

Single-trial decoding of auditory novelty responses facilitates the detection of residual consciousness



J.R. King^{a,b,c,*}, F. Faugeras^{c,d}, A. Gramfort^{b,e,f}, A. Schurger^{a,b}, I. El Karoui^c, J.D. Sitt^{a,b,c}, B. Rohaut^{c,g}, C. Wacongne^{a,b}, E. Labyt^{a,b}, T. Bekinschtein^h, L. Cohen^{c,d,i}, L. Naccache^{c,d,h}, S. Dehaene^{a,b,j,k}

^a Cognitive Neuroimaging Unit, Institut National de la Santé et de la Recherche Médicale, U992, F-91191 Gif/Yvette, France

^b NeuroSpin Center, Institute of Biomedicine Commissariat à l'Energie Atomique, F-91191 Gif/Yvette, France

^c Institut du Cerveau et de la Moelle Épineuse Research Center, Institut National de la Santé et de la Recherche Médicale, U975 Paris, France

^d AP-HP, Groupe hospitalier Pitié-Salpêtrière, Department of Neurophysiology, Paris, France

^e Institut Mines-Télécom, Télécom ParisTech, CNRS LTCI, F-75014 Paris, France

^f INRIA Saclay – Ile de France, PARIETAL INRIA, Neurospin, CEA Saclay, Bâtiment 145, 91191 Gif-sur-Yvette Cedex, France

^g AP-HP, Groupe hospitalier Pitié-Salpêtrière, Department of Neurology, Intensive Care Unit, Paris, France

^h Cognition and Brain Sciences Unit, Medical Research Council, Cambridge, UK

ⁱ Faculté de Médecine Pitié-Salpêtrière, Université Paris 6, Paris, France

^j Université Paris 11, Orsay, France

^k Collège de France, F-75005 Paris, France

ARTICLE INFO

Article history:

Accepted 4 July 2013

Available online 13 July 2013

Keywords:

Vegetative state

Unresponsive wakefulness syndrome

Minimally conscious state

Multivariate pattern classifier

Decoding

EEG

MEG

ABSTRACT

Detecting residual consciousness in unresponsive patients is a major clinical concern and a challenge for theoretical neuroscience. To tackle this issue, we recently designed a paradigm that dissociates two electro-encephalographic (EEG) responses to auditory novelty. Whereas a local change in pitch automatically elicits a mismatch negativity (MMN), a change in global sound sequence leads to a late P300b response. The latter component is thought to be present only when subjects consciously perceive the global novelty. Unfortunately, it can be difficult to detect because individual variability is high, especially in clinical recordings. Here, we show that multivariate pattern classifiers can extract subject-specific EEG patterns and predict single-trial local or global novelty responses. We first validate our method with 38 high-density EEG, MEG and intracranial EEG recordings. We empirically demonstrate that our approach circumvents the issues associated with multiple comparisons and individual variability while improving the statistics. Moreover, we confirm in control subjects that local responses are robust to distraction whereas global responses depend on attention. We then investigate 104 vegetative state (VS), minimally conscious state (MCS) and conscious state (CS) patients recorded with high-density EEG. For the local response, the proportion of significant decoding scores ($M = 60\%$) does not vary with the state of consciousness. By contrast, for the global response, only 14% of the VS patients' EEG recordings presented a significant effect, compared to 31% in MCS patients' and 52% in CS patients'. In conclusion, single-trial multivariate decoding of novelty responses provides valuable information in non-communicating patients and paves the way towards real-time monitoring of the state of consciousness.

© 2013 Elsevier Inc. All rights reserved.

Introduction

Despite a recent flurry of experimental discoveries, the neuronal mechanisms that support conscious perception remain a major challenge for neuroscience. Numerous studies have associated conscious perception with macroscopic neurophysiological phenomena such as synchronous activity across distant cortical regions and sustained fronto-parietal activations (Fahrenfort et al., 2007, 2012; Gaillard et al., 2009; Melloni et al., 2007; Sergent et al., 2005) which are thought

to reflect global information integration (Dehaene and Changeux, 2011; Dehaene and Naccache, 2001; Fisch et al., 2009; Lamme, 2006; Melloni et al., 2011; Rees et al., 2002; Seth et al., 2011; Tononi and Edelman, 1998; Tononi and Sporns, 2003). However, identifying the neuronal signatures of conscious processing is not just a theoretical exercise. Every year, severe brain injuries lead thousands of patients to lose their communication abilities and fall into a variety of clinical conditions ranging from coma, to vegetative state (VS), minimally conscious state (MCS) or conscious but paralyzed patients (locked-in syndrome). Clinically, separating disorders of consciousness (DOC) from communication deficits can be difficult. In particular, VS (also known as “unresponsive wakefulness syndrome”, Laureys et al. 2010) present moments of arousal, during which they open their eyes and produce

* Corresponding author at: Cognitive Neuroimaging Unit, Institut National de la Santé et de la Recherche Médicale, U992, F-91191 Gif/Yvette, France.

E-mail address: jeanremi.king@gmail.com (J.R. King).

complex behavioral reflexes. Yet, they show no clear signs of intentional behavior, even after careful clinical examination performed by experienced teams (Bruno et al., 2011a). “Minimally conscious state” (MCS) patients, present some intentional behaviors but seem unable to establish any long-lasting functional communication (Giacino et al., 2002).

To facilitate clinical diagnosis, brain imaging techniques may play an important role (Laureys and Schiff, 2011; Laureys et al., 2004; Monti et al., 2010; Owen, 2008; Owen et al., 2006). By directly detecting the neural activity associated with conscious processing, they could be used to circumvent communication deficits and thus provide crucial information for the diagnosis of these patients. The Local–Global paradigm (Bekinschtein et al., 2009) was designed for this purpose (Fig. 1). This experimental setup allows the isolation of two event related potentials (ERPs) elicited by two types of auditory novelty. First, a change in pitch within a five-sound sequence (hereafter referred to as local deviancy) typically leads to a frontal mismatch negativity (MMN) ~150 ms after stimulus onset. Second, a change in auditory sequence in a fixed global context generates a late P300b response over centro-posterior electrodes. Crucially, these auditory changes can be arranged to create a 2 × 2 design in which local deviancy and global deviancy are orthogonally manipulated (Bekinschtein et al., 2009).

There is now growing evidence that the MMN reflects a prediction error signal elicited whenever the incoming sound does not fit with a prediction constructed on the basis of previous local auditory regularities (Garrido et al., 2007, 2008; Näätänen et al., 1978, 2010; Wacongne et al., 2011, 2012; Winkler, 2007). Moreover, manipulations of attention, sleep and anesthesia show that the MMN may persist even in unconscious states (Atienza and Cantero, 1997; Bekinschtein et al., 2009; Brázdil et al., 2001; Garrido et al., 2008; Heinke et al., 2004; Muller-Gass et al., 2007; Näätänen et al., 2010; Rohaut et al., 2009; Tzovara et al., 2013). By contrast, the P300b is thought to reflect a higher-order violation of subjects’ expectations of a given rule, constructed over a longer time period and has thus been closely linked to working memory (Goldstein et al., 2002; Polich, 2007) and conscious access (Dehaene and Changeux, 2011; Dehaene et al., 2006). Converging lines of evidence suggest that the dissociation between these two neural signatures could discriminate patients in vegetative state (VS) from those in conscious (CS) or minimally conscious states (MCS) (Faugeras et al., 2011, 2012; Fischer et al., 2010; Naccache et al., 2005; Rohaut et al., 2009; Wijnen et al., 2007).

However, isolating these event-related responses in single patients can unfortunately be difficult for several reasons. First, clinical recordings often present a low signal-to-noise ratio (SNR) because of numerous physiological (movements, eye blinks...) and environmental artifacts (presence of auditory noise, no Faraday cage...). Moreover, patients often present severe brain and even skull damages which can alter the

scalp electrical projections of their cortical activity. This topographical variability can be made worse by temporal delays and inter-trial variability caused by processing impairments or white matter damage (Newcombe et al., 2010; Tshibanda et al., 2009). In other words, unlike control recordings, a patient’s MMN and P300b may not be optimally observed over the frontal and parietal channels at ~150 and ~350 ms respectively. While more liberal analyses testing a greater number of EEG channels and time samples could be implemented, correction for multiple comparisons would largely diminish either sensitivity or confidence in the presence of a given brain response.

To overcome these common electrophysiological issues, we evaluate in the present research the potential of a single-trial multivariate pattern (MVP) analysis. We implemented, separately for each subject, a MVP classifier that aims at maximally extracting information from each trial by combining evidence from multiple EEG channels and multiple time-samples. After training on an independent dataset, the classifier estimates the probability that each trial contains a local or a global response to auditory novelty. This prediction can be compared to trials’ effective classes. Classification scores can thus indicate whether a given subject is able to detect the corresponding type of novelty.

To optimize the detection of single-trial local and global novelty responses, the present research followed a strict logic. First, to optimize our methods, we tested them with EEG, MEG and intracranial EEG recordings acquired from control subjects. Second, we then applied it to 158 high-density EEG recordings from 104 distinct patients whose state of consciousness (VS, MCS, or CS) was assessed immediately prior the experiment. We consider three successive questions: (1) What level of accuracy can be achieved from each type of recordings and does our method present genuine improvements as compared to traditional analyses? (2) Can decoders be formed to generalize the detection of novelty from one experimental context to another? (3) Is our method sensitive enough to be applied to the detection of residual novelty processing in non- or poorly communicating patients?

Methods

Procedure, material & apparatus

The data analyzed here come from four different experimental settings, using either scalp EEG, MEG, or intracranial EEG, which together enables the direct comparison of the utility of each approach with regard to single trial decoding. Events Related Potentials (ERPs) and Events Related Fields (ERFs) have been partially reported elsewhere (Bekinschtein et al., 2009; Faugeras et al., 2012; Wacongne et al., 2011). All experiments were approved by the relevant regional ethical committees (Comité pour la Protection des Personnes Pitié-Salpêtrière and Bicêtre hospitals). Healthy volunteers received a financial compensation for their participation. Unless specified otherwise, the procedure used in our experiments exactly followed the Local–Global paradigm (Bekinschtein et al., 2009) which enables the comparison of effects engendered by physically identical but contextually different auditory stimuli. In the standard Local–Global design, subjects are asked to count the global deviant trials and report this number at the end of each block, in order to ensure that they pay attention to the task.

The auditory stimuli were 50 ms-duration sounds composed of 3 sinusoidal tones (either 350, 700, and 1400 Hz, hereafter sound A; or 500 Hz, 1000 Hz, and 2000 Hz, hereafter sound B), with 7-ms rise and 7-ms fall times. Sequences were composed of five stimuli presented at a Stimulus Onset Asynchrony (SOA) of 150 ms, and were separated by a variable silent interval of 1350 to 1650 ms (50 ms steps). The sequences could comprise five identical tones (xxxxx) or four identical tones followed by a distinct one (xxxxy, where x can be sound A or sound B and Y the other sound). Following the original design, in a given block, 80% of trials consisted in one type of sequence (e.g. aaaaB) and 20% of trials were global deviants (aaaaa in this example), pseudo randomly distributed at least one and at

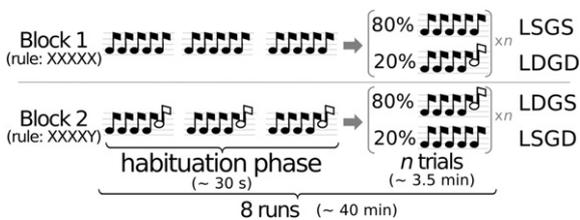


Fig. 1. The Local–Global paradigm. The Local–Global paradigm (Bekinschtein et al., 2009) is an auditory odd-ball experimental setup that implicitly tests subjects on their ability to detect two orthogonal types of auditory novelty. Each trial is composed of five successive tones (SOA = 150 ms). The first four sounds are always identical. Local-deviant (LD) trials differ from local-standard trials (LS) because their fifth sound deviates in pitch. Global-deviant (GD) trials correspond to the presentation of a sequence of five sounds which is rare in a given block, compared to the frequent “global standard” sequence (GS). Both local and global novelties depend entirely on the fifth sound, which therefore serves as the origin of time scales in all subsequent graphs.

most six global-standard trials apart (Fig. 1). In all experiments, trials immediately following a global deviant were removed from the analyses. Each block started with a 30 s habituation phase during which the frequent sound sequences were repeatedly presented to establish the global regularity, before the first infrequent stimulus was heard. In Experiments 1, 3–5, sounds were presented via headphones with an intensity of 70 dB, using E-prime v1.2 (Psychology Software Tools Inc.). Trials from the habituation phase were not included in the analyses. In Experiment 2, sounds were directly presented from the computer's speakers because the intracranial recording apparatus was incompatible with headphones.

