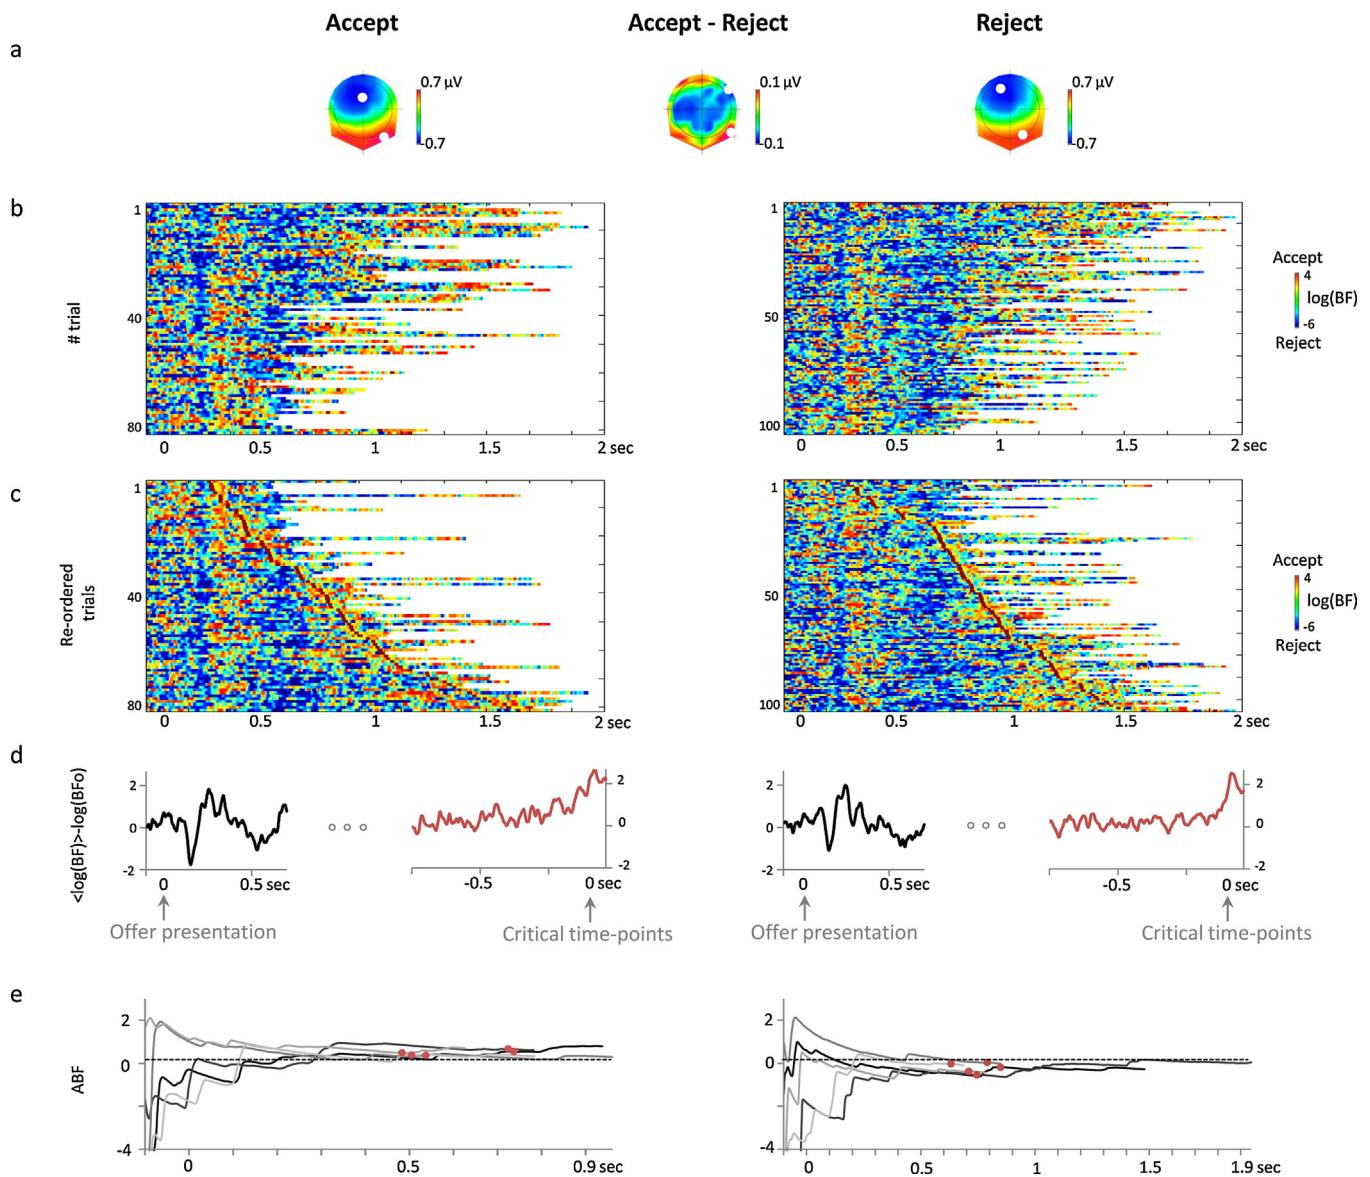


Fig. 2. Results of single-trial analysis for an exemplar subject. (a) Template maps for accept and reject decisions. The middle map corresponds to the difference between accept and reject, and highlights the pattern of the different responses across the scalp, for display purposes. White dots indicate the position of the maximum and minimum values across all electrodes. (b) The value of the BF across time and trials reflects how likely is one time-point to be represented by the template map extracted for accept decisions (warmer colors) vs. the one for reject decisions (cold colors). The onset of the trial (0 s) corresponds to the presentation of the offer. (c) BF across trials when re-ordering single trials according to their critical time-points (t^* , red points). For display purposes, we only show trials for which we could detect a critical time-point before the end of the trial, in both panels b and c. (d) Average BF across trials ($\langle \log(BF) \rangle$), aligned to the offer presentation (black lines) or to the critical time-points (red lines), for the two types of decisions. When time-locking at the beginning of the trial we could observe an early component, peaking at ~150 ms after offer presentation for both accept and reject decisions (black plots in left and right panels). By contrast, when realigning single-trials according to their critical time-points and re-computing the average BF, a late component appeared, related to the critical time-points (red plots in left and right panels). The time-course of the average log(BF) across trials is shown for display purposes with respect to its initial value at time 0 $\langle \log(BF_0) \rangle$. (e) Accumulation of the BF over time (gray lines) and t^* (red points) for five representative trials for each level of difficulty (darker lines correspond to harder decisions). As we consider more evidence over time, the ABF of single trials initially tend to increase or decrease over time, until they reach a plateau, after which they remain relatively stable. The black dashed lines highlight an indicative threshold, above which the ABF corresponds to accept and below which to reject decisions. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Based on these average ERPs, we calculated the Global Dissimilarity (Lehmann and Skrandies, 1980) for every time point, which computes differences between electric fields, independently of their strength. Statistically, this was evaluated by applying a Monte Carlo bootstrapping procedure (TANOVA; Murray et al., 2008), with two factors (decision and difficulty). This analysis can reveal periods of topographic difference between experimental conditions, which forcibly correspond to activations of different types of neural generators for each condition.

Additionally, we analyzed differences in the strength of the electric field, independently of its topographic distribution, as measured by the Global Field Power (GFP; Murray et al., 2008). The GFP corresponds to the standard deviation across electrodes at every instant in time. We assessed differences in the GFP using non parametric randomization statistics in a two-by-two design, performed for every time-point. For both contrasts of GFP and Dissimilarity, we used RAGU software (Koenig et al., 2011). Temporal autocorrelation was corrected by applying a criterion of at

least 15 contiguous time-points for the presence of a statistical effect.

2.5.3. Comparison with other single-trial techniques

We compared the level of decoding performance we can obtain with the single-trial topographic analysis with two other popular single-trial techniques. The first one is logistic regression (Parra et al., 2002, 2005), an algorithm which has been widely applied for studying perceptual decision-making (Philiastides and Sajda, 2006; Ratcliff et al., 2009) and in single-trial decoding in general (Brandmeyer et al., 2013; Farquhar and Hill, 2013). The second is Support Vector Machines (SVM), a powerful machine learning technique which has been used in various decoding applications (Rieger et al., 2008; Salvaris and Sepulveda, 2009; Schulz et al., 2012; Taghizadeh-Sarabi et al., 2014).

