

Measuring states of pathological (un)consciousness: research dimensions, clinical applications, and ethics[†]

Athena Demertzi,^{1,2,*} Jacobo Diego Sitt,^{1,3} Simone Sarasso,⁴ and Wim Pinxten^{5,‡}

¹Brain and Spine Institute- Institut du Cerveau et de la Moelle épinière (ICM), Hôpital Pitié-Salpêtrière, 47, bd de l'Hôpital - 75013 Paris, France; ²Coma Science Group, GIGA Research, CHU Sart Tilman B34-Quartier Hôpital, Avenue de l'Hôpital, 11 4000 Liège, Belgium; ³INSERM, U 1127, F-75013, Paris, France; ⁴Department of Biomedical and Clinical Sciences "L. Sacco", University of Milan, Via G.B. Grassi, 74. 20157, Milano, Italy and; ⁵Faculty of Medicine and Life Sciences, Hasselt University, Martelarenlaan 42, 3500 Hasselt, Belgium

[†]Paper based on Symposium held at ASSC20, Buenos Aires, Argentina, 17 June 2016

[‡]Author list according to the order of presentations during the Symposium

*Correspondence address. Institut du Cerveau et de la Moelle épinière (ICM), Hôpital Pitié-Salpêtrière, 47, bd de l'Hôpital, 75013 PARIS, France.

Tel: + 33 (0)1 57 27 40 00; Email: a.demertzi@ulg.ac.be

Abstract

Consciousness is a multidimensional construct with no widely accepted definition. Especially in pathological conditions, it is less clear what exactly is meant by (un)consciousness, how it can be reliably observed or measured. Here, we aim at (i) bringing together state of the art approaches to classification of single patients suffering from disorders of consciousness by means of motor-independent assessment of consciousness states with electrophysiology and functional neuroimaging, (ii) showing how each proposed metric translates into clinical practice and (iii) raising a discussion on the ethical aspects of consciousness measurements. We realize that when dealing with patients some issues commonly pertain to each methodology discussed here, such as the overall clinical condition, clinical heterogeneity, and diagnostic uncertainty. When predicting patients' diagnosis, though, each method adopts a different approach to determine (a) a "gold standard" of the benchmark population upon which the metric is computed and (b) the generalization and replicability in the attempt to avoid overfitting. From an applied ethics perspective, the focus is, hence, on knowing what one is measuring and on the validity of measurements. We conclude that, when searching for consciousness in pathological conditions, confident diagnosis can be based on the use of probabilistic predictions as well as on accumulative evidence stemming from multiple non-overlapping assessments with different modalities. A framework which will regulate the application order of these techniques (balancing their availability, sensitivity, and specificity, based on underlying clinical assumptions about a patient's conscious state), is expected to ameliorate clinical management and further inform on the critical patterns of (un)consciousness.

Key words: disorders of consciousness; classification; electrophysiology; TMS; functional neuroimaging; neuroethics

Received: 30 December 2016; Revised: 2 February 2017. Accepted: 18 April 2017

© The Author 2017. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact journals.permissions@oup.com

Introduction

Defining consciousness and its disorders

Consciousness is a multidimensional construct for which there is no universal definition (Baars 2015). Although there is a common sense of what conscious experience feels like in healthy conditions, descriptions of such experiences may vary in focus (e.g. perceptual, visual, being conscious of redness). However, in pathological conditions, it is less clear what exactly is meant by (un)consciousness, how it can be reliably observed or measured. In order to overcome prominent philosophical challenges around consciousness (e.g. Demertzi et al. 2009a), we here adopt an operational definition coming from clinical neurology, which evaluates consciousness based on two dimensions: wakefulness and awareness (Posner et al. 2007; Giacino et al. 2014). Wakefulness refers to the level of vigilance and is supported by the function of the subcortical arousal systems in the brainstem, midbrain, and thalamus; clinically, it is indicated by eyes-opening. Awareness refers to the contents of consciousness and it is thought to be supported by the functional integrity of the cerebral cortex and its subcortical connections; clinically, it is assessed by evaluating command following and by observing nonreflex motor behaviors, such as eye tracking and oriented movements to pain. Based on this definition, patients in coma and under anesthesia are not conscious because they cannot be awakened. An interesting dissociation between these two components comes from patients with disorders of consciousness, i.e. those in vegetative state/unresponsive wakefulness syndrome (VS/UWS) and in minimally conscious state (MCS). The VS/UWS is characterized as a “state of arousal without awareness” (Jennett and Plum 1972) because, although patients show intermittent wakefulness (manifested as eyes-open/eyes-closed periods), preserved hypothalamic and brainstem autonomic functions (permitting survival with medical and nursing care) as well as variably preserved cranial nerve and spinal reflexes, they nevertheless exhibit no evidence of awareness of self or environment, they are unable to interact with others and they show no evidence of sustained, reproducible, purposeful, or voluntary behavioral responses to visual, auditory, tactile or noxious stimuli, and no evidence of language comprehension or expression (The Multi-society Task Force on PVS 1994). Historically this state has been referred to as the “vegetative state” (VS). Recently, it has been proposed that, because the term “VS” might carry pejorative connotations, the more neutral term “unresponsive wakefulness syndrome” might be more appropriate to refer to patients showing a number of clinical signs of unresponsiveness (no conscious behaviors) in the presence of wakefulness (Laureys et al. 2010). Interestingly, some of these unresponsive patients may show inconsistent, but discernible signs of behavioral activity that is more than reflex. These patients are then said to be in a MCS (Giacino et al. 2002). Patients in MCS manifest at least one of the following: purposeful behavior, including movements or affective behavior contingent to relevant environment stimuli which are not due to reflexive activity (visual pursuit or sustained fixation occurring in direct response to moving or salient stimuli, smiling or crying in response to verbal or visual emotional but not neutral stimuli, reaching for objects demonstrating a relationship between object location and direction of reach, touching or holding objects in a manner that accommodates the size and shape of the object, and vocalizations or gestures occurring in direct response to the linguistic content of questions), they can follow

simple commands (gestural or verbal yes/no response, regardless of accuracy), and show intelligible verbalizations.