Experiment 1: Counting task (EEG)

In the first experiment (partially reported in (Faugeras et al., 2012)), ten healthy adults (Age: $M = 23.0$ years old, $SD = 0.67$ years, 3 females) performed the standard Local–Global paradigm, while EEG was continuously recorded using a 256-channel EEG geodesic net (EGI) sampled at 250 Hz.

Each subject was recorded for approximately 45 min, comprising 8 blocks of 3–4 min duration. Each block was interleaved with resting periods of a few minutes. In each block, subjects were instructed to mentally count the global deviant trials. This experiment will therefore be referred to as 'counting EEG'.

Bad sensors, defined as those showing no signal at all, constant white noise, or presenting intermittent signals, were interpolated. Trials during which more than 20% of the sensors were bad, the EEG voltages exceeded $\pm 150 \mu\text{V}$, EEG transients exceeded $\pm 100 \mu\text{V}$ or electrooculogram activity exceeded $\pm 80 \mu\text{V}$ were excluded from the analyses. All signals were digitally low-pass filtered at 40 Hz, and referenced with a common average. Trials were then segmented from -800 ms to 736 ms after the onset of the fifth sound, and were baseline corrected over a 200 ms window before the onset of the first of the five sounds. All EEG processing stages were performed in the EGI Waveform Tools Package and with the Fieldtrip toolbox (Oostenveld et al., 2011) and MATLAB 2009b.

Experiment 2: Counting task (iEEG)

In the second experiment, nine patients (Age: $M = 33$ years old, $SD = 11$ years, 5 women) suffering from drug-resistant epilepsy, and who had consequently undergone electrode implantation for pre-surgical purposes, performed a standard Local–Global paradigm (two patients were already reported in (Bekinschtein et al., 2009)) in 8 consecutive blocks. It should be noted that, although epileptic patients cannot be considered as healthy controls, their attention and their state of consciousness were relatively comparable to those of healthy subjects. Therefore, their intracranial EEG recordings may shed additional light on the EEG and MEG responses observed in healthy subjects. On average, patients had 56 intracranial electrodes (iEEG) sampled at 400 Hz or 1024 Hz (depending on the system) placed in various cortical areas including the temporal, the occipital, and the frontal lobes. Electrode locations are reported in Supplementary Fig. 2. The task was identical to Experiment 1s.

After removing channels showing inter-ictal activity, analyses were performed by pooling over the 401 channels recorded across all subjects. All signals were digitally low-pass filtered at 40 Hz and down-sampled to 256 Hz. Trials were then segmented from -800 ms to 700 ms after the critical stimulus onset, and were corrected for baseline over a 200 ms window before the onset of the first of the five sounds. All EEG processing stages were performed with MATLAB 2009b and the Fieldtrip toolbox (Oostenveld et al., 2011).

Experiment 3: Attentive task (MEG)

In the third experiment (partially reported in Wacongne et al., 2011), ten healthy adults (Age: $M = 25$ years old, $SD = 4.7$ years, 5 females) performed a modified version of the Local–Global paradigm (see below) while their brain activity was measured with MEG (Elekta

Neuromag® MEG system, Helsinki, Finland, comprising 204 planar gradiometers and 102 magnetometers in a helmet-shaped array) and EEG (built-in 64 electrodes system). Scalp EEG electrodes were not analyzed in the present study. Data were sampled at 1 KHz with on-line analog low-pass filtering at 330 Hz, and on-line analog high-pass filtering at 0.1 Hz. The head position with respect to the sensor array was determined by four head position indicator coils attached to the scalp. The locations of the coils and EEG electrode positions were digitized with respect to three anatomical landmarks (nasion and preauricular points) with a 3D digitizer (Polhemus Isotrak system®). Then, head position with respect to the device origin was acquired before each block of MEG/EEG recording.

Each recording session lasted 1 h, comprising 14 blocks of 3–4 min duration with resting periods between each block. Subjects were asked to keep their eyes opened and to avoid eye movements by focusing on a fixation cross displayed in the center of the screen. Subjects were instructed to pay attention to the auditory stimuli.

The MEG task differed from Experiment 1 in the following respects. First, although the Local Standard and Local Deviant sequences were mainly identical, 10% of the trials were omission trials composed of only four sounds. Furthermore, in two block blocks, the frequent auditory sequence was also made of only four sounds. These conditions were applied in order to test the brain responses to expected and unexpected omissions. All trials made of only four sounds, or in a block where the frequent sequences were composed of only four sounds were excluded from the present analyses, but are reported in detail in Wacongne et al. (2011). Second, subjects performed more trials than in the previous experiment (780 trials instead of 500). Third, subjects were not asked to count the number of globally deviant trials, but were only required to pay attention to the auditory stimuli. This 'attentive MEG' experiment aimed at demonstrating that the previously identified neurophysiological signatures associated with local and global deviant trials were independent of the counting task. This control can thus validate the applicability of the Local–Global paradigm in a clinical setup, in which patients may not be able to perform complex instructions such as counting global deviant trials. At the end of the recording, a list of questions was submitted to the subject to check that they had detected the various auditory regularities.

Signal space separation (SSS, Taulu et al., 2004) was applied to suppress unwanted magnetic interferences (e.g. outside disturbances, limb movements), to interpolate noisy MEG sensors and to realign MEG data into a subject-specific head position. This reference head position was determined from all head position measurements done at the beginning of each recording session. This data transformation helps the direct comparisons of MEG data between conditions and blocks.

Except if explained otherwise, eye blinks and cardiac artifact were corrected separately for each type of channel (gradiometer and magnetometers) by decomposing the average artifacts into principal components, and regressing out those principal components from the continuous recording – a technique known as signal space projection (SSP). Noisy MEG sensors were removed with Maxfilter in the SSS preprocessing step. All signals were digitally low-pass filtered at 40 Hz and down-sampled to 256 Hz. Trials were then segmented from -800 ms to 700 ms after the critical stimulus onset, and were corrected for baseline over a 200 ms window before the onset of the first of the five sounds. Segmentation was done with Fieldtrip (Oostenveld et al., 2011). Trials with more than 20% of bad sensors were rejected.

Experiment 4: Distracting task (EEG)

In the fourth experiment (partially reported in Bekinschtein et al., 2009), 9 subjects (6 females aged between 21 and 33 years old) were actively distracted by a continuous speeded visual detection task simultaneous with the presentation of the Local–Global auditory stimuli, while EEG was continuously recorded at 250 Hz using a

256-channel EEG geodesic net (EGI) referenced to the vertex. Participants were instructed to detect a visual target in a rapid stream of successive letters presented at the fovea and were explicitly asked to neglect the unrelated auditory stimuli. The visual stimuli were twelve $1^\circ \times 1^\circ$ colored upper or lower case letters, and were maximally presented for 1000 ms, using E-prime v1.1 (Psychology Software Tools Inc.). To check that subjects did not consciously perceive the global structure of the auditory stimuli, they were asked at the end of the experiment whether they had perceived any global regularity or novelty. EEG preprocessing was identical to Experiment 1.

Experiment 5: Levels of consciousness in patients (EEG)

One hundred and fifty eight (158) EEG recordings were acquired from 104 distinct patients (Age = 48 years old, SD = 17 years, 35 females) while they performed the Local–Global paradigm. The first 65 recordings were partially reported in [Faugeras et al. \(2012\)](#). Immediately before the EEG recording, each patient was carefully examined by a trained neurologist (FF, BR, LN) and was assessed with the French version of the Coma Recovery Scale – Revised (CRS-R, [Schnakers et al., 2008](#)). This scale allows a clinical categorization of each patient into one of three states of consciousness. The vegetative state (VS) refers to awake patients who fail to perform very simple tasks such as visual fixation and localization to noxious stimulations. The minimally conscious state (MCS) refers to responsive but non-communicating patients ([Giacino et al., 2002](#)). Recently, [Bruno et al. \(2011b\)](#) proposed to sub-categorize this category into MCS+ and MCS– classes to further separate MCS patients who can follow commands from those who cannot. As explained in the supplementary materials, MCS– can for instance track a visual stimulus while remaining unable to demonstrate intentional communication. Finally, conscious state (CS) refers to patients who present a functional communication or a functional capacity to intentionally use objects. With the exception of one recording (#140, see Supplementary Table 1), the formal classification of consciousness states is formally based on the CRS-R. In total, 70 EEG recordings were acquired from VS patients, 65 from MCS patients (in which 25 were acquired from MCS– patients) and 23 from CS patients. As can be seen in Supplementary Table 1, some patients ($n = 29$) were recorded several times for clinical purposes. Most of them ($n = 18/29$) were recorded in at least two distinct state of consciousness. Patients' etiologies was heterogeneous: 30 had suffered from an anoxia, 33 from a stroke, 20 from a traumatic brain injury (TBI) and 21 suffered from another etiology. On average, EEG recordings were acquired 178 days after patients' injury. All patients had been without sedation for at least 24 h prior to the recording session. All clinical details are presented in Supplementary Table 1.

After clinical examination, each participant was asked to perform a task identical to Experiment 1 (count the global deviants). Patients were verbally stimulated between each block (~4 min) to ensure stable arousal, and instructions were explicitly repeated by the experimenter before each run. Although we present results demonstrating that global effects are independent of the instruction to count, we reasoned that such instructions could help patients pay attention to the auditory stimulations.

Recording high-density scalp ERPs from non-communicating patients in the intensive care unit or a similar environment is very challenging for technical reasons. First, the electro-magnetic environment is noisy, and patients were not recorded in a shielded room but at bedside. Second, many patients presented physiological artifacts such as EMG, eye-movements and blinks, or other involuntary movements. Therefore, it is particularly important to systematically evaluate the technical quality of data before statistical analysis. Recordings including at least one block with more than 50% of rejected trials were discarded from further analyses in order to avoid possible biases across experimental conditions. EEG preprocessing was identical to Experiment 1.

Analyses

Classification

In the following experiments, we aimed at evaluating, for each subject, the number of trials that could be accurately classified in two distinct and orthogonal types of classes: (1) as local standard (LS, i.e. an xxxxx sequence) versus local deviant (LD, i.e. an xxxxY sequence, 'local classification'); (2) as global standard (GS, frequent sequence) versus global deviant (GD, rare sequence in a given block, 'global classification', see [Fig. 1](#)). The Local Standard class corresponds to the union of LSGS and LSGD trials, and the Local Deviant class corresponds to the union of LDGS and LDGD trials. Similarly, the Global Standard class corresponds to LSGS and LDGS trials, and the Global Deviant class corresponds to LSGD and LDGD trials. These analyses contrast trials evenly distributed across blocks and are therefore less likely to be contaminated by block-design pitfalls ([Goldfine et al., 2012](#); [Lemm et al., 2011](#)).