These approaches are typically based on selecting a sliding window of predefined temporal length (Ratcliff et al., 2009). Even though the use of such a sliding window has proven useful for studying perceptual decision-making, in our case we trained a classifier over the whole trials' length, for two reasons: the first is that not all trials had the same length, so we would need to only consider a short period from the beginning of the trial up to the shortest reaction time and we would thus lose important information for the majority of the trials. The second reason is that our decision-making task had a very high variability and a given time-window would not necessarily reflect the same processes for easy and hard decisions. For these reasons, we considered single-trial EEG activity over the whole trials' length both for logistic regression and SVM.

For easing the comparison with both methods we kept the same cross-validation and validation data splits as for the original single-trial analysis. For each subject, we computed the necessary parameters for classification (see below) based on the cross-validation trials. The validation dataset was always kept aside for a final validation of our results. Chance level was always estimated by randomly permuting the labels of cross-validation trials and re-computing all the parameters 100 times. These random parameters were then used for classifying the validation trials. Classification performance between this random distribution and the true value obtained using the original parameters was compared with a Wilcoxon signed-rank test ($p < 0.001$). Below we present the basic steps for the two algorithms that we used.

2.5.3.1. Logistic regression. The logistic regression algorithm is a generative classifier which models the single-trial EEG data as a logistic function, parameterized by $\mathbf{w} = (w_1, \dots, w_N)$:

$$P(k=1|\mathbf{m}, \mathbf{w}) = \frac{1}{1 + \exp(-\mathbf{m}^T \mathbf{w})} \quad (8)$$

where k is 0 or 1, according to the true labels of the two classes (accept or reject decisions). It can be shown that assuming such a distribution leads to a linear hyperplane, separating the two classes:

$$y(t) = \mathbf{w}^T \mathbf{m}(t) = \sum_{i=1}^N w_i m_i(t) \quad (9)$$

In the present study, the computation of the parameter \mathbf{w} is based on a penalized likelihood estimator with an l2-regularization (Conroy et al., 2012). All the computations have been performed based on the matlab scripts provided by the Laboratory for Intelligent Imaging and Neural Computing (http://liinc.bme.columbia.edu/mainTemplate.htm?liinc_downloads.htm) and described in (Parra et al., 2002, 2005).

2.5.3.2. Support vector machine. Similar to the logistic regression, the SVM algorithm aims at computing a hyperplane which

separates the two experimental conditions (accept vs. reject decisions), although in this case no assumptions on the statistical distribution of the data is made. The most basic form of SVM is the linear one, where the SVM computes a linear function:

$$f(\mathbf{m}) = \mathbf{w}^T \Phi(\mathbf{m}) + b \quad (10)$$

where w is a weight vector, Φ is a fixed feature space transformation, \mathbf{m} a vector with the single-trial EEG data, averaged over the whole trial's length and b is a bias term. Here, we used a linear feature space transformation.

The support vector machine approaches the computation of the separating hyperplane expressed by introducing the concept of 'margin'. Specifically the margin is the smallest distance between the hyperplane and any of the samples in the data. The goal of the support vector machine is to maximize this distance which translates in its simplest formulation in minimizing a cost-function:

$$\frac{1}{2} \mathbf{w}^T \mathbf{w} = \frac{1}{2} \|\mathbf{w}\|^2 \quad (11)$$

with the constraint that

$$k_i(\mathbf{w}^T \Phi(\mathbf{m}_i) + b) \geq 1 \quad (12)$$

where i spans over all the training samples and k_i is the class label of the vector \mathbf{m}_i .

In practice, the support vector machine allows some of the training data points to be misclassified which translates into adding a regularization term to the Eq. (11) controlling the trade-off between minimizing the classification error and the model complexity. This term here was set to 1. For the SVM calculation, we used the LIBSVM toolbox (Chang and Lin, 2011). For the SVM algorithm we did not need to carry out a separate cross-validation for parameter selection and a validation phase of the selected parameters. For this reason, we computed the separating hyperplane based on the cross-validation trials and report the decoding performance on the validation dataset.