Measuring states of pathological unconsciousness

When it comes to bedside detection of consciousness, this needs to be inferred via the evaluation of motor activity, with the aim to disentangle reflex from nonreflex behavior (Giacino et al. 2014). To date, the Coma Recovery Scale-Revised (CRS-R) is among the most accurate scales to assess such behaviors (Seel et al. 2010). The scale evaluates 25 arranged items and is organized on 6 subscales, addressing auditory, visual, motor, oromotor, communication, and arousal function. Each item assesses the presence or absence of a specific physical sign that represents the integrity of brain function at one of four levels: generalized, localized, emergent, or cognitively mediated responsiveness. Scoring is based on the presence or absence of specific behavioral responses to sensory stimuli administered in a standardized manner (Giacino et al. 2004). Despite the systematic assessment, the evaluation of nonreflex behavior is not straightforward because patients can show fluctuating vigilance, may suffer from cognitive (e.g. aphasia, apraxia) and/or sensory impairments (e.g. blindness, deafness), from small or easily exhausted motor activity and pain. As a result, the presence of consciousness can be underestimated (Schnakers et al. 2009).

During the past two decades, the diagnosis of disorders of consciousness has been notably facilitated by means of technological modalities, such as electroencephalography (EEG) and functional neuroimaging (Gantner et al. 2013). Although most of such research has concerned patient groups (Laureys and Boly 2007), lately single-patient differentiation by means of automatic algorithms has been achieved (Noirhomme et al. 2017). Data-driven single-patient categorization in MCS and VS/UWS has been performed by combining different markers derived from EEG recordings (Sitt et al. 2014), by estimating values of particular index reflecting reactions of the EEG signal after perturbations with transcranial magnetic stimulation (TMS) (Casarotto et al. 2016) and by functional magnetic resonance imaging (fMRI) during resting state, estimating connectivity within and between brain networks (Demertzi et al. 2015).

Considering the advances in single-patient automatic classification by means of technology-based measurements, the aim of the present article is to (i) bring together state of the art approaches for motor-independent assessment of consciousness levels with EEG, TMS/EEG, and fMRI, (ii) show how each proposed metric translates into clinical practice, and (iii) raise a discussion on the neuroethical aspects of measuring consciousness. It should be clarified that what we intent to measure with these approaches is not necessarily patients' level of consciousness, in the sense that patients can be ordered on the basis of how conscious they are, which may imply that consciousness is by definition graded. In lack of a full understanding of the nature of consciousness, we rather aim to investigate global states of consciousness, namely states which characterize an organism's overall conscious condition. Global states of consciousness are distinguished from each other on cognitive, behavioral, and physiological grounds, such as wakefulness, variable degrees of sedation, dreaming, hypnosis, and absence seizures. As such, an organism can be only in one global state of consciousness at a time avoiding assumptions about the features which are represented in experience (Bayne et al. 2016).

Methodological and technical challenges

The estimation of clinical diagnosis by means of data-driven approaches can be done in several ways (Noirhomme *et al.* 2017). The approaches we discuss here have utilized machine learning techniques which derive a prediction model that is used on individual subjects, and receiver operating curve (ROC) analysis which determines an optimal threshold for separating the categories we are interested in.

When these approaches are applied, several methodological issues emerge (Noirhomme *et al.* 2017). Some of these issues commonly pertain across all adopted methods. In particular, patients' clinical condition may hinder optimal data collection because patients can fluctuate in vigilance, which can be difficult to be monitored by eye tracking and/or electrooculography. As a result, the obtained signal may represent lower level of vigilance and/or be contaminated by noise due to motion artifacts. Additionally, patients can be highly heterogeneous in terms of etiology and chronicity, which may confound the results of the test. Importantly, some of the diagnostic criteria for VS/UWS and MCS do not share international consensus, such is the case with visual fixation, which is considered as a sign of consciousness for some whereas others provide evidence for the opposite (Bruno *et al.* 2010; Naro *et al.* 2016). Such an issue could jeopardize the initial clinical diagnosis with consequences both on the outcome of a test as well as on the used labels which will train a machine learning algorithm.

Each approach, though, addresses the following two issues in different ways: First, it is the issue of establishing the “gold standard”, namely the population on which the metric should be validated. In general, a gold standard is the condition with the highest regarded validity, i.e. it corresponds accurately to what it is supposed to correspond to (Peterson 2016). Based on that, a proposed metric can be validated on healthy controls under the premise that subjective reports are a direct reflection of conscious experience. With this approach, however, the capacity for consciousness can be detected as either present or absent thus providing, at the very best, only a very coarse calibration of the proposed metric. A proposed metric can also be validated on patient population under the premise that the proposed metric will be specific for this group without strong priors coming from healthy conditions. With this approach, however, one runs the risk of reasoning in a circular manner, namely building a predictive algorithm destined to distinguish two patient categories for which the labels are based on behavioral assessment, which we already know is not an optimal measure (Harrison and Connolly 2013). Second, it is the issue of generalization and replicability to avoid overfitting of the used algorithm. Overfitting happens when a model describes noise in the data rather the underlying pattern of interest. As a result, overfitting characterizes very good classification performance on the observed data and very poor performance on unseen data (Arbabshirani *et al.* 2017). Therefore, we are in need of ways which will ensure that the obtained results are relevant for the prediction of newly assessed cases.

EEG signatures of awareness – Jacobo Diego Sitt

EEG recordings represent an important tool to evaluate the state-of-consciousness of brain-injured patients. The main advantage of EEG is the availability in most clinical settings and the applicability in both acute and chronic conditions. The virtual absence of contraindications of this methodology and the fact that it can be easily used at the bedside means that almost all patients can be evaluated using this technique. The

challenge is that EEG data are of a multidimensional space (i.e. several electrodes and time samples) and therefore a direct clinical interpretation of the data is difficult. A way to reduce this complication is to estimate two different types of clinically relevant EEG markers of consciousness: (i) markers of conscious access and (ii) markers of conscious state.

Markers of conscious access are linked to the brain activity associated with the processing of specific external stimuli that can later be reported by the subject. Detecting such activity in patients is important because this would provide direct evidence that the global state-of-consciousness of the patient allows the processing and report of information. These markers are typically extracted using event-related potentials (ERPs) in the framework of different types of stimulation protocols. These protocols are typically designed to evaluate three cognitive processes in patients: stimulus perception (i.e. Bekinschtein *et al.* 2009; Morlet and Fischer 2013), language processing (i.e. Cruse *et al.* 2014; Rohaut *et al.* 2015), or emotional processing (i.e. Perrin *et al.* 2006); see also Harrison and Connolly (2013) for a review of the different protocols used in patients. A difficulty with this approach is that the recoding conditions tend to be rather noisy (i.e. recordings performed in a clinical environment) and the population is highly heterogeneous, e.g. each patient has cerebral lesions in different locations). It is worth noticing that recent research has tried to overcome these limitations and enhance the signal-to-noise ratio of the extracted ERPs by using machine learning. The methodology used in that case is based on the maximization of detecting an activity using multivariate decoding to optimally combine information from several EEG electrodes at different time points (see King *et al.* 2013a) for an example and Noirhomme *et al.* (2017) for a review).