The decoding steps are summarized in Supplementary Fig. S6. A ten-fold stratified cross-validation was implemented for each within-subject analysis. Stratified cross-validation balances the proportion of each class across K folds ($K = 10$) in order to maximize the classifiers' ability to generalize to unknown data. Stratified cross-validation is thus particularly relevant for problem with unbalanced class frequencies. Cross-validation consists in repeatedly applying a series of computations fitted to a fraction of the dataset ($1 - 1/K$ percent of the trials, called the training set), and then evaluate the results on the remaining fraction ($1/K$ percent of the trials, called the test set). Within each fold, we estimated the mean (μ_{train}) and the standard deviation (σ_{train}) of each sample of each sensor across all training trials. The training set X_{train} and the test set (X_{test}) were then normalized ($X_{\text{train}} = (X_{\text{train}} - \mu_{\text{train}}) / \sigma_{\text{train}}$; $X_{\text{test}} = (X_{\text{test}} - \mu_{\text{train}}) / \sigma_{\text{train}}$). As detailed below, in some analyses a univariate feature selection was fitted on the training set and subsequently applied to both the training and the test set. Finally, a support vector classifier (SVC) with a linear kernel ([Chang and Lin, 2001](#)) was supplemented with a continuous output method providing, for each trial, an estimate of the probability of belonging to a given class ([Platt, 1999](#)). Amongst the several advantages of this continuous method, we note that it allows across-trial rank statistics and hence avoids the computationally expensive permutation analysis generally required with discrete and/or imbalanced classifiers (see [Fig. S5](#)). After fitting the SVC on the training set, classification scores were estimated on the independent test set. All preprocessing steps posterior to trials segmentation (feature selection, normalization, parameter selection etc) were fitted, within each cross-validation loop, on the training set only, and are thus immune to circular analysis issues ([Lemm et al., 2011](#)). Cross-validation was constructed from chronologically shuffled data which minimizes effects of non-stationarity that could have been observed between the beginning and the end of the recording session for instance ([Lemm et al., 2011](#)). To further minimize the possibility of a trial ordering confound, the cross-validation was applied four times for each patient. Each of these cross-validations was based on a different chronologically shuffled dataset thus implying different folds too. This step did not however fundamentally change the obtained results.

Classification scores were estimated with an empirical receiver-operative curve (ROC) analysis applied on trials' predicted probabilities. The ROC analysis is a standard 2-class (positive and negative) non-parametric statistical method which allows the estimation of the effect size of a Wilcoxon/Mann–Whitney U Test ([Mason and Graham, 2002](#)). It is based on the plotting of true positive rate as a function of false positive rate. The result of this function can be summarized by the area under its curve (AUC). An AUC of 50% implies that true positive predictions (e.g. trial predicted deviant and is deviant) and false positive predictions (e.g. trial predicted as deviant but is standard) are, on average, equally probable; an AUC of 100% indicates a perfect positive prediction with no false positive. Amongst the advantages of the ROC analysis, we note that, unlike mean accuracy,

it is robust to imbalanced problems and, as a non-parametric analysis, does not make any hypothesis about the distribution underlying the data. The reported AUCs correspond to the averaged within-subject effect size: the ability of the classifier to discriminate standard from deviant trials in a given subject. Note that the average within-subject AUC can be relatively low and yet remains significant across subjects. Such cases suggest that while the analysis robustly identifies a significant effect across an entire EEG session, it is unlikely to be useful to online decoding of single-trials.

Generalization across context

We implemented a set of generalization analyses which aimed at testing the invariance of the neurophysiological signatures of local and global novelty detection. These generalization analyses consisted in training the classifier using data from a fixed context, and subsequently testing it in a different context. Classifiers were for instance trained to discriminate local standard trials from local deviant trials using only global standard trials (LSGS *versus* LDGS). The discriminative hyper-plane (w) found was subsequently applied to trials from a different global context, in this case the global deviants (LSGD *versus* LDGD). This approach was also adapted to the contextual generalization of global effects. In this case, the classifier was first trained, for instance, to discriminate global effects in a fixed local context (LSGS *versus* LSGD), and then tested on a different local context (LDGS *versus* LDGD, see Figs. 1 & 4). This procedure was systematically applied in a symmetrical manner: a classifier was either trained in standard contexts and tested in deviant ones, or trained in deviant contexts and tested in standard ones, and the results were then averaged over both directions of generalization. This procedure thus aimed to identify the neural signatures of novelty processing that are robust and generic to different contexts, with the added difficulty of being based on only half of the training trials.

Time-windows of interest

We applied several multivariate pattern classifiers on different temporal regions of interest (Fig. S5). First, we repeatedly tested a classifier using all channels on a sliding window of a single time sample (~4 ms). Second, we implemented a set of classifiers using multiple time samples. These classifiers combined either all or a sub-selection of the time samples following the onset of the last sound ($t = [0,736]$ ms): the early time-window refers to $t = [0, 367]$ ms, and the late time window refers to $t = [367, 700]$ ms. As a control, supplementary analyses also tested the ability of classifiers to extract the electrophysiological information on a pre-stimulus window ($[-367, 0]$ ms relative the onset of the fifth sound). Finally, the last approach consisted in averaging brain signals within a given time windows (Fig. S3). The method, rationale and results are presented in the supplementary materials.

Support Vector Machine (SVM)

Each trial was transformed into a p -dimensional vector x , in which each coordinate (“attribute”) corresponds to a single data sample at a given sensor ($p = n_{\text{sensors}} \cdot n_{\text{samples}}$). The entire dataset can hence be represented as a matrix X in which each row i corresponds to one trial x_i , and each column corresponds to an attribute. Moreover, trial classes (standard or deviant) can also be represented as a binary vector y , in which standard trials are labeled as 1, and deviant trials as -1 . As for all linear classification analyses, the aim is to find a hyperplane (materialized by a p -dimensional vector of weights w) that discriminates the two classes (y_{standard} and y_{deviant}). For a new trial x_{new} , the sign of the dot product $x_{\text{new}} \cdot w$ is then generally used to predict its class. In the present case, a cumulative probability distribution function was fitted to the training set using Platt’s method (Platt, 1999). The signed distance computed from the dot product $x_{\text{new}} \cdot w$ could thus be transformed into a continuous value bounded between 0 and 1. This

value is directly representative of the probability of belonging to one of the two classes. Note that rank statistics are not affected by the transformation of the signed distance between a given trial and w by a monotonous function, such as a cumulative distribution function. When the number of trials is small compared to the number of attributes (here up to 56,000), there are infinitely many w that can fit the training data equally well. To limit this issue, we implemented within the cross-validation, a dimensionality reduction based on a univariate feature-selection step that selects a subset of attributes prior to the classification analysis. Feature selection was performed with an ANOVA nested within the cross-validation step. Various levels of the ANOVA p value threshold were explored in Experiment 1 (1%, 5%, 10%, 30%, 50%, 100%), and on this basis a fixed value of 10% was chosen for all other experiments. No feature selection was used for the decoding of each time point as dimensionality was relatively small (e.g. from 40 channels to 306 channels depending on the type of recording). SVC’s regularization parameter was calibrated by nested cross-validation in Experiment 1 (.01, .1, .3, .5, 1, 2) and remained arbitrarily fixed (1) for all other experiments as its values did not dramatically affect classification scores. Finally, sample weights were applied in proportion of the trial classes (LSGS, LSGD, LDGS, LDGD) so that each of the category would equally contribute to the definition of w , and this independently of the contrast of interest (local or global). This step is known to contribute to the minimization of imbalanced dataset artifacts and thus maximizes the number of trials that can be considered in the training process. In practice this step did not lead to a strong improvement of classification performance. All multivariate analyses were performed with the Scikit-learn toolbox (Pedregosa et al., 2011).

As a control, in supplementary analyses, we applied the same approach after randomly shuffling the label (y) of each trial and subsequently training the classifier with this incorrect class labeling. Ideally, this method would be applied 10,000 times per subject to allow the estimation of the null distribution of each subject. However, this approach would be too computationally expensive with the present dataset, because the classification of a single EEG recording session takes approximately 30 min of computer time. We thus applied shuffling only once per EEG recording, and thus only report across-subjects statistics.

Statistics

ROC analyses and AUCs are methods to estimate the size of a given effect, and in our case, the classifiers’ ability to discriminate two types of trials. To test for statistical significance within subjects, we performed Mann–Whitney U tests across trials. Because the classifier attributes a continuous estimate to each trial (its predicted probability of belonging to a given class C), we can indeed efficiently compare those predicted probabilities across the two levels of the trials’ true classes. For instance, for each trial, the classifier outputs a predicted probability of belonging to the standard class (S). Note that the probability of belonging to the other, deviant class (D) can be calculated as $P(D) = 1 - P(S)$. We can then compare the predicted probability of belonging to S depending on whether the trials truly are standard ($y = S$) or not ($y = D$) and thus apply a traditional Mann–Whitney U Test between $P(S|y = S)$ and $P(S|y = D)$.

Similarly, across-subjects statistics were performed using Wilcoxon Signed Rank Tests based on the mean predicted probability conditional of trials’ true classes. For each subject, the predicted probability of belonging to the standard class (S) is averaged across standard trials ($y = S$) and, separately, over deviant trials ($y = D$), yielding $P(S|y = S)$ and $P(S|y = D)$. For each subject i , we can thus compute the sum of positive ranks and perform a Wilcoxon test. A correction for multiple comparisons, either across classifiers or across subjects, was performed using the standard False Discovery Rate (FDR) correction, and is hereafter referred to as p_{FDR} . Statistical analyses were performed with R and MATLAB 2009b.

Results

The Local–Global paradigm enables the isolation of two types of neurophysiological activity: either the response to a change in pitch within a five-sound sequence (local effect), or the response to a rare auditory sequence within a block dominated by another frequently repeated sequence (global effect, Fig. 1).

Topographical analyses

Traditional ERP analyses revealed topographies and time courses similar to the ones observed in previous studies (Bekinschtein et al., 2009; Faugeras et al., 2012; King et al., 2011; Wacongne et al., 2011). EEG results from Experiment 1, summarized in Fig. 2, showed that local effects arose between approximately 130 ms and 350 ms and mainly evoked a frontal negativity (MMN) followed by a central positivity (P300a). In contrast, global effects were mainly observed from 200 ms onwards, and were characterized by a sustained centro-posterior positivity (P300b) peaking between 300 ms and 500 ms. These two effects replicate the EEG components previously reported in this type of paradigms. Interestingly, and as can be seen on Fig. 2, a vast amount of inter-individual variability can be observed in the single-subject data. This variability thus highlights the potential usefulness of tailoring the analyses to each subject.

Decoding across time circumvents individual variabilities

To maximize the detection of the neurophysiological responses elicited by local and global novelties, we applied, to each subject separately, a multivariate pattern (MVP) classifier. The classification scores of each classifier reported below refer to the area under the curve (AUC) estimated separately for each subject from a receiver operative curve analysis (see Methods).