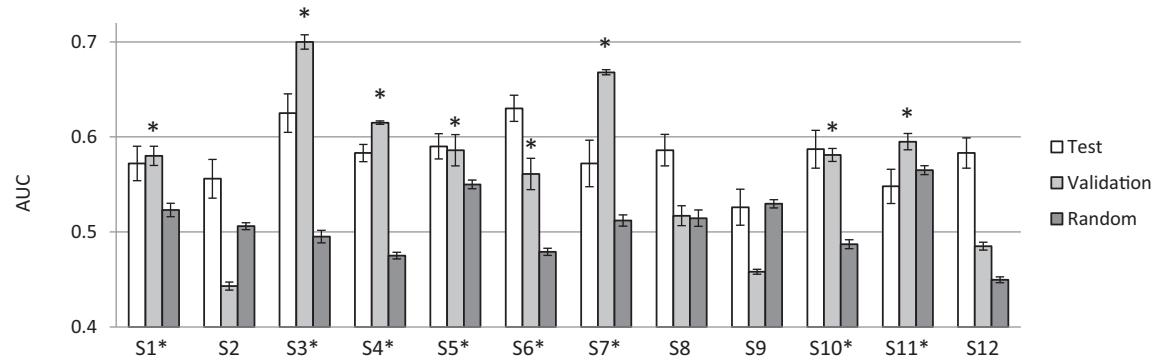
3. Results

3.1. Behavior

During the training phase of the experiment, subjects learned to associate colors and shapes to corresponding rewards and losses. Their performance increased in the three training blocks and reached an average maximum of $92 \pm 1\%$ correct responses for rewards and $94 \pm 2\%$ for losses. Average reaction times during training were 1076 ± 37 ms for rewards and 1051 ± 42 ms for losses and did not differ significantly for the two types of training (paired t -test, $t(11) = 0.73$, $p = 0.48$).

During the main task, reaction times decreased as the level of difficulty of the decision decreased (Fig. 1b), while performance (ratio of correct responses) increased (Fig. 1c). This behavior, both in terms of reaction times and accuracy, was in accordance to what has been shown in a previous study with a similar task (Basten et al., 2010). Subjects' behavior during the main task was modeled with a DDM whose parameters were optimized for each subject separately. Median drift rates across subjects followed a similar pattern and increased as the level of difficulty decreased (Fig. 1d), resulting in faster reaction times. The rest of the parameters of the DDM was kept constant across difficulty levels and included the mean boundary separation (a), which quantifies the distance between the two boundaries of a diffusion process, the starting point of the diffusion process (zo), as well as the total non-decision time (Ter). The mean boundary separation across subjects was 0.23 ± 0.02 , the mean starting point of the diffusion (zo) was 0.12 ± 0.01 and the total non-decision time (Ter) was 559 ± 4 ms. This last parameter reflects the time needed for processes that do not contribute in the

a. Decoding performance based on the whole trial



b. Decoding performance up to the critical time-points

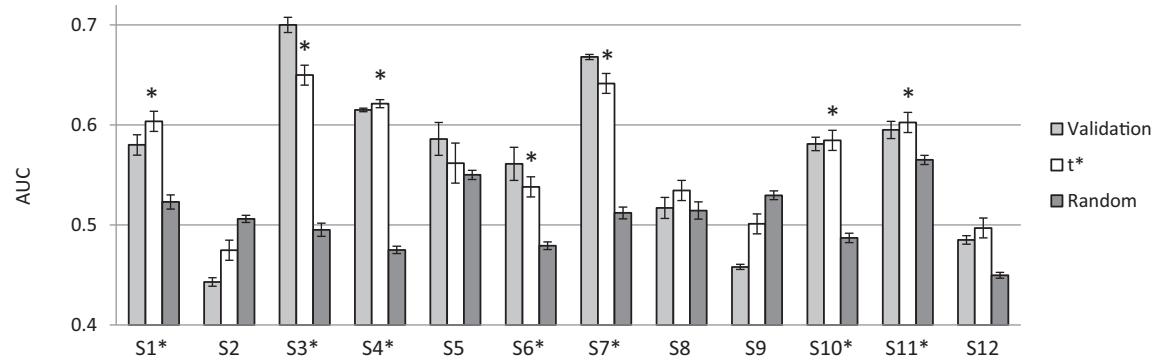


Fig. 3. Decoding performance for all subjects. (a) Decoding performance in test (white bars) and validation datasets (light gray bars), as well as chance levels (dark gray bars) for all subjects of the experiment. Subjects with significant decoding results (Wilcoxon signed rank test, $|z| \geq 3.3, p \leq 0.001$) in the validation dataset are denoted with an asterisk ($n=8$ subjects). Chance levels were computed by randomly permuting the labels of the training and test trials (cross validation dataset) and re-computing the template maps 100 times. (b) Decoding performance when taking into account information only up to the critical time-points (t^* , white bars) was above chance levels for 7 out of the 8 significant subjects, denoted here with asterisks (Wilcoxon test, $|z| \geq 4.02, p \leq 0.001$). An above chance decoding suggests that the neural generators responsible for making a decision have already started differentiating their responses for each trial at the critical time-points and provide further support that the t^* reflect time-points when the decision has already been made at a neural level. The error bars on panels (a) and (b) represent the standard errors.

decision itself, such as the encoding of visual stimuli, or the motor response (VandeKerckhove and Tuerlinckx, 2007).

3.2. Multivariate EEG decoding based on accumulation of topographic activity

By accumulating over time the BF (Fig. 2a and b for an exemplar subject), we were able to decode above chance levels subjects' decisions for 8 out of 12 subjects (Wilcoxon test, $|z| \geq 3.3, p \leq 0.001$). For these 8 subjects, decoding performance in the validation dataset was on average 0.61 ± 0.02 while chance level was 0.51 ± 0.01 (Fig. 3a, subjects with an asterisk). Decoding performance across all subjects was 0.58 ± 0.03 and 0.57 ± 0.08 in the test and validation datasets (Fig. 3a). An above chance level decoding for the majority of the subjects suggests that the extracted template maps were indeed informative of the underlying decisions (accept vs. reject) and able to generalize to new trials.