Markers of conscious state aim at detecting conscious information processing independently of the actual content. These markers are typically computed from resting state or nonstimulus-locked activity. The approach here is to quantify the combination of neuronal markers that, according to predictions of current theories of consciousness, are sufficient for the conscious processing of information. To this aim, in a series of works we have introduced novel methods to determine the conscious state of brain-injured patients. In particular, we developed EEG markers to test the hypothesis that cortical information-sharing could distinguish those patients showing signs of conscious behavior (King *et al.* 2013b). In that work we first demonstrated that the introduced measure of information-sharing systematically varied in accordance to the state-of-consciousness, particularly in parietal regions and for long-distance connections. Next, we tested whether the simultaneous evaluation of several EEG markers (either markers of conscious access or conscious state) could provide complementary information to determine patients' state-of-consciousness. To that end, we studied 92 different measures in a cohort of 181 EEG patient recordings (Sitt *et al.* 2014). Figure 1 shows a topographical comparison of a subset of these measures in the different clinical groups. In line with the “multidimensional global states of consciousness” (Bayne *et al.* 2016), the proposed EEG markers quantify different dimensions that are useful to distinguish between different clinical states. This concept of multidimensional EEG evaluation has been recently explored using cognitive ERPs (Sergent *et al.* 2016).

We then examined whether these different EEG measures could be combined to enhance the discrimination performance between the different clinical states-of-consciousness, in particular between VS/UWS and MCS patients. To this aim we trained,

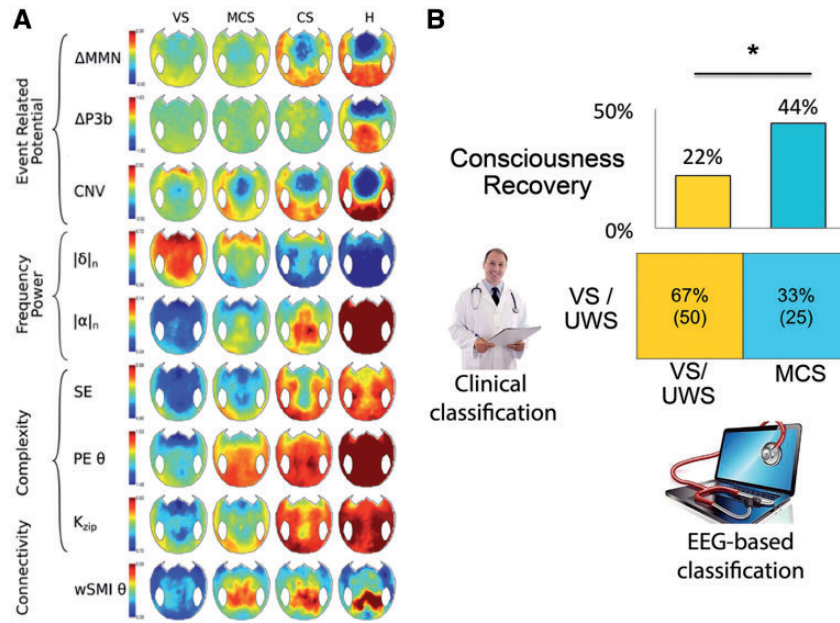


Figure 1. (A) Scalp topography of markers of event-related potentials, spectral decomposition, complexity, and connectivity. The topographical 2D projection (top = front) of each marker is plotted for each state of consciousness (columns). While some markers can better distinguish VS/UWS from MCS (i.e. delta power, permutation entropy, or wSMI connectivity) others are less efficient (i.e. MMN) (B) Comparing clinical and data-driven diagnosis. (B-bottom) Seventy-five patients clinically classified in VS/UWS were evaluated using the EEG-based classification. While in 50 of these recordings the two classifications matched, in 25 the EEG-based system classified the patients as in MCS. (B-top) The bar charts show the clinical outcome of these subgroups of VS/UWS patients. The probability of recovery was higher ($p = 0.02$) for patients classified into a higher state of consciousness than for patients predicted to be actual VS/UWS. AMMN: mismatch negativity, $\Delta P3b$: P300b, $|\delta|_n$: normalized power in delta band, $|\alpha|_n$: normalized power in alpha band, SE: spectral entropy, PE_θ: permutation entropy in theta band, K: Komolgorov-Chaitin Complexity, wSMI_θ: weighted symbolic mutual information (modified from Sitt et al. 2014).

with cross-validation, a support vector machine using the scikit-learn library for machine learning (Pedregosa et al. 2012) and contrasted the performance of the best marker to the optimal combination of markers. The results showed that by using the best marker, the VS/MCS discrimination performance reached an area under the curve (AUC; in classification analysis it determines which of the used models predicts the classes best) value of $71 \pm 4\%$. In contrast, when using all EEG markers the AUC was significantly higher and reached a value of $78 \pm 4\%$ (Sitt et al. 2014). These results suggest that different EEG markers carry partially independent information and that their combination can better inform of patients' state-of-consciousness.

ISSUE 1. How to account for the gold standard

When establishing a clinical tool, the ultimate objective is to be useful for routine clinical practice. Under this premise, our algorithm was trained and tested on data whose labels were based on behavioral evaluation of all patients admitted in neurological service for the evaluation of their conscious state. In that sense, we did not limit our dataset to specific characteristics. Behavioral assessment of consciousness, however, is not entirely accurate to diagnose patients and approximately 15% of the patients behaviorally classified in VS/UWS are expected to retain some conscious processing (Stender et al. 2014). In that case, one should consider that a 100% prediction rate of the clinical labels probably represents overfitting of the used dataset.

In keeping with this limitation, we automatically classified subjects as belonging to a given clinical group, and showed that

EEG markers could be combined to improve the discrimination of the patients' state of consciousness (Sitt et al. 2014). We found that most patients classified as in VS/UWS on both clinical and EEG-based criteria showed no signs of regaining consciousness in the 6 weeks following EEG recording. In clear contrast, for the clinically VS/UWS patients who were classified as MCS based on their EEG activity, the proportion of recovery significantly increased (Fig. 1B). Hence, within a behaviorally indistinguishable group of clinically unresponsive patients, neurophysiological measures provided substantial information about the future improvement of consciousness, suggesting a better functional status at the time of recording.