We first aimed at characterizing the dynamics of classification scores across time, and thus trained a different classifier for each time sample. As expected, during the time period preceding the onset of the last sound, both local and global decoding remained at chance level (AUC did not differ from 50%). Accuracy of the local decoder exceeded chance level earlier in the epoch than did the

accuracy of the global decoder. As can be seen on Fig. 3a, two peaks of local decoding were observed in each experiment between 150 ms and 300 ms after stimulus onset. Mean single-trial local classification scores across subjects varied from 62.0% to 77.8% of trials depending on the experimental apparatus (all $p_{FDR} < .05$).

The dynamics of global effects were more variable across experiments. Experiment 1, in which subjects were instructed to count rare global deviant trials, revealed a quick rise of EEG-based classification scores around ~130 ms, followed by a sustained period during which global classification remained above chance almost throughout the epoch (between 300 ms and 730 ms). However, in intracranial EEG (Experiment 2), despite similar instructions, only a transient period of global decoding was seen, peaking at ~260 ms, then dropping back to chance level. Finally, Experiment 3, in which subjects were recorded with MEG and were instructed to merely pay attention to the sounds, demonstrated sustained global classification scores from 200 ms onwards (AUC = 66%). These results suggest that the neurophysiological signatures of novelty detection seem relatively independent from task instructions.

Interestingly, a sharp increase in global classification scores was observed as early as 150 ms in the attentive MEG condition (Experiment 3). Although much smaller (AUC = 55.0%, $p_{FDR} < .05$) a similar trend was also apparent in the counting EEG condition (Experiment 1). Finally, intracranial recordings, mainly taken from the temporal cortices (Supplementary Fig. 2) presented significant global classification performance between 100 ms and 300 ms (max AUC = 66.1%, $p_{FDR} < .05$). Taken together, these results suggest, contrary to what was previously suggested (Bekinschtein et al., 2009), that global novelty can affect early processes too.

Finally, subjects who were distracted by a concurrent visual task (Experiment 4) presented lower but still significant local classification scores (AUC = 63.7%, $p_{FDR} < .05$). Moreover, distraction dramatically impaired global classification scores, which consequently failed to reach significance across subjects (AUC = 54.0%, $p_{FDR} > .05$).

Overall decoding local and global effects across time showed that multivariate pattern classifiers could extract qualitatively similar effects as the ones observed from ERPs (for ERFs see Wacongne et al., 2011) but could also reveal subtle effects such as an early global effect unreported in previous studies.

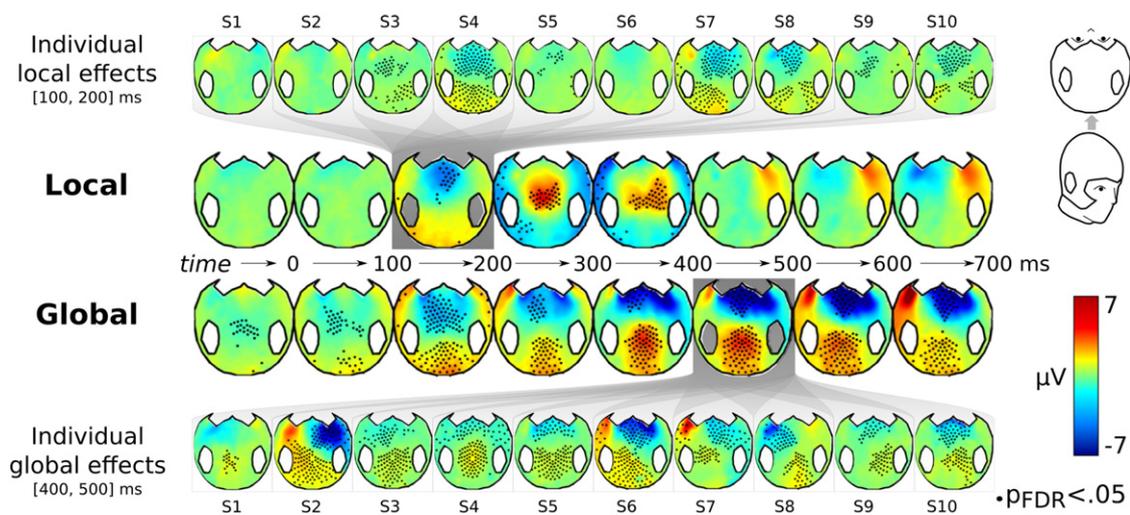


Fig. 2. Inter-individual variability in the topography of EEG effects in Experiment 1. In the two central rows, control EEG topographies of the mean local effect (LD > LS) and of mean the global effect (GD > GS) are plotted as a function of time following the onset of the fifth sound. Black dots correspond to EEG channels presenting significant differences between standard and deviant conditions, corrected for multiple comparisons with FDR. The top and bottom rows show the topography in individual subjects, separately for the local and global effects in their respective time windows of interest. Local deviant trials elicit significant differences over anterior regions with an initial negativity peaking around 150 ms followed by a short central positivity between 200 and 300 ms. Global deviant trials elicit a centro-posterior sustained positivity mainly from 300 ms onwards. While the group statistics replicate the traditional MMN and P3b associated with the Local–Global paradigm (Bekinschtein et al., 2009), individual analyses reveal substantial topographical variability in healthy subjects, and thus highlight the usefulness of tailoring the analyses to each subject.

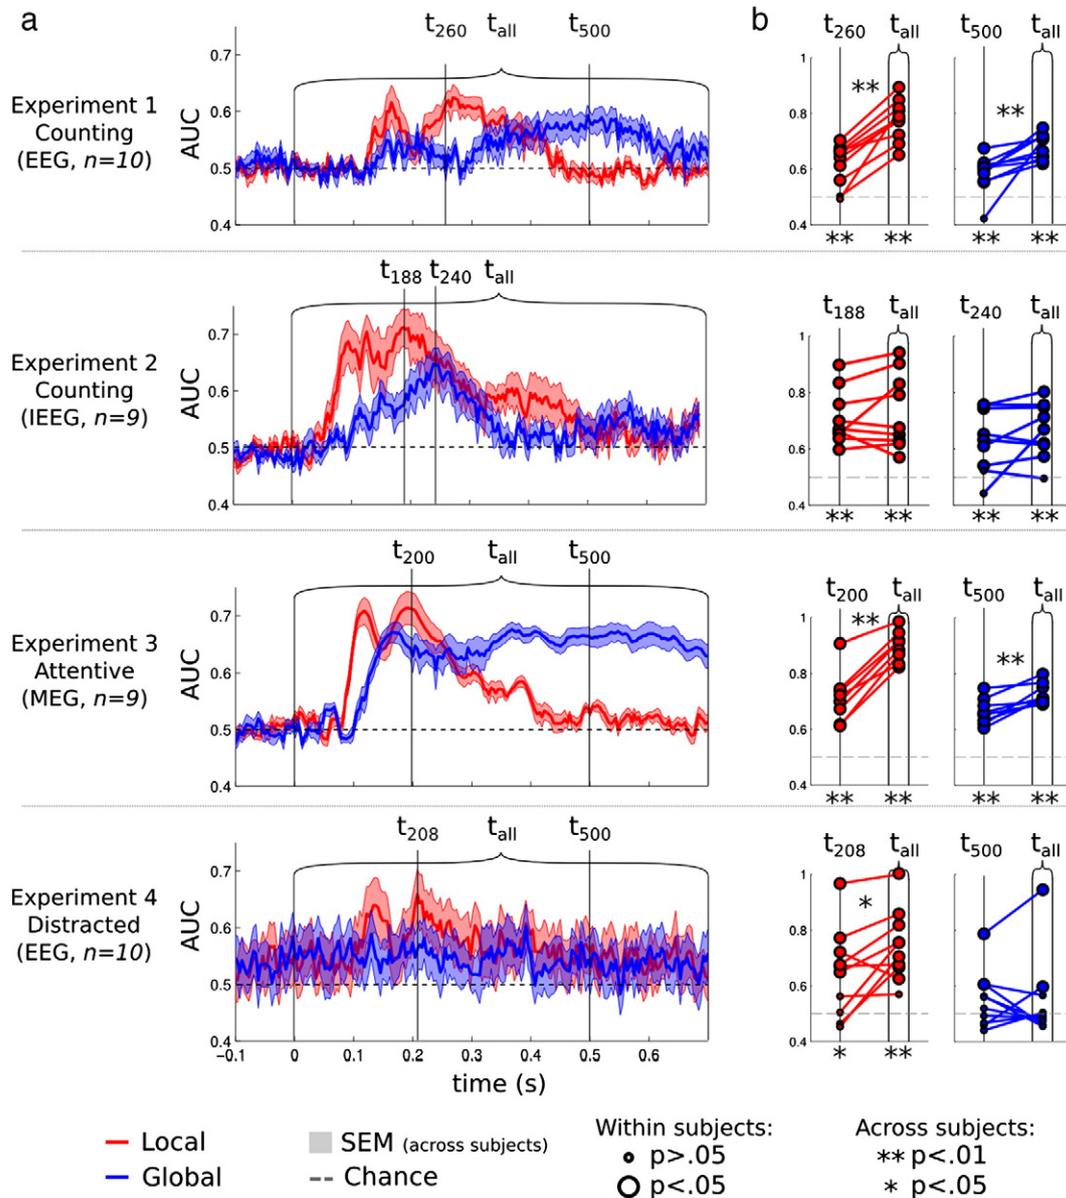


Fig. 3. Local and global multivariate decoding scores. a. Decoding of the local effect (red) and of the global effect (blue) was applied successively to each time sample of the scalp EEG, MEG and intracranial EEG recordings obtained under different experimental conditions: subjects either counted the global deviant trials (Experiments 1–2), were attentive to the sounds (Experiment 3), or were distracted by a visual task (Experiment 4). Results demonstrate a non-sustained decoding of the local effect (~130–400 ms), with above-chance performance regardless of recording apparatus and experimental condition. Decoding of the global effect appeared later and was more sustained in time, but was absent in the distracted condition. Local and global decoding scores based on the dynamics of the electrophysiological signal across multiple time samples (t_{all}) led to a significant improvement of decoding scores in most scalp recordings. b. Comparison of the decoding scores based on the best time point (left dot in each graph) or the full trial dynamics (right dot). Each dot represents the decoding score of an individual subject. All subjects but one presented significant local decoding scores. All experiments led to significant global decoding scores, except for the distracted condition. Local and global decoding scores based on the dynamics of the electrophysiological signal across multiple time samples (t_{all}) led to a significant improvement of decoding scores in most scalp recordings.

Using signal dynamics facilitates decoding

Decoding a spatial topography at each time sample, as was done in the previous section, can overcome individual topographical variability but remains dependent on the precise timing of a given effect. To maximize the extraction of local and global effects, we thus trained another decoder using the full dynamics of brain signals on a given trial. For each subject separately, and for local and global effects, this decoder was trained to distinguish the standard and deviant trials based on all the information available on a given trial (*all channels* \times *all time samples*). Results are presented in Fig. 3b.

Overall, this dynamic approach provided better classification scores than in the previous section. To prove this, we compared it to the performance of the best time sample of the previous section. Note that this is a

very conservative test. Indeed, the a posteriori selection of the best time sample is not cross-validated, and thus likely overestimates the actual scores one could hope to obtain if an independent data sample was available (e.g. Vul et al., 2009). Still, an ANOVA across subjects, contrasts (local and global) and type of classifier (dynamic versus best single time point) revealed a significant main effect of classifier type, $F(1,37) = 47.9$, $p < 10^{-10}$, indicating better classification overall with the dynamic approach (mean AUC = 72.0%) than with the previous approach (mean AUC = 63.7%). This improvement was robust across recording methods (all $p < .05$) except for intracranial EEG ($p > .33$), and therefore suggests that the classifier manages to exploit the time course of brain signals to provide information at the single-trial level (See Supplementary Fig. 1 for a depiction of single trials' local and global predictions in Experiment 3.). Supplementary analyses confirmed that the observed improvement

of decoding scores reflects an ability of the classifier to extract the dynamics of the brain signals (Fig. S3).