3.3. Temporal evolution of single decisions

The accumulated BF over time, (ABF), tended in general to increase or decrease as we considered more evidence, until it reached some plateau (Fig. 2e for few representative trials). After this plateau, the ABF remained relatively stable, possibly because the decision had already been taken at some point before (Fig. 2). Based on the first derivative of the ABF we detected the onset of these plateaus (critical time-points, t^* ; Fig. 2e red points). Fig. 2c displays the detected critical time-points (t^*) across trials for an

exemplar subject (Fig. 2, red points). These t^* were detected on average at ~ 590 ms after the offer display, or at ~ 340 ms before the subsequent response (Fig. 4a, left panel for the distribution of the t^* across all subjects). We were able to detect such t^* for 75% of the single-trials on average (range 61 up to 90% across subjects).

Importantly, when considering time-locked activity across trials based on the BF, we observed an early component peaking at ~ 150 ms after the display of the payoff for both accept and reject decisions (Fig. 2b and d, black line). Re-ordering single trials according to the detected t^* revealed an additional late component locked to the t^* which was previously masked (Fig. 2c and d, red lines). These two components closely match previous findings in perceptual decision-making, where early components appeared around 170 ms and were linked to encoding of sensory information, while late components were also shifted across trials and linked to the decision process (Philiastides and Sajda, 2006).

In order to validate the role of the detected t^* , we re-computed accuracy of decoding in the validation dataset for the significant subjects by taking into account only information up to the critical time-points for each trial. For each subject separately, we compared the decoding performance based on the whole trial vs. only up to the t^* , based on a Wilcoxon test. Average decoding performance when only considering activity up to the t^* was 0.60 ± 0.01 and at similar levels to the decoding results based on the activity extracted from the whole trials' length (0.61 ± 0.02 ; Fig. 3b). Importantly, the decoding performance up to the t^* was significantly above chance levels for 7 out of the 8 subjects that had significant results when using the whole trial (Wilcoxon signed rank test,

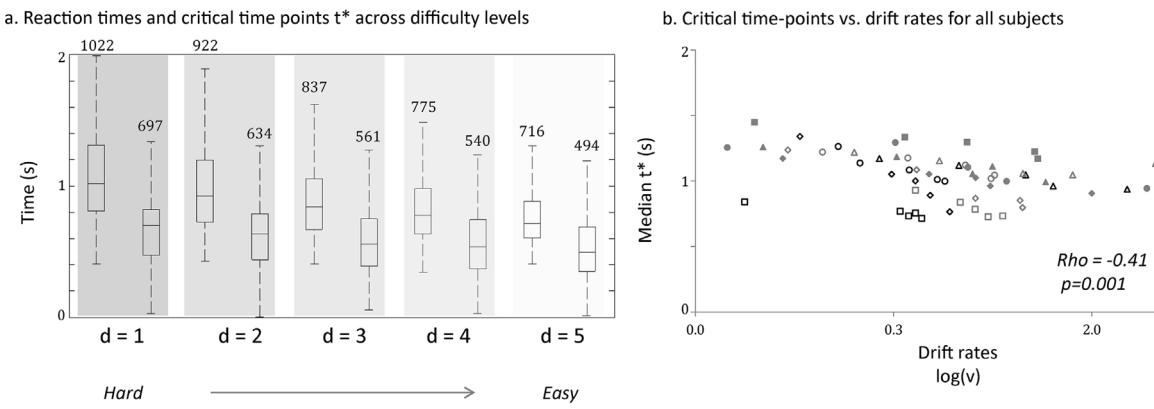


Fig. 4. Reaction times and critical time-points (t^*) across difficulty levels. (a) The boxplots show for each level of difficulty the median reaction times (on the left) and t^* (on the right), across all subjects. As the level of difficulty increases (from $d = 5$ to $d = 1$) both reaction times and t^* are longer and variability increases. (b) The median drift rates (v) per subject and difficulty level significantly correlated with the corresponding median t^* (Spearman's $\rho = -0.41$, $p = 0.001$). This negative correlation suggests that longer t^* are associated with a slower accumulation process, or smaller drift rates. For display purposes only, the drift rates are plotted in a semi-logarithmic scale and different subjects are shown with different symbols.

$|z| \geq 4.02$, $p < 0.001$, Fig. 3b, subjects with an asterisk). These results suggest that already at or before the detected t^* , the neural generators of the subsequent decisions have started differentiating their responses.

3.4. Modulation of t^* by decision difficulty

We tested whether the amount of evidence needed by our classifier to accurately decode subjects' decisions, as indicated by the t^* , was modulated by the level of difficulty by computing the correlation between the t^* and the corresponding reaction time for each trial. We found a positive and significant correlation for all subjects (Spearman's $\rho > 0.24$, $p < 0.001$), suggesting that greater t^* are associated with longer reaction times.

An additional link between t^* and subjects' behavior was revealed when examining individual subjects' drift rates. We considered the median t^* and the median drift rates across subjects and difficulty levels and we found that they significantly correlated (Spearman's $\rho = -0.41$, $p = 0.001$; Fig. 4b). This negative correlation shows that smaller drift rates (or slower decision-times) correspond to longer t^* and vice versa. We should note here that these results might be partly driven by the fact that the trials' lengths vary across trials and are equal to the corresponding reaction times. It is therefore more likely to detect a higher t^* in a longer trial, than in a shorter one. This confound might also partly explain the correlation of t^* with drift rates, because the parameters of the behavioral model were based on the subjects' reaction times.

Moreover, we also observed that the variability of t^* (as measured by its standard deviation) is modulated by the difficulty of the decision. The distribution of t^* across all subjects (Fig. 4a) revealed both longer latencies and greater variability of t^* for 'hard' (median $t^* = 697$ ms, standard deviation SD = 274 ms) compared to 'easy' decisions (median $t^* = 494$ ms, SD = 241 ms). This difference in t^* variability was found to be highly significant using Levene's test for equality of variances, with five levels ($F = 3.88$, $p = 0.004$). This suggests that 'hard' decisions are more noisy than 'easy' ones, even though the contextual factors influencing the decision remain the same across trials and difficulty levels. It is worth noting here that tracking this variability at the neural level across levels of difficulty was only possible because we could determine the time of decision at the level of each single trial.