More recently, and in order to increase clinical application, we implemented a web-based automated solution to provide an EEG-based clinical diagnostic of patients' state-of-consciousness. Its goal is to estimate, in each new patients' recording, the probability that the patient belongs to either the VS/UWS or MCS clinical groups. For this purpose, we developed a flexible and scalable data analysis workflow that automates processing of EEG recordings, the extraction of EEG measures and the communication of results (Fig. 2). The proposed solution is merely based on open source software and is scalable on multiple levels (Engemann et al. 2015).

ISSUE 2. Generalization and replicability

Using this tool, we validated our results using different recording conditions that match the clinical EEG. The classification performance remained unchanged to different conditions, such as when varying the sampling frequency (125 and 250 Hz), the

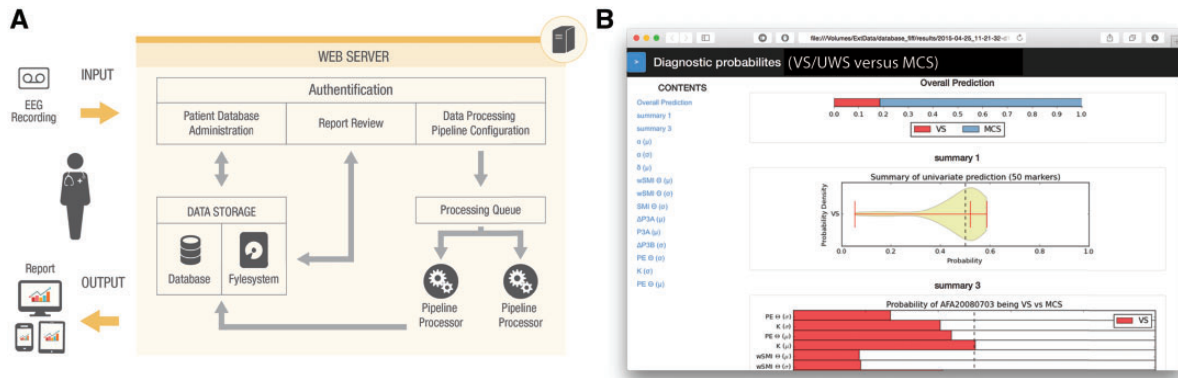


Figure 2. Overview of the automated system for measuring and diagnosing the clinical state of consciousness using EEG makers. Panel (A) illustrates the overall workflow. The operator enters EEG data into the system, the automated pipeline is launched on a web-server, and summary reports are dispatched to the operator. Panel (B) is a screenshot of a diagnostic report that presents the estimated probability of the patient being in a MCS (modified from Engemann et al. 2015).

total number of number of trials (10–500), and channels (8–256). A similar procedure was followed to test the validity under different recording conditions, this time by training the classifier using auditory stimulation recordings and testing on pure resting-state recordings (Engemann et al., n.d., manuscript in preparation).

In summary, these findings show that EEG can produce generalizable predictive models, sensitive to different recording conditions. The capacity of integrating data across a wide range of situations facilitates the evaluation of the state-of-consciousness in both nonspecialist multisite clinical settings and specialised research facilities.

Stratification of unresponsive patients by an independently validated index of brain complexity – Simone Sarasso

Validating and calibrating an objective index of the brain’s capacity for consciousness when behavioral signs of consciousness are unreliable or inconsistent represents a formidable challenge. Here we consider the Perturbational Complexity Index (PCI), a theoretically inspired (Tononi 2004; Tononi et al. 2016) measure that gauges the ability of thalamocortical circuits to integrate information irrespectively of the integrity of sensory processing, motor behavior, and subject’s participation (Casali et al. 2013). PCI gauges the amount of information contained in the integrated response of the thalamocortical system to a direct perturbation. In this context, the notion of information is very different from its classical definition (Shannon 1948). Indeed, in Shannon’s formulation, information is extrinsic and observational; it is assessed from the extrinsic perspective of an observer and it quantifies how accurately input signals can be decoded from the output signals transmitted across a noisy channel. On the contrary, here, information is intrinsic and causal; it is assessed from the intrinsic perspective of a system based on the cause–effect repertoire generated by its internal mechanisms.

Practically, the idea is that brain complexity could be estimated empirically by perturbing the cortex (“zapping”) to engage distributed causal interactions and measuring the information content of the ensuing responses by algorithmic compressibility (“zipping”) (Massimini et al. 2009). This is made possible by the introduction of an electrophysiological technique, based on the combination of navigated TMS and high-

density EEG (hdEEG) (Ilmoniemi et al. 1997). By means of TMS-hdEEG, it is possible to measure the dynamics of the cortical response to TMS, thus inferring on the complexity of the underlying network (Massimini et al. 2009). Indeed, networks in which functional integration is lost will react to TMS with a response that is simple because it is local. On the other hand, networks in which functional specialization is lost will react to TMS with a response that is simple, because it is redundant. Only complex networks, where functional specialization and functional integration are balanced, will react to TMS with a complex response where a large number of integrated areas react in a differentiated way.

In a recent work (Casarotto et al. 2016), we first validated the PCI on a benchmark population of 150 subjects, who could provide a report about the absence or presence of conscious experience either during the TMS-hdEEG assessment or retrospectively upon awakening from sleep (Siclari et al. 2013) and anesthesia (Sanders et al. 2016). From this validation we derived an empirical cutoff (PCI*) that was able to distinguish between the presence and the absence of consciousness as assessed through subjective reports with 100% accuracy (Fig. 3). In particular, PCI was invariably higher in the conscious as compared to the unconscious conditions.

ISSUE 1. Determining the gold standard: reportability

In the lack of an objective independent measure of consciousness, relying on subjective reports – either immediate or delayed – currently represents a reliable, although very coarse, way to assess the presence/absence of conscious experience (Noreika et al. 2011) and therefore to validate objective measures of consciousness (Sanders et al. 2016). As mentioned above, considering subjective reports as the true state of affairs, if a brain-based test of consciousness is positive in subjects who are fully unresponsive at the time of measurement but provide a delayed report of a vivid dream upon awakening, the result should be considered a true positive. Here, the benchmark population included data from subjects under ketamine anesthesia and REM (rapid-eye-movement) sleep conditions in order to specifically validate PCI in subjects who are, at the same time, conscious but behaviorally unresponsive and disconnected from the external environment (Siclari et al. 2013). Moreover, the benchmark population also included a large group of responsive brain-injured patients (with locked-in syndrome, subcortical