In more detail, using the dynamic approach, the local classification scores reached an average AUC of 77.8% ($p < .01$) in the counting EEG condition (Experiment 1), 73.2% ($p < .01$) in the counting intracranial condition (Experiment 2), 73.5% ($p < .01$) in the distracted EEG condition (Experiment 4) and 89.8% ($p < .01$) in the attentive MEG condition (Experiment 3). The latter performance, interestingly, was higher than the one obtained from high-density EEG (all $p < .01$) and even than intracranial recordings ($p < .05$). Importantly, across four experiments, all but one subject (from the distracted group) presented significantly above chance local classification scores.

Similar results were obtained for global classification scores across Experiments 1–3. The counting EEG condition (Experiment 1) led to an average AUC of 67.9% ($p < .01$), the counting intracranial condition (Experiment 2) led to an average AUC of 66.1% ($p < .01$) and the attentive MEG condition (Experiment 3) led to an average AUC of 72.3% ($p < .01$). Amongst the 28 subjects who were paying attention to the sounds, only one subject failed (from Experiment 1) to present a significant global decoding score.

Crucially, global classification scores were dramatically reduced in distracted subjects ($p < .01$ as compared to EEG and MEG recordings – Experiments 1 and 3; $p < .05$ as compared to intracranial recordings – Experiment 2). Not only was the global decoding score not-significantly different from chance for distracted subjects (AUC = 55.0%, $p > .05$), but 8 out of 10 subjects did not present any significant global decoding.

In summary, by maximizing the extraction of information from individual recordings, these results demonstrate that the human brain response to auditory deviancy, both local and global, can be detected in single trials and a fortiori in individual subjects. It also confirms that this information is not entirely dependent on the fact that subjects are asked to count the rare global deviant trials, but remains present under the instruction to merely attend to the sounds. Finally, the comparison between the active and the distracted conditions confirms the automaticity of processes underlying local novelty detection, and the necessity of attention in the detection of global novelty.

Generalization across contexts isolates rule-specific effects

The above analysis, while sensitive, could be partially affected by a contextual modulation of the local mismatch effects (MMN). Indeed, the MMN elicited by an xxxxY sequence is stronger when this sequence is rare than when it is frequent (Wacongne et al., 2012). This effect could be used by the global decoder to provide early above-chance classification of global deviants from global standards, even though it does not reflect a genuine response to global novelty. To address this issue, we reasoned that a cross-context generalization would provide a stricter criterion for a physiological signature of the brain's response to global novelty. For instance, a neuronal process responding to rare sequences, if generic, should be found whenever a rare global deviant sequence is presented, whether this sequence is xxxxY (in blocks where xxxxx is the frequent stimulus) or xxxxx (in blocks when xxxxY is the frequent stimulus). Such generalization analysis could sort out the genuine signatures of global effects from the modulations of local novelty effects. With this idea in mind, we investigated the generalization of global classification across different local contexts (and *vice-versa* see [Methods](#)), and restricted the local and the global analyses to the early (0–367 ms) and to the late (367–700 ms) time windows respectively. As intracranial classification scores were not significant in the late time window (Fig. 3), we tested this analysis on Experiments 1, 3 and 4. Results are presented in Fig. 4.

An ANOVA across subjects, contrasts (local or global) and type of classifier (using both contexts or within-context decoding) showed that classification scores resulting from within-context decoders were smaller than those obtained from decoders using all trials ([Introduction](#)

section): $F(2,28) = 26.56$, $p < 10^{-6}$. This is expected as restricting the decoder to trials performed in a given context automatically reduces the number of trials available for the training set. Note however that the within-context classification scores remained significantly above chance in all experimental conditions: Local: Experiment 1: 69.1% ($p < .01$), Experiment 3: (71.2% $p < .01$), and Experiment 4: 68.8% ($p < .01$); Global: Experiment 1: 66.0% ($p < .01$), Experiment 3: 63.2% ($p < .01$), and Experiment 4: 56.7% ($p < .05$).

Cross-context generalization scores appeared significantly worse than within-context classification scores (Local: Experiment 1: $p < .01$, Experiments 3–4: $p < .05$; Global: Experiment 1: $p < .01$, Experiments 3–4: $p < .10$). As can be seen in Fig. 4a presenting the MEG generalization scores across time, the differences between within-context and across-context scores were mainly visible i) in early time windows; and ii) in the global contrast. This result suggests that part of the early global effect previously described was due to an early modulation of novelty responses that did not generalize across contexts. However, the overall reduction of cross-context generalizations suggests that this conceptually more suitable analysis may, in practice, be less appropriate when the signal-to-noise ratio is low.

Crucially, whereas all local generalization scores remained significantly above chance (Experiment 1: 57.1%, Experiment 3: MEG: 66.3%, Experiment 4: 61.5%, all $p < .01$), global generalization scores were now exclusively significant in subjects who were paying attention to the sounds (Experiment 1: 57.9%, Experiment 3: 60.2%, both $p < .01$). Distracted subjects did not present any significant global generalization scores (48.1%, $p > .10$), and the latter were significantly smaller than the ones obtained in the counting (Experiment 1: $p < .001$) and the attentive (Experiment 3: $p < .0001$) conditions.

Finally, it is interesting to note that global generalization scores steadily increased from 180 ms after stimulus onset, until they eventually reached, at the end of the epoch, scores that resembled those obtained with the within-context decoding analysis in MEG (Fig. 4a). This pattern of results suggests that distinct neurophysiological signatures were initially elicited by the two forms of global deviancy tested here. Most likely, the rare xxxxY sequence was easily detected in an xxxxx context, whereas, conversely, the rare xxxxx sequence took a longer time to be detected in an xxxxY context. Late in the epoch, however, the neuronal activity that they evoked eventually converges to the same state, thus permitting similar levels of within-block classification and cross-block generalization.

Patients: Decoding across time and dynamics

The above results demonstrate that multivariate classifiers efficiently capture the dynamics of the brain responses to local and global deviancy. In particular, we confirmed that the processing of local novelty was relatively independent of attention, while the late effects elicited by global novelty only emerged and generalized across contexts in attentive subjects. These analyses allowed us to formulate a simple hypothesis: in non-communicating patients, cerebral activity specific to local and to global novelty may provide useful markers of automatic versus conscious processing of auditory regularities. We thus applied the present method to 158 EEG recordings acquired from patients diagnosed in vegetative state (VS), minimally conscious state (MCS) or conscious state (CS). Individual and group results are summarized in Fig. 5.

First, it appeared that the decoding of the EEG dynamics of the patients was qualitatively similar to what was observed in healthy controls: local classification scores were maximal between 150 ms and 300 ms after stimulus onset, whereas global effects, if present, appeared later and were relatively sustained between 250 ms and 700 ms (see Fig. 5a). Classifiers based on the full dynamics of the EEG signals (see [Methods](#)) were again superior to selecting the best time point obtained from healthy controls (local: $p < .10^{-5}$, global: $p < .01$) or directly from

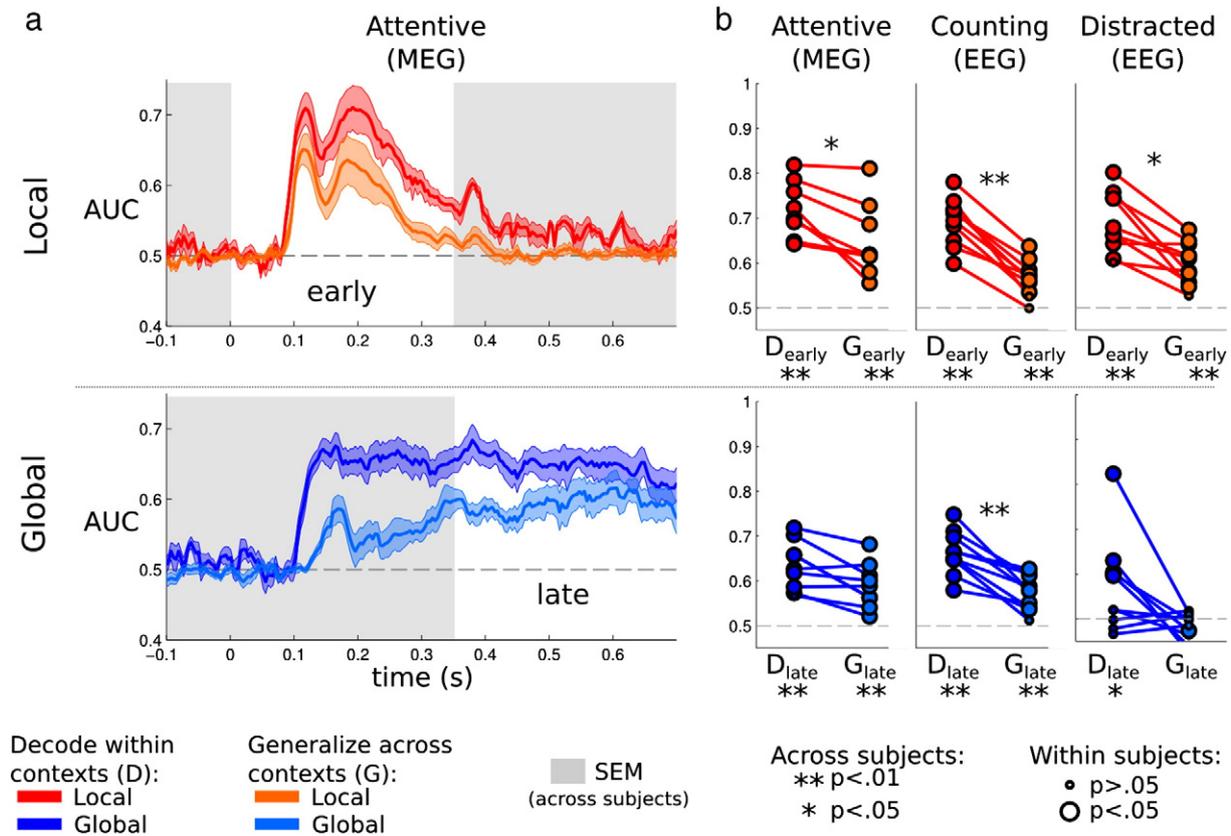


Fig. 4. Generalization of the multivariate decoding across experimental contexts. Generalization analyses consist in training the classifier in a given context (e.g. local standard trials only) and testing the ability of the classifier to generalize to another context (e.g. local deviant trials only). For example, we train the classifier to discriminate global deviants from global standards, within the context of local standard trials, and then we test the performance of the classifier on the same task, but within the context of local deviant trials. a. Local (red) and global (blue) scores obtained from within-context (D) and cross-context (G) generalization analyses are plotted as a function of time. Local effects generalized early on (orange) and rapidly vanished. Global effects generalized increasingly well over time (cyan) until generalization performance did not differ significantly from decoding performance (blue). b. Each dot represents the decoding and generalization score of an individual subject. Although the generalization of local effects was significantly smaller than decoding performance, it remains above chance in all experiments. Generalization of global effects was however only significant in the attentive and counting condition.

the patients (local: $p < .0001$, Global: positive trend $p > .10$). These results therefore confirm the utility of our decoding approach.

Despite similar qualitative results, quantitative comparisons with healthy controls (Experiment 1) showed that VS, MCS and CS patients presented, overall, lower local classification scores (AUC = 56.3%, $p < 10^{-5}$) and lower global classification scores (AUC = 51.7%, $p < 10^{-5}$). Interestingly, when comparing patients to healthy distracted subjects (Experiment 4), results demonstrated significant differences only in local classification scores ($p < 10^{-4}$) but not in global classification scores ($p > .1$).