3.5. Comparison with average ERP analyses

For reasons of comparison with 'classic' ERP analysis, we additionally considered EEG data time-locked to the offer presentation,

at the grand average level in the group of twelve subjects. Based on time-point by time-point permutation tests we separately analyzed differential responses as a function of difficulty level and decision type in terms of response strength (measured by the GFP at average level for each subject) and in terms of configurations of voltage topographies.

Visual inspection of the GFP revealed a first peak at ~ 135 ms after offer presentation, which was similar for all conditions (Fig. 5a). The GFP permutation statistics revealed a significant interaction in response strength (decision * difficulty, $p < 0.05$) later in time, at 394–426 ms post-stimulus onset (Fig. 5a). No main effect of difficulty or decision was detected. By contrast, the TANOVA analysis revealed a main effect of decision (bootstrapping procedure, $p < 0.05$) on the configuration of voltage topographies during two distinct periods, the first occurring at 243–285 ms and the second at 395–437 ms after the offer display (Fig. 5b). This second period also overlapped with the period over which a significant interaction of decision * difficulty was found on the GFP. Both of these time-locked periods were identified well before the typical latencies of the t^* (Fig. 4a), suggesting that they could likely reflect processes linked to encoding of the presented offers.

Finally, we examined the link between results obtained from average ERP analysis and from single-trial decoding. In a similar way to our previous study (Tzovara et al., 2012a), we examined whether the two time-periods revealed by the TANOVA also carry predictive power of single decisions. To test this, we restricted the single-trial decoding of accept vs. reject decisions only to the periods identified from the average ERP analysis. In this case, decoding performance was lower both on test and on validation datasets. For the early period, ranging between 243 and 285 ms post-stimulus onset, decoding performance in the test dataset was on average across subjects 0.51 ± 0.01 and in the validation dataset 0.51 ± 0.02 . For the late period, ranging between 395 and 437 ms decoding in the test dataset was 0.53 ± 0.01 and in the validation 0.52 ± 0.04 . Only 5 out of 12 subjects were significantly above chance levels for the first period and 6 out of 12 during the second (Wilcoxon test, $|z| \geq 4.8$, $p < 0.001$), in both cases less subjects than when accumulating information during the whole trial (i.e. 8 out of 12 subjects).

3.6. Comparison with other single-trial techniques

We compared the decoding performance obtained with the single-trial topographic analysis with logistic regression (Parra

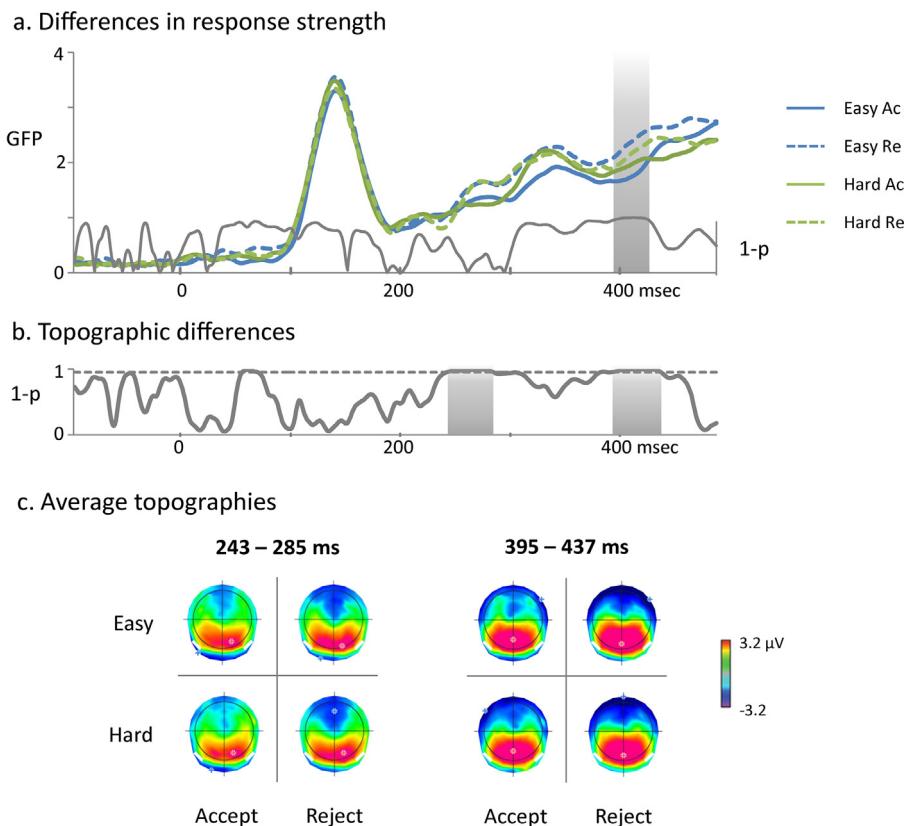


Fig. 5. Time-locked average ERP analysis. (a) Global field power of the grand average ERPs for different types of decisions and results of permutation statistics (highlighted period). A significant interaction decision * difficulty was observed in response strength at 394–426 ms post-stimulus onset (axis on the right $p < 0.05$). The labels 'Ac' and 'Re' correspond to 'accept' and 'reject', respectively. (b) Topographic differences at the grand average level. The TANOVA analysis revealed a main effect of decision ($p < 0.05$) during an early period, at 243–285 ms and also during a late one, at 395–437 ms post-stimulus onset. (c) Average topographies during the two periods identified by the TANOVA in (b) for all experimental conditions.

et al., 2005; Ratcliff et al., 2009) and SVM (Chang and Lin, 2011; Rieger et al., 2008).

The average decoding performance with the logistic regression across all subjects was 0.56 ± 0.02 for the cross-validation and the validation datasets. This performance across subjects was at similar levels as with the single-trial topographic analysis (paired t -test, $|t| = 0.92$, $p = 0.38$ for the cross-validation and $|t| = 0.19$, $p = 0.85$, for the validation datasets). When using the SVM, the average decoding performance in the validation dataset was 0.57 ± 0.02 . This average performance did not differ significantly from the one obtained with our method (paired t -test, $|t| = 0.16$, $p = 0.87$).