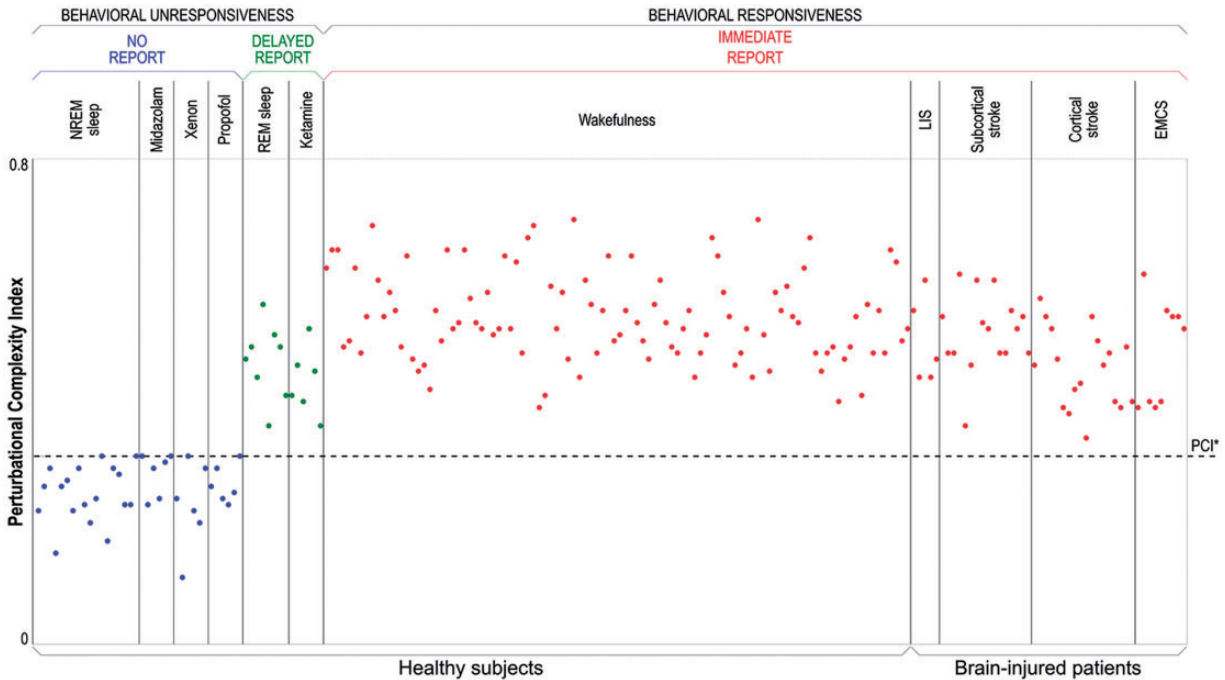


Figure 3. The PCI differentiates between conscious and unconscious subjects. Each circle represents the PCI value computed from the cortical responses to TMS of one stimulation site. PCI values are computed from TMS-evoked potentials recorded in healthy subjects and conscious brain-injured patients during different conditions. Individuals are grouped by condition, and within each condition are sorted by increasing age. During non-rapid eye movement (NREM) sleep and anesthesia with midazolam, xenon, and propofol, subjects were behaviorally unresponsive and did not provide any report upon awakening. During dreaming and ketamine anesthesia, subjects were behaviorally unresponsive but provided subjective reports upon awakening. During wakefulness, both healthy subjects and conscious brain-injured patients (with locked-in syndrome-LIS, subcortical and cortical stroke, patients who emerged from the MCS-EMCS) could immediately report their subjective experience. The dashed horizontal line represents the empirical cutoff PCI^* obtained by applying ROC curve analysis using the presence/absence of report as the gold standard (modified from Casarotto et al. 2016).

and cortical stroke, patients who emerged from the MCS). This latter feature allowed to account for the presence of brain lesions in the validation of PCI, which is ultimately aimed at a target population of MCS and VS/UWS patients with severe neurological damage. Then, we applied this externally validated PCI cutoff to evaluate patients with disorders of consciousness. Slicing through the PCI distribution of a MCS cohort ($n=38$) with the empirical cutoff derived from the benchmark population, we further challenged the sensitivity of PCI in detecting patients showing minimal but unequivocal behavioral signs of consciousness. Crucially, PCI was higher than PCI^* in 36 of 38 patients, leading to an unprecedented sensitivity value of 94.7% in detecting minimal signs of consciousness (Fig. 4).

Finally, by applying the externally validated cutoff to a cohort of 43 VS/UWS patients, we found that 9 unresponsive patients had a PCI value higher than PCI^* , raising the issue of whether these are false or true positives. Since PCI was always higher than PCI^* only when consciousness was present and never when consciousness was absent across the 200 measurements pertaining to the benchmark population, it is parsimonious to assume that these high-complexity unresponsive patients may retain a capacity for consciousness that is not expressed in behavior.

Taken together, using subjective reports as the gold standard, we found that PCI optimally discriminates between the conscious and the unconscious conditions in a benchmark population, irrespectively of behavioral unresponsiveness and the presence of brain lesions. Testing our metric on an independent dataset of subjects able to report – immediately or

retrospectively – the presence/absence of consciousness clearly offers the advantage of avoiding the issue of reasoning in a circular manner, i.e. assessing the performance of a measure of consciousness using the target population as a gauge for the accuracy of the measure in the absence of a veridical benchmark of the true state-of-affairs. Furthermore, the inclusion in our benchmark population of subjects that were behaviorally unresponsive and disconnected from the external environment while conscious, as well as conscious brain-injured patients, allows for the best available approximation of the challenging conditions presented by the unresponsive brain-injured patients.

ISSUE 2. Generalization and replicability

The sensitivity obtained by means of an independent validation of PCI, confirmed on MCS patients, allows for a sufficiently unbiased, and therefore generalizable, condition to slice through a population of patients in which the absence of behavioral signs of consciousness *per se* cannot be considered a proof of the absence of consciousness. More importantly, our approach allows for the identification of three subgroups of VS/UWS patients with both pathophysiological and management practical implications. Specifically, we observed (i) a no-response VS/UWS subgroup in which TMS targeted over different cortical areas failed to engage any significant cortical response (black traces in Fig. 4), (ii) a low-complexity subgroup where TMS triggered a local and stereotypical positive-negative response (blue traces in Fig. 4) similar to the one observed in healthy controls during