Despite being reduced, the patients' classification scores were not at chance. Local AUCs ranged on average from 55.3% amongst VS patients, 56.1% amongst MCS patients, to 60.0% amongst CS patients. These values were above chance within each group of patients (all $p < .0001$, Fig. 5). Local classification scores of VS and MCS patients were not significantly different from one another ($p > .10$). Further analyses showed that VS patients did not differ from either MCS- patients ($p = .688$, AUC = 52.7%) or from MCS+ patients ($p = .239$, AUC = 56.8%). An ANOVA using distinct patients as a random factor only revealed a marginal interaction between the state of non- or poorly-communicating patients (1: VS, 2: MCS) and the local class (1: standard, 2: deviant): $F(1,105) = 3.63$; $p = .059$. However, CS patients presented significantly higher local classification scores than both MCS ($p < .05$) and VS ($p < .01$) patients. Crucially, for global decoding, the scores were above chance for MCS (AUC = 51.7%, $p < .01$) and CS (AUC = 56.2%, $p < .001$) patients, whereas VS patients, on average, did not present any significant global decoding effects: AUC = 50.2%, $p > .10$. Moreover, the global decoding AUCs of VS patients were significantly smaller than both those of CS patients ($p < .0001$) and those of MCS patients

($p < .05$). Note that the latter also presented lower global classification scores than CS patients ($p < .01$). Further analyses revealed that whereas VS patients did not differ from MCS- ($p = .274$, AUC = 57.4%), their scores were significantly smaller than those of MCS+ patients ($p = .038$, AUC = 62.0). An ANOVA using distinct patients as a random factor confirmed this effect by demonstrating a significant interaction between the state of non- or poorly-communicating patients (1: VS, 2: MCS) and the global class (1: standard, 2: deviant): $F(1,105) = 4.14$; $p = .044$. The above group differences could be due to a minority of subjects, or to an overall group effect. An advantage of the present single-trial MVP analysis is that it provides within-subject significance. We thus quantified the proportion of patients who presented significant local and global classification scores. The results, summarized in Fig. 5b showed that 51% of VS patients, 65% of MCS patients and 70% of CS patients presented significant local classification scores. This suggests the presence of an MMN amongst these subjects. The proportions of patients with significant Local classification scores did not differ across groups, neither with full (VS versus MCS versus CS) nor with pair-wise (VS versus MCS, MCS versus CS, and CS versus VS) chi-square tests (all $p > .10$).

The results were quite different for the global decoding, as the proportion of patients presenting significant global classification scores nearly doubled across states of consciousness. First, only 14% of VS subjects showed significant global decoding. Amongst these 14%, half presented scores which resisted a FDR correction for multiple comparisons across subjects. Like other studies, this result suggests that a small proportion of carefully diagnosed VS patients may still have residual consciousness (Monti et al., 2010). Notably, significant local and global decoding were not associated with etiologies

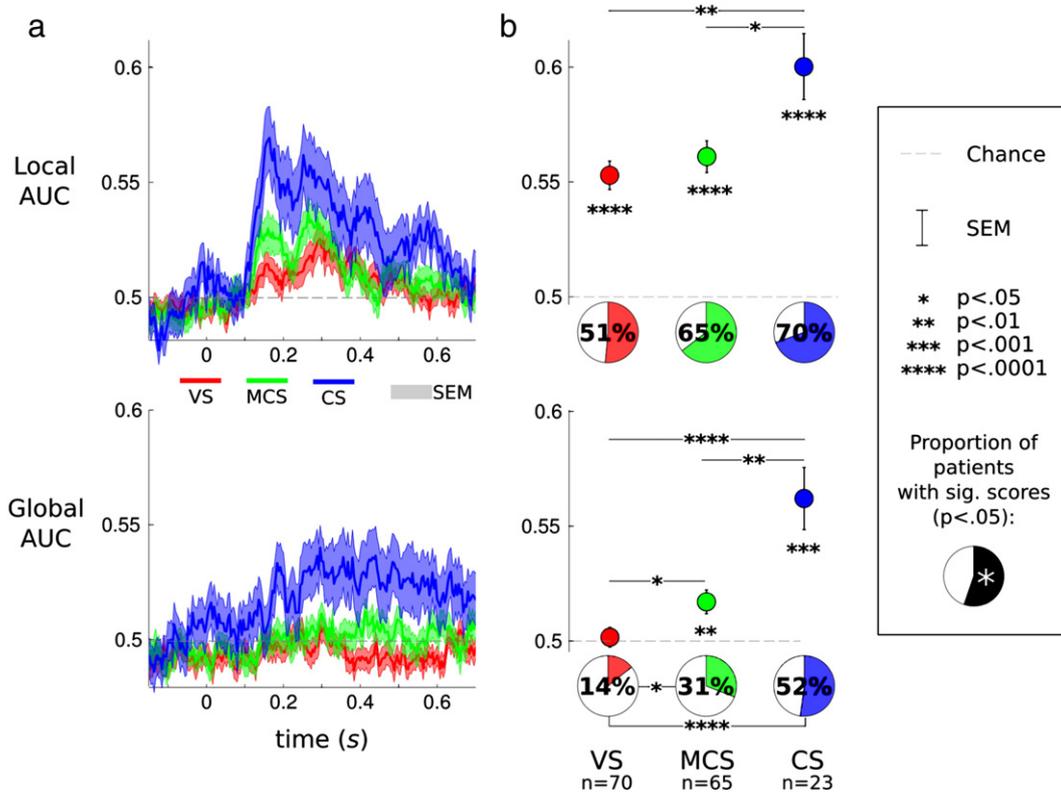


Fig. 5. Local and global decoding in patients whose state was diagnosed as vegetative (VS), minimally conscious (MCS) and conscious (CS). a. Local (top) and global (bottom) decoding scores are plotted as a function of time. Overall, local and global decoding scores follow the same qualitative trends as the ones observed in healthy subjects (Fig. 2). CS patients (blue) presented higher local and global decoding scores over time than either MCS (green) or VS (red) patients. b. Decoding scores obtained from the EEG dynamics of each trial are summarized for each state of consciousness. The graphs give the mean and standard error of the decoding scores in each group, its significance relative to chance level, and the significance of pair-wise group comparisons. Pie charts summarize the proportion of patients who presented significant decoding scores.

(local: $\chi^2(3, N = 70) = 2.30, p = .513$; global: $\chi^2(3, N = 70) = .66, p = .882$) nor with the delay separating the insult from the EEG recording (both $p > .669$). Second, in contrast, 31% of the MCS patients and 52% of the conscious patients presented significant global scores. The proportions of significant global scores across the three states of consciousness were significantly different from one another: $\chi^2(2, N = 158) = 13.73, p < .001$. Pair-wise chi-square tests confirmed this finding by showing a smaller proportion of patients with global decoding in VS relative to MCS ($\chi^2(1, N = 135) = 4.39, p = .036$) and to CS ($\chi^2(1, N = 93) = 11.74, p = .0001$). No differences were observed between the proportions of CS and MCS patients with significant global decoding ($\chi^2(1, N = 88) = 2.50, p = .113$).

To increase the specificity of the present results, we also applied the generalization method detailed above. Scores obtained from within-context classification were similar to but lower than the classification scores obtained from the joint analysis of both contexts (unsurprisingly given that the number of training trials was halved). For the local effect, while not differing from one another (VS-MCS: $p = .252$), local scores from VS (AUC = 56%, $p < .0001$) and MCS (AUC = 57%, $p < .0001$) were significantly smaller than those obtained than CS patients (AUC = 61%, $p < .0001$, CS-VS: $p = .005$, CS-MCS: $p = .039$). Global within-context classification scores presented a similar pattern: VS (AUC = 54%, $p < .0001$) and MCS (54%, $p < .0001$) scores did not differ from each other ($p = .969$), but were significantly smaller from CS (59%; both $p = .001$). Finally, cross-context generalization scores were low (all below 53%). Local cross-context generalization performance was lower for VS patients than for MCS ($p = .038$) and CS ($p = .003$) patients, but none of the global cross-context generalization scores differed from one another (all $p > .287$).

In summary, by quantifying the proportion of patients with significant local and global effects, we confirmed the earlier finding that DOC patients may still exhibit similar local mismatch effects, relatively

independently of their conditions (Bekinschtein et al., 2009; Faugeras et al., 2012; Fischer et al., 1999; Naccache et al., 2005; Tzovara et al., 2013), and found that the capacity to show global effects was modulated by the state of consciousness. A classifier based on the whole set of trials and their temporal dynamics proved the most sensitive tool to detect these local and global effects. Cross context generalization was unable to dissociate the three types of patients, presumably because of a halving of the number of training trials.

Discussion

In this study, we investigated, at the single trial level, the neuronal response following the violation of two embedded auditory regularities structured across two different time ranges (“local” and “global”). We implemented, for this purpose, a series of multivariate pattern (MVP) analyses extracting this information from the temporal dynamics of the neurophysiological activity recorded with high-density EEG (Experiment 1 and 4, $n = 20$), intracranial EEG (Experiment 2, $n = 9$) or MEG (Experiment 3, $n = 9$). Analyses were performed on attentive and distracted healthy control subjects, as well as on three types of patients, namely patients in vegetative state (VS, $n = 70$), minimally conscious state (MCS, $n = 65$) and conscious state (CS, $n = 23$), all recorded with high-density EEG for clinical purposes.

Results from control experiments revealed that when subjects were attentive, single-trial classification could lead to AUCs between 73% and 90% for local novelty, and between 66% and 72% for global novelty, depending on the recording apparatus (Experiments 1–3). Although EEG, MEG and IEEG record different types of signals and noise and are therefore not easily comparable, it is remarkable to note that MEG recordings achieved classification scores comparable to, or even higher than intracranial EEG data. Moreover, we showed that providing the decoder with multiple time samples systematically

improved classification as compared to a decoder trained on the single best time sample. This result therefore demonstrates the utility of MVP classifiers in the present context, and confirms that this method can reliably and automatically extract neuronal dynamics specific to each subject in order to efficiently classify each trial.

It should be noted that our method, although different in technical details, follows a similar approach to that taken in a recent study (Tzovara et al., 2013). Tzovara and collaborators decoded, from EEG recordings of coma patients, the difference between regular and irregular trials, that is, the equivalent of the local effect in the present study. To classify each trial, the authors used a different method which modeled the ERPs with a mixture of Gaussians. Advantages and disadvantages can be identified in each of these classification methods. Our approach does not make any Gaussian assumptions, can make use of imbalanced training dataset and can extract a large number of different types of topographies by searching across all channels – which therefore are not restricted to EEG. In contrast, Tzovara and collaborators first transform the ERPs into components. This computational step reduces the dimensionality of the data, and thus likely improves the efficiency of the ensuing classification process. However, their classifier is unable to optimally use imbalanced datasets, which are necessarily encountered in this type of novelty paradigm. Taken together, these differences may explain why our approach provided slightly better classification scores (AUC = 77.8%) than theirs (AUC = 71%) in similar control subjects recorded with EEG. Future efforts should capitalize on a combination of these methodological technicalities.

Crucially, our MVP classifiers replicated and extended previous observations on the key electrophysiological properties of local and global effects (Bekinschtein et al., 2009). As expected, local novelty mainly affected the early part of the neural signal (<300 ms) and remained unaffected by visual distraction. In sharp contrast, global novelty elicited late and stable effects in both counting and attentive subjects, but most distracted subjects presented dramatically reduced, and in fact, non significant global classification scores.