However, when examining the significance of results at the single-subject level, the topographic analysis outperformed both the logistic regression and the SVM, as the total number of subjects showing significant results was higher: we obtained above-chance level decoding for 8/12 subjects based on the topographic analysis, versus 7/12 and 7/12 subjects based on the logistic regression and SVM respectively. The results of these comparisons are summarized in Table 1.

4. Discussion

We presented a novel approach for studying the temporal evolution of single value-based decisions. This consists of a multivariate EEG decoding analysis which extracts time-unlocked topographic activity, related to the underlying decisions. Using this approach, we were able to quantify, for the first time, a correlate of the time when single decisions are made and we showed that it is possible to predict decisions much earlier than the behavioral response (~340 ms before).

4.1. Comparison with other single-trial methods

We compared the performance of the single-trial topographic analysis with two widely used single-trial classification techniques: logistic regression (Parra et al., 2002, 2005; Ratcliff et al., 2009) and SVM (Rieger et al., 2008; Taghizadeh-Sarabi et al., 2014). Despite similar levels of decoding performance across all subjects, the single-trial topographic analysis gave significant results at the single-subject level for more subjects than either of the two techniques. Moreover, it has the additional advantage of extracting interpretable features that can be tracked across time (i.e. the template maps for each type of decision and the BF).

Indeed, the main novelty of our study consists of extracting discriminant patterns of time-unlocked activity at the single-trial level. Our approach differs not only from 'classic' average ERP analysis, but also from the majority of previous single-trial techniques. Indeed, most of the existing approaches for studying decision-making at the single-trial level rely mainly on preselected

Table 1

Summary of decoding results obtained with single-trial topographic analysis, logistic regression and support vector machine. The results shown correspond to the average AUC values across 12 subjects, obtained in the cross-validation and validation datasets. The third row shows the total number of subjects with above-chance level results.

	Single-trial topographic analysis	Logistic regression	Support vector machine
Cross-validation	0.58 ± 0.01	0.56 ± 0.02	–
Validation	0.57 ± 0.02	0.56 ± 0.02	0.57 ± 0.02
# Significant subjects	8	7	7

time-intervals, or on sliding windows, whose length and latency has to be identified beforehand (Billeke et al., 2013; Bode et al., 2012; Philiastides and Sajda, 2006; Zhang et al., 2014). Other techniques consist of model-based approaches, such as regression of single-trial EEG activity, mostly using behavioral models or subjects' responses as regressors (Boehm et al., 2014; Fischer and Ullsperger, 2013; van Vugt et al., 2012; Wyart et al., 2012).

While model-based approaches can be used to test specific hypotheses, our study presents a data-driven alternative for analyzing single-trial data. This analysis does not require strong assumptions about the experimental factors influencing neural responses, or about the latency/location of the expected effects. Importantly, it is suitable for assessing time-unlocked EEG activity, in a dynamic way as it gives the possibility to study trials with different lengths (as for example (van Vugt et al., 2012) also did, by modeling events occurring in the continuous EEG signal). This is potentially of interest in paradigms where subjects' reaction times are strongly modulated by experimental conditions and therefore the periods of interest might change significantly across the experiment (Bode et al., 2012; Boehm et al., 2014).

This approach could also be potentially useful for online applications, such as in the field of brain-computer interfaces (BCI; Millan et al., 2010). BCI applications usually require a much higher decoding performance than the one we obtained here, but nevertheless, the detection of the critical t^* for each single decision could provide a realistic upper bound of when to stop accumulating EEG activity for accurately predicting a decision. Our results indicate that an accurate prediction can be made in the order of few hundreds of milliseconds before the actual decision. This information could be used for improving the speed of BCI applications. Even though the speed of calculations was not the main goal of the present study, future experiments could evaluate and optimize the performance of the decoding algorithm in real-life situations. In addition, future investigations could aim at increasing the accuracy of our algorithm, possibly by optimizing the extracted spatiotemporal features.

4.2. The timing of single decisions

From a neurophysiological point of view, the results we obtained using this novel approach are in line with recent developments in the field of decision-making, which suggest that perceptual and value-based decisions may rely on similar mechanisms of evidence accumulation (Basten et al., 2010; Hunt et al., 2012). This claim comes mainly from model-based analyses of neuroimaging data, which have identified brain regions whose activity correlates with parameters of diffusion models assessed at behavioral level (Basten et al., 2010; Gluth et al., 2012; Hunt et al., 2012). However, the temporal dynamics of evidence accumulation have been mainly studied in the field of perceptual decision-making (O'Connell et al., 2012; Philiastides and Sajda, 2006; van Vugt et al., 2012; Wyart et al., 2012).

In particular, EEG studies have revealed two discrete components, an early one, peaking at ~ 170 ms which is related to the encoding of sensory information and a late one, which is shifted across trials, according to reaction time and is related to the perceptual decision and amount of available evidence (Philiastides and Sajda, 2006). Our analysis at the single-trial level parallels those previous findings, as it also identified an early, time-locked component at ~ 150 ms after offer presentation and a late one, shifted across trials and depending upon the level of difficulty (Fig. 2d). Recent studies in perceptual decision-making have also identified neural correlates of an accumulation mechanism, mainly in low frequency bands of EEG (van Vugt et al., 2012) and have shown that perceptual decisions are influenced by the phase of ongoing delta band oscillations during stimulus encoding (Wyart et al., 2012).

Other studies have identified EEG signals corresponding to various stages of perceptual decisions, irrespective of the sensory modality and even in the absence of a behavioral response (O'Connell et al., 2012) and have shown that EEG activity over centro-parietal electrodes correlates with subjects' reaction times (Kelly and O'Connell, 2013).

In the case of value-based decisions, fine-grained temporal information has been coupled with behavioral models, in order to isolate various stages of decision-making, from value encoding to formation of a response (Gluth et al., 2013; Hunt et al., 2012). In particular, Hunt et al. (2012) have assessed information flow among different brain regions, within pre-defined regions of interest (Hunt et al., 2012). Another related recent study has shown that amplitude modulations on central electrodes reflect emerging value-based decisions, in terms of readiness potentials (Gluth et al., 2013). Our approach complements this work, as it takes advantage of the whole electrode montage, providing a global measure of neural responses preceding the formation of single decisions.