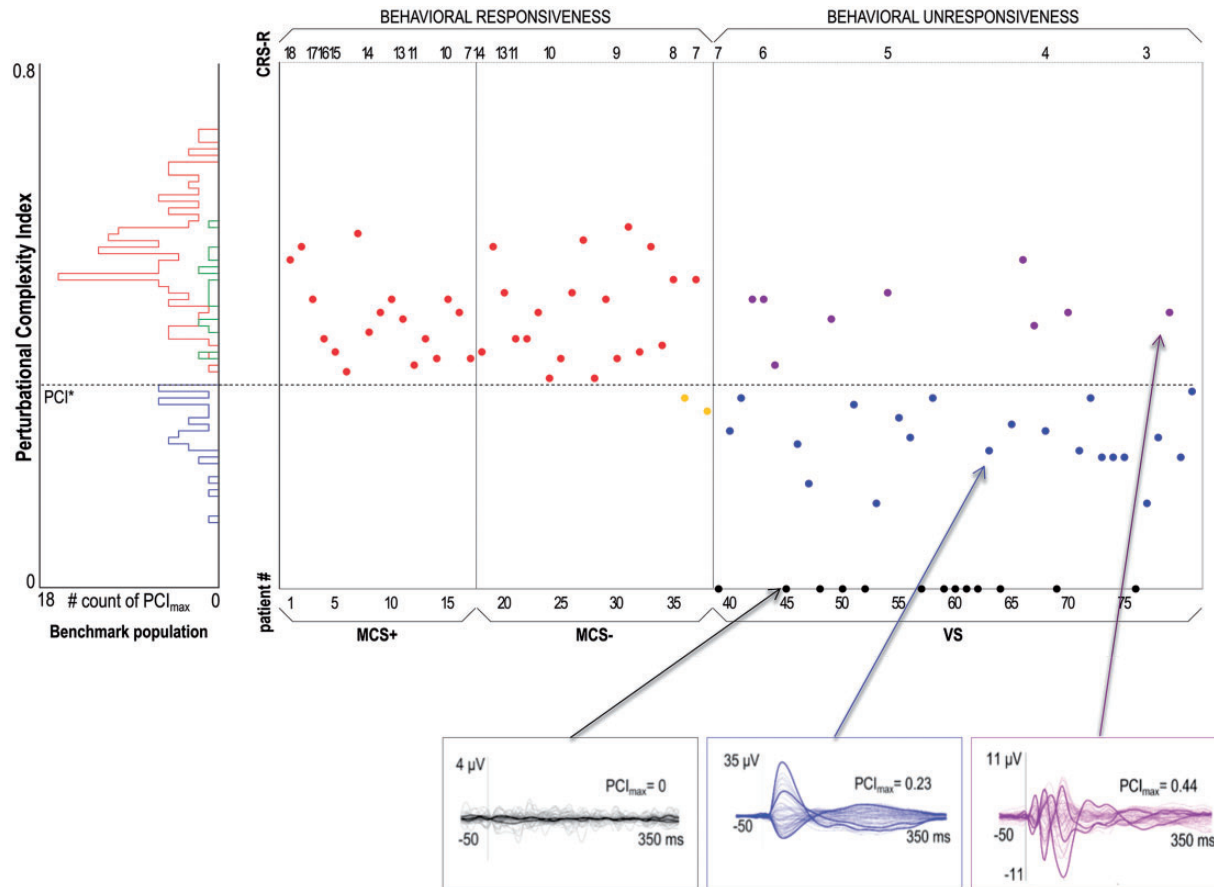


Figure 4. PCI^* performance tested on behaviorally responsive and unresponsive patients. The histogram on the left summarizes the distribution of PCI in the benchmark population, specifically obtained in the absence of subjective report (blue) and in the presence of subjective report (delayed, green; immediate, red) conditions. The dashed horizontal line highlights the optimal cutoff (PCI^*) computed from ROC curve analysis on the benchmark population. The scatter plot on the right shows all the PCI values obtained in minimally conscious state (MCS+/MCS-) and unresponsive patients (VS). Within each diagnostic group, patients are sorted by the CRS-R total score in decreasing order. The three boxes at the bottom show the average TMS-evoked potentials (all channels superimposed, with three illustrative channels highlighted in bold), together with the PCI values for three representative unresponsive patients (VS) with individual value of PCI respectively=0 (left, in black), lower than PCI^* (center, in blue), and higher than PCI^* (right, in purple) (modified from Casarotto et al. 2016).

unconscious NREM sleep and anesthesia, and (iii) a high complexity subgroup, in which TMS engaged a rapidly changing and spatially differentiated cortical response (purple traces in Fig. 4), similar to the one observed in MCS patients and in responsive wakefulness or unresponsive conscious controls, i.e. during REM sleep and ketamine anesthesia.

Practically, by approximating an optimal trade-off between sensitivity and specificity, the PCI test may represent an important step within a hierarchical diagnostic flow of patients with DOC. For example, one may apply the PCI after a first screening with tests characterized by high sensitivity, such as positron emission tomography assessment of cortical metabolic rates (Stender et al. 2014) in search of preserved cortical and subcortical metabolic activations that may have escaped the TMS probing. Then, patients with $PCI > PCI^*$ could be selected to confirm the presence of covert consciousness through more demanding tests characterized by maximal specificity, such as functional MR imaging active paradigms. Finally, patients with $PCI < PCI^*$ should be directed toward neuromodulation with medications or brain stimulation techniques (Fridman and Schiff 2014) aimed at restoring complex patterns of activity.

Single patients are classified in discrete diagnostic categories based on fMRI default brain activity – Athena Demertzi

Even when the mind is free to rest and do nothing, spontaneous cognition tends to gravitate toward thoughts and feelings. This means that the apparently idle brain is, instead, constantly active (Gusnard and Raichle 2001). The resting state paradigm is particularly appealing for clinical applications because it does not require sophisticated experimental setup and surpasses the need for subject's collaboration (Demertzi and Whitfield-Gabrieli 2016). Therefore, it is suitable for studying subjects who are unable to communicate in a functional manner, including patients suffering from disorders of consciousness.

Analyses of fMRI data during resting state indicate that the brain can be organized in reproducible networks of cognitive significance (Smith et al. 2009; Laird et al. 2011). The most robustly identifiable network is the default mode network (DMN) which classically encompasses posterior cingulate cortex and adjacent precuneus as well as anterior cingulate cortex and mesiofrontal areas (Raichle et al. 2001). When DMN functional connectivity (i.e. the temporal synchronization between

spontaneous BOLD signal; Biswal et al. 1997) was estimated in patients with disorders of consciousness, it was shown that it was relatively preserved in patients in MCS and significantly reduced in patients in a VS/UWS; in contrast, there was preserved connectivity in healthy controls and in a patient with locked-in syndrome (Vanhaudenhuyse et al. 2010). Furthermore, DMN connectivity was absent in a brain dead patients (Boly et al. 2009). Such results suggest that DMN functional connectivity correlates with consciousness states and, at least to a certain degree, can be used as a proxy to study residual cognitive function in these patients.