The present decoding approach allowed us to investigate whether the neuronal activities elicited by local and global novelties are specific and reproducible enough to generalize from one context to the other. Specifically, we asked which aspects of the neuronal response to global novelty could generalize from a context in which the xxxxy sequence was rare (in xxxxx blocks) to another in which the xxxxy sequence was rare (in xxxxy blocks), and *vice-versa*. Generalizations of the global effects confirmed a progressive increase in the similarity of these two types of global novelty responses from 180 ms onwards, until they eventually fully resembled each other at the end of the trial. Conversely, local novelty detection showed a significant generalization across contexts only in the early part of the event-related response. Taken together, these results reveal a double dissociation. Local effects are attention independent, context dependent and are not maintained across time. By contrast, late global effects are attention-dependent, context-independent and stable for a prolonged temporal duration.

One of the major motivations of the Local–Global paradigm is to provide a minimal design that dissociates two types of auditory novelty processing – one generating the MMN followed by a P300a, both likely being automatic, and the other generating the P3b which depends on working memory and conscious access. This goal is of particular importance for DOC patients. Despite an increasing interest for this clinical population (Laureys et al., 2004; Owen et al., 2009), unambiguously distinguishing patients suffering from communication disorders from those with a genuine loss of conscious processing remains a challenging task (Laureys and Schiff, 2011; Owen et al., 2006). In this context, our present capacity to dissociate, from EEG alone, an automatic process of local novelty detection from a later process that depends on conscious processing thus opens up the possibility of detecting states of consciousness independently of the patient's ability to communicate. Remarkably, Bekinschtein et al.

(2009) have applied the Local–Global paradigm to four MCS and four VS patients, and shown that the latter were less likely to present late global effects than the former. More recent studies testing larger groups of patients however failed to identify significant global effects in these two groups of patients (Faugeras et al., 2012).

We here applied our MVP classifier to 158 EEG recordings acquired at bedside from awake but non- or poorly communicating patients, whose state of consciousness (VS, MCS or CS) was assessed immediately before the experiment. The state of consciousness is here determined clinically based on the CRS-R (Giacino et al., 2004; Schnakers et al., 2008). This immediate assessment differs from clinical practice in which several assessments are often conducted to maximize the chance of detecting residual consciousness. Here, we aim at identifying a neural marker of consciousness state. It is therefore crucial to determine the patients' state of consciousness at the time of their EEG recording.

The results demonstrated that, at the group level, our method could accurately distinguish neuronal responses elicited by local deviants relative to local standard sounds: 51% of VS, 65% of MCS and 70% of CS patients presented significant local classification scores. Moreover, whereas CS patients presented significantly higher classification scores, VS and MCS patients did not differ from one another, suggesting that the cerebral processes that detect local novelty are not unique to conscious processing and can remain functional in all states of consciousness. These results are in line with a series of studies demonstrating the presence of the MMN in many DOC patients (Bekinschtein et al., 2009; Faugeras et al., 2012; Fischer et al., 2010; Wijnen et al., 2007), including coma patients (Fischer et al., 1999; Kane et al., 1996, 2000; Naccache et al., 2005; Tzovara et al., 2013). Indeed, they confirm that the effects elicited by local changes in pitch, as observed in the auditory odd-ball paradigm, are poor predictors of the state of consciousness (Bekinschtein et al., 2009; Tzovara et al., 2013). The lack of a detectable local effect in some patients could be due to a variety of causes, including poor signal-to-noise, excessive number of artifacted trials, or an impairment to auditory pathways.

The situation was quite different for the global effect. Crucially, the detection of global deviants was considerably reduced in patients as compared to healthy controls, but as predicted it remained weakly but significantly above chance amongst MCS and CS patients. These results demonstrate that the present approach outperforms traditional methods which have, so far, failed to identify a significant global effect in these two groups of patients (Faugeras et al., 2012). Importantly, as a group, VS patients did not present any significant global effects, and their classification scores were smaller than those of MCS and CS patients. Taking advantage of the fact that the present analyses allow for single-subject predictions, we showed that only a small proportion (14%) of VS patients exhibited a significant global decoding score, and that this proportion was significantly smaller than the proportions of MCS (31%) and CS (52%) patients in whom a global effect could be significantly detected. This finding confirms that global effects provide a reliable, although partial, index of the state of consciousness.

Further studies will need to assess the sensitivity and specificity of the global effect as a test of consciousness. Although a small fraction of VS patients still showed a significant global effect, this needs not necessarily imply a failure of our test, but rather may indicate that the clinical label of VS may not be fully reliable. Indeed, fMRI studies also reveal that a fraction of VS patients, continue to present complex cortical responses that suggest preserved consciousness (Monti et al., 2010). We have previously reported that two VS patients with a significant global effect moved to the MCS category within the next few days (Faugeras et al., 2011).

In the converse direction, unfortunately, our test clearly lacks sensitivity since about half of CS patients and two-third of MCS patients did not present any significant global classification scores. Further research should determine i) if this problem can be remedied, for instance, by using a larger number of trials, different days of testing,

or better noise reduction techniques, or ii) if it reflects a genuine cognitive limit, whereby patients are conscious but too cognitively impaired to successfully detect global novelties. Another important issue is that the global effect is known to vanish under conditions of inattention (Bekinschtein et al., 2009). Yet, our auditory stimuli are monotonous and devoid of interest. Some patients are thus likely to lose focus during the 40 min recording session. In the future, special efforts should thus be dedicated to enhance the patients' motivation and attention towards the stimuli, as well as to improve the quality of EEG recordings.

These difficulties are not unique to our test. Most other tests of consciousness currently require patients to understand and maintain, for several minutes, a complex instruction such as imagining playing tennis (Owen et al., 2006), or retrieving the answer to a spoken question (Monti et al., 2010). These tests, like ours, are therefore asymmetrical: when positive, they are highly indicative of preserved consciousness, but they may also fail to detect residual consciousness if the patient suffers from hearing, linguistic, attentional or working memory deficits. For instance, in the Monti et al. (2010)'s study, 30 out of 31 MCS patients, who therefore gave occasional behavioral signs of consciousness upon clinical examination, showed no sign of response to command via imagined tennis playing or spatial navigation. Relative to these fMRI studies, the current approach presents at least two advantages. First, it relies on a method (EEG) which is easy to implement, available in all clinics and applicable at bedside. Second, it only depends on subjects' attention to the stimuli and does not seem to require complex task instructions.

While the ability of the classifiers to discriminate single trials remained significant at the group level – as well as within several patients – the single trials were often poorly discriminated. Consequently, while the present method may be useful for an overall detection of consciousness, it is unlikely to be an efficient tool for monitoring of consciousness at bedside. Some patients however presented relatively high decoding accuracies (up to 83.1% in the local classification and up to 76.5% in the global classification) and could thus be potential subjects for these types of online paradigms.

Alternative ways of investigating consciousness, which bypass entirely the need to attend to external stimuli, are also being developed. For instance, Massimini and collaborators have developed a TMS–EEG apparatus that tests the complexity and the functional connectivity of brain responses to TMS pulses. The results demonstrate that this artificial probe differentiates well the states of consciousness in sleep (Massimini et al., 2010), anesthesia (Ferrarelli et al., 2010) and DOC patients (Rosanova et al., 2012). Our own research, again using only high-density EEG, also suggests that the intrinsic complexity of the EEG and, especially, the amount of information shared across distant electrode sites, provides an index that usefully complements the present approach. Ultimately, a combination of simple experimental paradigms with sophisticated signal post-processing, as attempted here, may prove crucial for the automatic detection of conscious processing in non-communicating subjects.

Acknowledgments

We thank Gael Varoquaux for useful discussions, and the Neurospin, the ICM and the hospital infrastructure for their generous help and administrative support. This work was supported by a Direction Générale de l'Armement (DGA) grant to J.R.K., by an 'Equipe FRM 2010' grant of Fondation pour la Recherche Médicale (FRM) to L.N., by Institut National de la Santé et de la Recherche Médicale (INSERM), Commissariat à l'Energie Atomique (CEA), and a European Research Council (ERC) senior grant "NeuroConsc" to S.D., by Journées de Neurologie de Langue Française and FRM (Fondation pour la Recherche Médicale) to F.F., by INSERM to B.R. and by AXA Research Fund to I.E.K. The Neurospin MEG facility was sponsored by grants from INSERM, CEA, FRM, the Bettencourt-Schueller Foundation, and the Région île-de-France.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.neuroimage.2013.07.013>.

Conflict of interest

The authors declare no conflict of interest.