Importantly, the present study builds upon a previous work (Tzovara et al., 2012a), where we also detected timing information based on an exploration/exploitation task. This previous study showed that decisions had already been made at ~ 520 ms after reward presentation of the previous trial. However, this timing was only detected at an across-trials level, leaving under-explored single-trial variability. Our present results generalize this approach by detecting critical time points at the level of single-trials. This has also allowed quantifying the variability in the time of decisions at a neural level. Indeed, we showed that as the level of difficulty increases, the variability in the latency of decisions also increases, despite otherwise identical experimental conditions. This result complements previous findings in perceptual decision-making, studying endogenous and external influences on an evidence accumulation mechanism (Kelly and O'Connell, 2013). Future studies could investigate the sources of this variability and their link to behavioral performance and cognitive state.

4.3. Value encoding phase

In addition to the single-trial analysis, the average ERPs identified two distinct time periods over which neural activity prior to accept vs. reject decisions differed in terms of voltage topographies. Both of them occurred before the majority of the critical time-points: the first one at ~ 240 – 285 ms and the second at ~ 395 – 440 ms after offer presentation. The timing of these distinct periods is compatible with the existence of an encoding phase of the incoming stimuli, a comparison of their values and a third successive process in time related to the decision which varies from trial to trial (as proposed in Basten et al., 2010). The low decoding performance obtained when considering only the first two periods instead of the whole trial supports this interpretation.

It is worth noting here, that another important point of the present study is the analysis of EEG activity from the value encoding phase. The vast majority of M/EEG studies in value-based decision making focused on the period after subjects receive feedback about their decisions (for example Cohen and Ranganath, 2007; Frank et al., 2005; Pedroni et al., 2011; Yeung et al., 2005) and examine the effect of this feedback on subsequent behavior. Only few studies have assessed the temporal dynamics of the value encoding phase in humans, right after an offer presentation (Gluth et al., 2013; Hunt et al., 2012; Minati et al., 2012) but these were either based on ERP responses at single electrodes (Gluth et al., 2013; Minati et al., 2012), or were restricted within specific brain regions (Hunt et al., 2012). Our study instead examines global measures of EEG activity and provides further insight into the processes taking place while

subjects see an offer and encode its value by quantifying the time of the single decision for each subject.

5. Conclusions

In summary, we presented a novel approach for studying trial-by-trial modulations in value-based decision-making and quantifying the time when single decisions are made at a neural level. To the best of our knowledge, this is the first study to extract, in a data-driven way, differential patterns of EEG activity which are compatible with an accumulation mechanism. Such fine-grained temporal information fills a critical gap in combining the theoretical diffusion model of decision making and direct neuroimaging evidence at the level of the single decision.