The quantification of fMRI resting state functional connectivity in patients can be challenging. When employing data-driven independent component analysis (ICA), for example, a dataset is divided into maximally different statistical components and, thus, it is able to isolate cortical connectivity maps from non-neuronal signals (Beckmann et al. 2005). This method allows us to evaluate and compare the coherence of activity in multiple distributed voxels and hence identify more brain networks other than the DMN (Heine et al. 2012). However, ICA does not provide any ordering of the independent components, which can make the data representation difficult to interpret. A common solution to identify the network of interest is to calculate the fit between the components' spatial pattern and the pattern of a pre-defined template which represents the network (Greicius et al. 2004). Importantly, in cases of deformed brains, results should be interpreted with caution. This is because in the extreme conditions, where a patient shows only one component of neuronal origin, the use of a single template-matching method would lead to identify this component as the network we are investigating each time (Heine et al. 2012; Soddu et al. 2012). In that case, one endangers to identify an unsuitable component as the network of interest. We overcame this issue by means of a multiple template-matching procedure, where the goodness of fit was calculated among several pairs of component-template comparisons. The pair with the maximal goodness of fit value was eventually selected. Hence, the chances to select the "proper" network of interest increased (Demertzi et al. 2014). As a second network selection criterion, we isolated the components which were not characterized by neuronal properties based on the information of how the BOLD signal fluctuates. For the neuronality check, non-neuronal were those components showing, among others, activation/deactivation in peripheral areas, in the cerebrospinal fluid and white matter and which were characterized by high-frequency fluctuations (>0.1 Hz). Conversely, neuronal were considered those components when at least 10% of the activations/deactivations were found as gray matter clusters. With this "neuronality test", we ensured that the selected networks were within healthy boundaries both in spatial and temporal terms. Following this approach, we found that, compared to controls, patients showed fewer neuronal components. Also, four networks were less identifiable in the clinical population, namely the left and right executive control network, the DMN and the auditory network. Supervised machine learning was able to separate healthy controls from patients with 85% accuracy based on information about the neuronality of the DMN and the auditory network. The separation between MCS and VS/UWS, however, was more challenging (Demertzi et al. 2014). These results imply that systems-level resting state fMRI can be informative of consciousness states. However, the discriminative characteristic (here the neuronality of the networks) appeared as a non-sensitive feature to capture the separation between these two

classes. Also, due to the selection criteria for identifying the networks, many patient data had to be excluded.

In order to increase power and determine a more sensitive metric across all patients, we sought to go hypothesis-driven by means of seed-based correlation analysis. We investigated seed-based functional connectivity in three large-scale networks (DMN, frontoparietal, and salience) and three sensory networks (auditory, sensorimotor, and visual) (Demertzi et al. 2015). We first found that the clinical scores of the CRS-R correlated with functional connectivity of all networks, highlighting their contribution to consciousness state (note: as a control network we included the cerebellum which did not show connectivity changes as a function of CRS-R scores). We further aimed at determining the capacity of each network to differentiate between patients in MCS and VS/UWS. We found group-level differences in all networks. Moving toward single-patient classification, we identified that all networks were able to differentiate patients in MCS and VS/UWS with high accuracy ($>86\%$). Such accuracy could be partially attributed to the fact that the network ranking was based on features extracted from the same population for which between-group differences were already known. To avoid double-dipping, we validated the most highly ranked network on two independent clinical datasets and across healthy controls scanned in different centers. Single-patient classification was performed based on the connectivity strength of the "auditory network", which was ranked most highly among all studied systems in separating MCS and VS/UWS patients. Based on this network's connectivity, 20 of the 22 new patients were classified congruently, namely the clinical diagnosis matched the classifier's decision (Fig. 6A). It should be noted that this network was referred to as "auditory" for the sake of consistency with the literature. In fact, apart from temporal cortices, this network encompasses regions in occipital cortex, pre- and postcentral areas and insula. Hence, the discriminative feature of this system was connectivity in areas beyond audition-related, such as occipital and bilateral opercular/insular cortex (Fig. 5).

This crossmodal connectivity has been previously identified in normal conscious subjects during rest (Eckert et al. 2008). Also in healthy subjects, preserved fMRI activity in temporal and occipital areas has been shown during mental counting of auditory temporal irregularities, only when they were attentive and aware of the auditory violations (Bekinschtein et al. 2009). Inversely, when subjects were scanned under pharmacologically induced anesthesia, there was decreased crossmodal interaction, which was restored after recovery of consciousness (Boveroux et al. 2010). Crossmodal connectivity is considered relevant for multisensory integration (Clavagnier et al. 2004), i.e. a facilitator for top-down influences of higher order regions to create predictions of forthcoming sensory events (Engel et al. 2001). Interestingly, top-down connections were found to be present in patients in MCS but absent in patients in VS/UWS (Boly et al. 2011). Taken together, our results indicate that connectivity in occipital and bilateral opercular/insular cortex captures a dimension of the MCS conscious state which goes beyond the classical DMN connectivity and is able to separate these patients from those in VS/UWS.

ISSUE 1. Determining the gold standard: multimodal assessment

In order to ensure that the proposed metric reflects the construct which is supposed to reflect (i.e. a connectivity pattern




Figure 5. The fMRI functional connectivity pattern which was used as a feature to differentiate MCS from VS/UWS patients. Areas in occipital and bilateral opercular/insular cortex were more functionally connected in patients in MCS compared to those in VS/UWS. A support vector machine classifier was trained on this pattern with data collected in one center, and it was generalized on an independently dataset of patients evaluated at different sites. The algorithm's classification was congruent with the clinical evaluation of consciousness in 20/22 assessed patients.