References

- Atienza, M., Cantero, L., 1997. The mismatch negativity component reveals the sensory memory during REM sleep in humans. *Neurosci. Lett.* 237, 21–24.
- Bekinschtein, T. a, Dehaene, S., Rohaut, B., Tadel, F., Cohen, L., Naccache, L., 2009. Neural signature of the conscious processing of auditory regularities. *Proc. Natl. Acad. Sci. U. S. A.* 106, 1672–1677.
- Brázdil, M., Rektor, I., Daniel, P., Dufek, M., Jurák, P., 2001. Intracerebral event-related potentials to subthreshold target stimuli. *Clin. Neurophysiol.* 112, 650–661.
- Bruno, M.-A., Gosseries, O., Ledoux, D., Hustinx, R., Laureys, S., 2011a. Assessment of consciousness with electrophysiological and neurological imaging techniques. *Curr. Opin. Crit. Care* 17, 146–151.
- Bruno, M.-A., Vanhauwenhuyse, A., Thibaut, A., Moonen, G., Laureys, S., 2011b. From unresponsive wakefulness to minimally conscious PLUS and functional locked-in syndromes: recent advances in our understanding of disorders of consciousness. *J. Neurol.* 258, 1373–1384.
- Chang, C., Lin, C., 2001. LIBSVM: a library for support vector machines. *Computer* 2, 1–30.
- Dehaene, S., Changeux, J.-P., 2011. Experimental and theoretical approaches to conscious processing. *Neuron* 70, 200–227.
- Dehaene, S., Naccache, L., 2001. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cogn. Neurosci. Conscious.* 1–37.
- Dehaene, S., Changeux, J.P., Naccache, L., Sackur, J., Sergent, C., 2006. Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends Cogn. Sci.* 10, 204–211.
- Fahrenfort, J.J., Scholte, H.S., Lamme, V. a F., 2007. Masking disrupts reentrant processing in human visual cortex. *J. Cogn. Neurosci.* 19, 1488–1497.
- Fahrenfort, J.J., Snijders, T.M., Heinen, K., Van Gaal, S., Scholte, H.S., Lamme, V.A.F., 2012. Neuronal integration in visual cortex elevates face category tuning to conscious face perception. *Proc. Natl. Acad. Sci.* 109, 21504–21509.
- Faugeras, F., Rohaut, B., Weiss, N., Bekinschtein, T. a, Galanaud, D., Puybasset, L., Bolgert, F., Sergent, C., Cohen, L., Dehaene, S., Naccache, L., 2011. Probing consciousness with event-related potentials in the vegetative state. *Neurology* 77, 264–268.
- Faugeras, F., Rohaut, B., Weiss, N., Bekinschtein, T., Galanaud, D., Puybasset, L., Bolgert, F., Sergent, C., Cohen, L., Dehaene, S., Naccache, L., 2012. Event related potentials elicited by violations of auditory regularities in patients with impaired consciousness. *Neuropsychologia* 50, 403–418.
- Ferrarelli, F., Massimini, M., Sarasso, S., Casali, A., Riedner, B. a, Angelini, G., Tononi, G., Pearce, R. a, 2010. Breakdown in cortical effective connectivity during midazolam-induced loss of consciousness. *Proc. Natl. Acad. Sci. U. S. A.* 107, 2681–2686.
- Fisch, L., Privman, E., Ramot, M., Harel, M., Nir, Y., Kipervasser, S., Andelman, F., Neufeld, M.Y., Kramer, U., Fried, I., Malach, R., 2009. Neural "ignition": enhanced activation linked to perceptual awareness in human ventral stream visual cortex. *Neuron* 64, 562–574.
- Fischer, C., Morlet, D., Bouchet, P., Luaute, J., Jourdan, C., Salord, F., 1999. Mismatch negativity and late auditory evoked potentials in comatose patients. *Clin. Neurophysiol.* 110, 1601–1610.
- Fischer, C., Luaute, J., Morlet, D., 2010. Event-related potentials (MMN and novelty P3) in permanent vegetative or minimally conscious states. *Clin. Neurophysiol.* 121, 1032–1042.
- Gaillard, R., Dehaene, S., Adam, C., Clémenceau, S., Hasboun, D., Baulac, M., Cohen, L., Naccache, L., 2009. Converging intracranial markers of conscious access. *PLoS Biol.* 7, 1–21.
- Garrido, M.I., Kilner, J.M., Kiebel, S.J., Stephan, K.E., Friston, K.J., 2007. Dynamic causal modelling of evoked potentials: a reproducibility study. *Neuroimage* 36, 571–580.
- Garrido, M.I., Friston, K.J., Kiebel, S.J., Stephan, K.E., Baldeweg, T., Kilner, J.M., 2008. The functional anatomy of the MMN: a DCM study of the roving paradigm. *Neuroimage* 42, 936–944.
- Giardino, J.T., Ashwal, S., Childs, N., Cranford, R., Jennett, B., Katz, D.I., Kelly, J.P., Rosenberg, J.H., Whyte, J., Zafonte, R.D., 2002. The minimally conscious state: definition and diagnostic criteria. *Neurology* 58, 349.
- Giardino, J.T., Kalmar, K., Whyte, J., 2004. The JFK Coma Recovery Scale–Revised: measurement characteristics and diagnostic utility. *Arch. Phys. Med. Rehabil.* 85, 2020.
- Goldfine, A.M., Victor, J.D., Conte, M.M., Bardini, J.C., Schiff, N.D., 2012. Authors' reply *Lancet* 379, 1702.
- Goldstein, A., Spencer, K.M., Donchin, E., 2002. The influence of stimulus deviance and novelty on the P300 and novelty P3. *Psychophysiology* 39, 781–790.
- Heinke, W., Kennner, R., Gunter, T.C., Sammler, D., Olthoff, D., Koelsch, S., 2004. Sequential effects of increasing propofol sedation on frontal and temporal cortices as indexed by auditory event-related potentials. *Anesthesiology* 100, 617.
- Kane, N.M., Curry, S.H., Rowlands, C.A., Manara, A.R., Lewis, T., Moss, T., Cummins, B.H., Butler, S.R., 1996. Event-related potentials – neurophysiological tools for predicting emergence and early outcome from traumatic coma. *Intensive Care Med.* 22, 39–46.

- Kane, N.M., Butler, S.R., Simpson, T., 2000. Coma outcome prediction using event-related potentials: P(3) and mismatch negativity. *Audiol. Neurootol.* 5, 186–191.
- King, J.-R., Bekinschtein, T., Dehaene, S., 2011. Comment on “Preserved feedforward but impaired top-down processes in the vegetative state”. *Science* (N. Y., N. Y.) 334, 1203.
- Lamme, V.A.F., 2006. Towards a true neural stance on consciousness. *Trends Cogn. Sci.* 10, 494–501.
- Laureys, S., Schiff, N.D., 2011. Coma and consciousness: paradigms (re)framed by neuroimaging. *Neuroimage* 61, 478–491.
- Laureys, S., Owen, A.M., Schiff, N.D., 2004. Brain function in coma, vegetative state, and related disorders. *Lancet Neurol.* 3, 537–546.
- Laureys, S., Celesia, G.G., Cohadon, F., Lavrijsen, J., León-Carrión, J., Sannita, W.G., Szabon, L., Schmutzhard, E., Von Wild, K.R., Zeman, A., Dolce, G., 2010. Unresponsive wakefulness syndrome: a new name for the vegetative state or apallic syndrome. *BMC Med.* 8, 68.
- Lemm, S., Blankertz, B., Dickhaus, T., Müller, K.-R., 2011. Introduction to machine learning for brain imaging. *Neuroimage* 56, 387–399.
- Mason, S.J.J., Graham, N.E.E., 2002. Areas beneath the relative operating characteristics (ROC) and relative operating levels (ROL) curves: statistical significance and interpretation. *Q. J. R. Meteorol. Soc.* 128, 2145–2166.
- Massimini, M., Ferrarelli, F., Murphy, M., Huber, R., Riedner, B., Casarotto, S., Tononi, G., 2010. Cortical reactivity and effective connectivity during REM sleep in humans. *Cogn. Neurosci.* 1, 176–183.
- Melloni, L., Molina, C., Pena, M., Torres, D., Singer, W., Rodriguez, E., 2007. Synchronization of neural activity across cortical areas correlates with conscious perception. *J. Neurosci.* 27, 2858–2865.
- Melloni, L., Schwiedrzik, C.M., Müller, N., Rodriguez, E., Singer, W., 2011. Expectations change the signatures and timing of electrophysiological correlates of perceptual awareness. *J. Neurosci.* 31, 1386–1396.
- Monti, M.M., Vanhaudenhuyse, A., Coleman, M.R., Boly, M., Pickard, J.D., Tshibanda, L., Owen, A.M., Laureys, S., 2010. Willful modulation of brain activity in disorders of consciousness. *N. Engl. J. Med.* 362, 579–589.
- Muller-Gass, A., Macdonald, M., Schröger, E., Sculthorpe, L., Campbell, K., 2007. Evidence for the auditory P3a reflecting an automatic process: elicitation during highly-focused continuous visual attention. *Brain Res.* 1170, 71–78.
- Näätänen, R., Gaillard, A.W.K., Mäntylä, S., 1978. Early selective-attention effect on evoked potential reinterpreted. *Acta Psychol.* 42, 313–329.
- Näätänen, R., Astikainen, P., Ruusuvirta, T., Huottilainen, M., 2010. Automatic auditory intelligence: an expression of the sensory-cognitive core of cognitive processes. *Brain Res. Rev.* 64, 123–136.
- Naccache, L., Puybasset, L., Gaillard, R., Serve, E., Willer, J.C., 2005. Auditory mismatch negativity is a good predictor of awakening in comatose patients: a fast and reliable procedure. *Clin. Neurophysiol.* 116, 988.
- Newcombe, V.F.J., Williams, G.B., Scoffings, D., Cross, J., Carpenter, T.A., Pickard, J.D., Menon, D.K., 2010. Aetiological differences in neuroanatomy of the vegetative state: insights from diffusion tensor imaging and functional implications. *J. Neurol. Neurosurg. Psychiatry* 81, 552–561.
- Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.-M., 2011. FieldTrip: open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput. Intell. Neurosci.* 2011, 156869.
- Owen, A.M., 2008. Disorders of consciousness. *Ann. N. Y. Acad. Sci.* 1124, 225–238.
- Owen, A.M., Coleman, M.R., Boly, M., Davis, M.H., Laureys, S., Pickard, J.D., 2006. Detecting awareness in the vegetative state. *Science* (N. Y., N. Y.) 313, 1402.
- Owen, A.M., Schiff, N.D., Laureys, S., 2009. A New Era of Coma and Consciousness Science. Elsevier, *Progress in Brain Research*.
- Pedregosa, F., Weiss, R., Brucher, M., 2011. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Platt, J.C., 1999. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Adv. Large Margin Classif.* 10, 61–74.
- Polich, J., 2007. Updating P300: an integrative theory of P3a and P3b. *Clin. Neurophysiol.* 118, 2128–2148.
- Rees, G., Kreiman, G., Koch, C., 2002. Neural correlates of consciousness in humans. *Nat. Rev. Neurosci.* 3, 261–270.
- Rohaut, B., Faugeras, F., Bekinschtein, T.A., Wassouf, a., Chausson, N., Dehaene, S., Naccache, L., 2009. Prédiction du réveil et détection de la conscience: intérêt des potentiels évoqués cognitifs. *Reanimation* 18, 659–663.
- Rosanov, M., Gosseries, O., Casarotto, S., Boly, M., Casali, A.G., Bruno, M.-A., Mariotti, M., Boveroux, P., Tononi, G., Laureys, S., Massimini, M., 2012. Recovery of cortical effective connectivity and recovery of consciousness in vegetative patients. *Brain* 135, 1308–1320.
- Schnakers, C., Majerus, S., Giacino, J., Vanhaudenhuyse, A., Bruno, M.-A., Boly, M., Moonen, G., Damas, P., Lambermont, B., Lamy, M., Damas, F., Ventura, M., Laureys, S., 2008. A French validation study of the Coma Recovery Scale-Revised (CRS-R). *Brain Inj.* 22, 786–792.
- Sergent, C., Baillet, S., Dehaene, S., 2005. Timing of the brain events underlying access to consciousness during the attentional blink. *Nat. Neurosci.* 8, 1391–1400.
- Seth, A.K., Barrett, A.B., Barnett, L., 2011. Causal density and integrated information as measures of conscious level. *Philos. Transact. A Math. Phys. Eng. Sci.* 369, 3748–3767.
- Taulu, S., Kajola, M., Simola, J., 2004. Suppression of interference and artifacts by the Signal Space Separation Method. *Brain topogr.* 16 (4), 269–275. <http://www.ncbi.nlm.nih.gov/pubmed/15379226>.
- Tononi, G., Edelman, G.M., 1998. Consciousness and complexity. *Science* (N. Y., N. Y.) 282, 1846.
- Tononi, G., Sporns, O., 2003. Measuring information integration. *BMC Neurosci.* 4, 31.
- Tshibanda, L., Vanhaudenhuyse, A., Galanaud, D., Boly, M., Laureys, S., Puybasset, L., 2009. Magnetic resonance spectroscopy and diffusion tensor imaging in coma survivors: promises and pitfalls. *Prog. Brain Res.* 177, 215–229.
- Tzovara, A., Rossetti, A.O., Spierer, L., Grivel, J., Murray, M.M., Oddo, M., De Lucia, M., 2013. Progression of auditory discrimination based on neural decoding predicts awakening from coma. *Brain* 136 (1), 81–89.
- Vul, E., Harris, C., Winkelman, P., Pashler, H., 2009. Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspect. Psychol. Sci.* 4, 274–290.
- Wacongne, C., Labyt, E., Van Wassenhove, V., Bekinschtein, T., Naccache, L., Dehaene, S., 2011. Evidence for a hierarchy of predictions and prediction errors in human cortex. *Proc. Natl. Acad. Sci. U. S. A.* 108, 20754–20759.
- Wacongne, C., Changeux, J.-P., Dehaene, S., 2012. A neuronal model of predictive coding accounting for the mismatch negativity. *J. Neurosci.* 32, 3665–3678.
- Wijnen, V.J.M., Van Boxtel, G.J.M., Eilander, H.J., De Gelder, B., 2007. Mismatch negativity predicts recovery from the vegetative state. *Clin. Neurophysiol.* 118, 597–605.
- Winkler, I., 2007. Interpreting the mismatch negativity. 21, 147–163.