References

- Basten U, Biele G, Heekeren H, Fiebach C. How the brain integrates costs and benefits during decision making. *Proc Natl Acad Sci U S A* 2010;107:21767–72.
- Bernasconi F, De Lucia M, Tzovara A, Manuel AL, Murray MM, Spierer L. Noise in brain activity engenders perception and influences discrimination sensitivity. *J Neurosci* 2011;31:17971–81.
- Billeke P, Zamorano F, Cosmelli D, Aboitiz F. Oscillatory brain activity correlates with risk perception and predicts social decisions. *Cereb Cortex* 2013;23:2872–83.
- Bishop CM. Neural networks for pattern recognition; 1995.
- Bode S, Sewell DK, Lilburn S, Forte JD, Smith PL, Stahl J. Predicting perceptual decision biases from early brain activity. *J Neurosci* 2012;32:12488–98.
- Boehm U, van Maanen L, Forstmann B, van Rijn H. Trial-by-trial fluctuations in CNV amplitude reflect anticipatory adjustment of response caution. *Neuroimage* 2014;96:95–105.
- Brandmeyer A, Sadakata M, Spyrou L, McQueen JM, Desain P. Decoding of single-trial auditory mismatch responses for online perceptual monitoring and neurofeedback. *Front Neurosci* 2013;7:265.
- Brunet D, Murray MM, Michel CM. Spatiotemporal analysis of multichannel EEG: CARTOOL. *Comput Intell Neurosci* 2011;2011:813870.
- Chang CC, Lin CJ. LIBSVM: a library for support vector machines. *ACM Trans Intell Syst Technol* 2011;2.
- Cohen MX, Ranganath C. Reinforcement learning signals predict future decisions. *J Neurosci* 2007;27:3711–8.
- Conroy B, Sajda P. Fast exact model selection and permutation testing for l-regularized logistic regression. *JMLR Workshop Conf Proc* 2012;22:246–54.
- Cossy N, Tzovara A, Simonin A, Rossetti AO, De Lucia M. Robust discrimination between EEG responses to categories of environmental sounds in early coma. *Front Psychol* 2014;5:155.
- De Lucia M, Tzovara A, Bernasconi F, Spierer L, Murray MM. Auditory perceptual decision-making based on semantic categorization of environmental sounds. *Neuroimage* 2012;60:1704–15.
- Dempster A, Laird N, Rubin D. Maximum likelihood from incomplete data via the EM algorithm (with discussion) 1977;B: 39:1–38.
- Farquhar J, Hill NJ. Interactions between pre-processing and classification methods for event-related-potential classification: best-practice guidelines for brain-computer interfacing. *Neuroinformatics* 2013;11:175–92.
- Fischer AG, Ullsperger M. Real and fictive outcomes are processed differently but converge on a common adaptive mechanism. *Neuron* 2013;79:1243–55.
- Frank M, Woroch B, Curran T. Error-related negativity predicts reinforcement learning and conflict biases 2005;47:495–501.
- Gluth S, Rieskamp J, Buchel C. Classic EEG motor potentials track the emergence of value-based decisions. *Neuroimage* 2013;79:394–403.
- Gluth S, Rieskamp J, Buchel C. Deciding when to decide: time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *J Neurosci* 2012;32:10686–98.
- Gold I, Shadlen M. The neural basis of decision making. *Annu Rev Neurosci* 2007;30:535–74.
- Green D, Swets J. Signal detection theory and psychophysics. New York: John Wiley and Sons Inc; 1966.
- Heekeren H, Marrett S, Bandettini P, Ungerleider L. A general mechanism for perceptual decision-making in the human brain. *Nature* 2004;431:859–62.
- Hunt L, Kolling N, Soltani A, Woolrich M, Rushworth M, Behrens T. Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci* 2012;15:470–6.
- Kelly SP, O'Connell RG. Internal and external influences on the rate of sensory evidence accumulation in the human brain. *J Neurosci* 2013;33:19434–41.
- Kim J, Shadlen M. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature* 1999;2:176–85.
- Kittler J, Hatief M, Duin RJM. On combining classifiers. *IEEE Trans Pattern Anal Mach Intell* 1998;20:226–39.
- Koenig T, Kottlow M, Stein M, Melie-Garcia L. Ragu: a free tool for the analysis of EEG and MEG event-related scalp field data using global randomization statistics. *Comput Intell Neurosci* 2011;2011:938925.
- Lehmann D, Skrandies W. Reference-free identification of components of checkerboard-evoked multichannel potential fields. *Electroencephalogr Clin Neurophysiol* 1980;48:609–21.
- Millan JD, Rupp R, Muller-Putz GR, Murray-Smith R, Giugliemma C, Tangermann M, et al., Cincotti F, Leeb R, Muller KR. Combining brain-computer interfaces and assistive technologies: state-of-the-art and challenges. *Front Neurosci* 2010;4.
- Minati L, Grisoli M, Franceschetti S, Epifani F, Granvillano A, Medford N, et al., Piacentini S. Neural signatures of economic parameters during decision-making: a functional MRI (fMRI), electroencephalography (EEG) and autonomic monitoring study. *Brain Topogr* 2012;25:73–96.
- Murray M, Brunet D, Michel C. Topographic ERP analyses: a step-by-step tutorial review. *Brain Topogr* 2008;20:249–64.
- O'Connell RG, Dockree PM, Kelly SP. A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat Neurosci* 2012;15:1729–35.
- Parra L, Alvino C, Tang A, Pearlmuter B, Yeung N, Osman A, et al. Linear spatial integration for single-trial detection in encephalography. *Neuroimage* 2002;17:223–30.
- Parra LC, Spence CD, Gerson AD, Sajda P. Recipes for the linear analysis of EEG. *Neuroimage* 2005;28:326–41.
- Pedroni A, Langer N, Koenig T, Allemand M, Jancke L. Electroencephalographic topography measures of experienced utility. *J Neurosci* 2011;31:10474–80.
- Pernet CR, Sajda P, Rousselet GA. Single-trial analyses: why bother? *Front Psychol* 2011;2:322.
- Perrin F, Pernier J, Bertrand O, Giard M, Echallier J. Mapping of scalp potentials by surface spline interpolation. *Electroencephalogr Clin Neurophysiol* 1987;66:75–81.
- Philastrides MG, Sajda P. Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cereb Cortex* 2006;16:509–18.
- Philastrides M, Biele G, Heekeren H. A mechanistic account of value computation in the human brain. *Proc Natl Acad Sci U S A* 2010;107:9430–5.
- Raftery A. Bayesian model selection in social research. *Sociological methodology*; 1995.
- Ratcliff R, Philastrides MG, Sajda P. Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proc Natl Acad Sci U S A* 2009;106:6539–44.
- Ratcliff R. A theory of memory retrieval. *Psychol Rev* 1978;85:59–108.
- Rieger JW, Reichert C, Gegenfurtner KR, Noesselt T, Braun C, Heinze HJ, et al. Predicting the recognition of natural scenes from single trial MEG recordings of brain activity. *Neuroimage* 2008;42:1056–68.
- Salvaris M, Sepulveda F. Visual modifications on the P300 speller BCI paradigm. *J Neural Eng* 2009;6:046011.
- Schulz E, Zherdin A, Tiemann L, Plant C, Ploner M. Decoding an individual's sensitivity to pain from the multivariate analysis of EEG data. *Cereb Cortex* 2012;22:1118–23.
- Shadlen MN, Newsome WT. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 2001;86:1916–36.
- Taghizadeh-Sarabi M, Daliri MR, Niksirat KS. Decoding objects of basic categories from electroencephalographic signals using wavelet transform and support vector machines. *Brain Topogr* 2014.
- Talmi D, Fuentemilla L, Litvak V, Duzel E, Dolan RJ. An MEG signature corresponding to an axiomatic model of reward prediction error. *Neuroimage* 2012;59:635–45.
- Tax D, van Breukelen M, Duin R, Kittler J. Combining multiple classifiers by averaging or by multiplying? *Pattern Recognit* 2000;33:1475–85.
- Tzovara A, Murray M, Bourdaud N, Chavarriaga R, Millán JDR, Lucia MD. The timing of exploratory decision-making revealed by single-trial topographic EEG analyses. *Neuroimage* 2012a;60:1959–69.
- Tzovara A, Murray M, Plomp G, Herzog M, Michel C, Lucia MD. Decoding stimulus-related information from single-trial EEG responses based on voltage topographies. *Pattern Recognit* 2012b;45:2109–22.
- van Vugt MK, Simen P, Nystrom LE, Holmes P, Cohen JD. EEG oscillations reveal neural correlates of evidence accumulation. *Front Neurosci* 2012;6:106.
- Vandekerckhove J, Tuerlinckx F. Fitting the Ratcliff diffusion model to experimental data. *Psychon Bull Rev* 2007;14:1011–26.
- Wetzels R, Matzke D, Lee MD, Rouder JN, Iverson GJ, Wagenmakers EJ. Statistical evidence in experimental psychology: an empirical comparison using 855 t tests. *Perspect Psychol Sci* 2011;6:291–8.
- Wyart V, de Gardelle V, Scholl J, Summerfield C. Rhythmic fluctuations in evidence accumulation during decision making in the human brain. *Neuron* 2012;76:847–58.
- Yeung N, Holroyd C, Cohen J. ERP correlates of feedback and reward processing in the presence and absence of response choice. *Cereb Cortex* 2005;15:535–44.
- Zhang D, Gu R, Broster LS, Jiang Y, Luo W, Zhang J, et al. Linking brain electrical signals elicited by current outcomes with future risk decision-making. *Front Behav Neurosci* 2014;8:84.