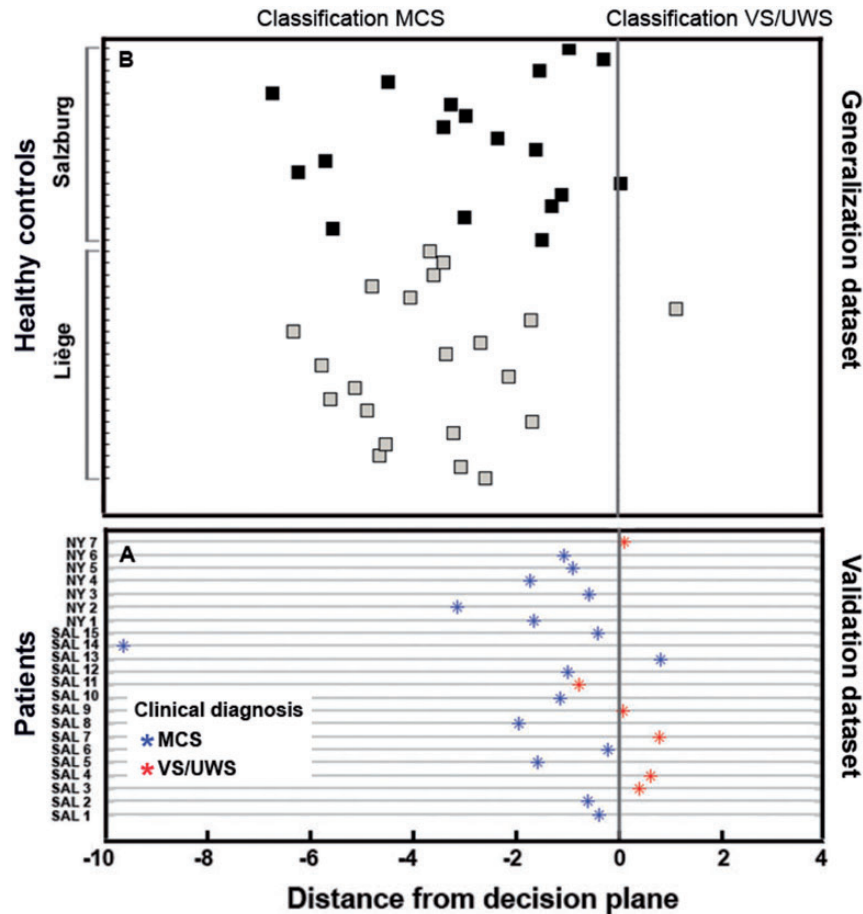


Figure 6. Validation and generalization of the fMRI resting state classifier for the separation between patients in minimally conscious state (MCS) and in vegetative state/unresponsive wakefulness syndrome (VS/UWS). Panel (A) The classifier was validated on a dataset of patients assessed in two sites (SAL: Salzburg, $n=15$; NY: NewYork, $n=7$) other than the center where the training dataset were collected (Liège, $n=45$). The classifier's decision was congruent with the clinical evaluation of MCS (blue) and VS/UWS (red) in 20/22 cases. Panel (B) The classifier was generalized on a set of healthy controls scanned in Liège and Salzburg ($n=39$; no data were available for New York). The majority of healthy controls were classified as in MCS (left-hand side of the decision plane), ensuring the classifier's ability of capturing properties of a fully conscious state (modified from Demertzi et al. 2015)

discriminating the conscious state of MCS from that of VS/UWS), a well-defined diagnostic baseline was essential. To that end, the classifier's trainset included data of patients who underwent a multimodal assessment of their conscious state by means of behavioral evaluation and neuroimaging. Patients were repeatedly examined with standardized clinical tools (an average of six assessments per patient with the CRS-R by multiple assessors). This clinical diagnosis was further cross-validated with FDG-PET imaging, which has been shown to have high sensitivity in identifying patients in MCS (Stender et al. 2014). Therefore, patients with an ambiguous profile on clinical assessment and neuroimaging data were not included in the analysis. In that way, we ensured that the training set included patients of both classes, whose profile was in accordance with behavioural and cognitive criteria of the VS/UWS and MCS.

ISSUE 2. Generalization and replicability

To test the robustness of the classifier, we evaluated whether it generalizes to healthy control subjects scanned in two centers (Fig. 6B). The hypothesis was that healthy controls would belong to the class of MCS, hence ensuring the classifier's ability of capturing properties of a fully conscious state. The hypothesis was verified in all but two healthy individuals. It should be noted that the present classifier was created on patients and subjects scanned in an awake condition. In this case, its clinical translation needs to be restricted to those patients scanned under similar conditions. Similarly, patients who received sedatives to minimize motion in the scanner were further excluded. The reason to exclude sedated patients was because of our limited understanding of the potential effect of anaesthetics on network connectivity (Heine et al. 2012). We here recognize the importance of increasing the classification power for patients scanned after receiving anaesthetics, given that many patients undergo anaesthesia not only to restrict scanner motion but also for neuroprotective reasons. In the future, a more generalizable algorithm applicable to patients scanned under anaesthetics (Kirsch et al. 2017) is expected to shed light on both the pathophysiology and the drug-related variance leading to consciousness alterations in patients.

On measuring consciousness: ethical reflections – Wim Pinxten

In principle, consciousness is impartial: humans are either conscious or not. Although a person may be conscious "of" more than another person, this does not imply that the person is "more" conscious (Bayne et al. 2016). Also the way in which we try to infer consciousness, e.g. by checking wakefulness and awareness, does not affect this impartiality. In practice, however, disorders of consciousness confront us with strong variations in the manifestation of consciousness among patients, which are translated into states of consciousness. It is not written in stone how such variations in the expression of consciousness should be interpreted and whether or how they could be translated into states or levels of consciousness. For example, some patients show more clinical signs of consciousness than others, according to which they are classified as either MCS or VS/UWS. But such clinical signs can be misleading, and high rates of misdiagnosis have been reported (Schnakers et al. 2009; Weijer et al. 2014). An extreme example is the misdiagnosis of patients with locked-in syndrome, in whom consciousness is intact but motor function has almost completely been lost (Laureys et al. 2005). New technologies, including advanced

neuroimaging, have extended the scope of parameters that can be employed to explore variations in the state of consciousness, and will continue to do so in the future (Laureys and Boly 2007). This opens up potential for further refining states of consciousness, or even for developing "consciousness meters".

Determining the extent to which patients are capable of conscious experiences, now and in the future, is very relevant for caregivers and decision-makers. For example, estimates of the remaining capability of conscious experiences (further referred to as residual consciousness) shape attitudes towards the appropriate clinical approach of patients with DOC (including pain management and end-of-life decisions). Failing to identify consciousness, however, could influence medical care in terms of limiting rehabilitation efforts, pain management, and end-of-life decisions (Demertzi et al. 2009b, 2011, 2013; Jox et al. 2012; Peterson et al. 2015). Distinguishing between different states of consciousness, hence, comes with important ethical issues.

Any quantification of residual consciousness is man-made, and therefore the meaning attributed to measurements of consciousness requires adequate ethical justification. In other words, measuring consciousness is a profoundly normative enterprise. In this article, we will adopt an applied ethics perspective, and reflect on ethical challenges related to the clinical management of patients. As such, we will focus on two ethical challenges: knowing what one is measuring and the validity of measurements.

When exploring the ethical issues in measuring consciousness, a first challenge lies in knowing what one is measuring. In clinical practice, states of consciousness (VS/UWS or MCS) are commonly determined by observing signs of awareness and wakefulness. While such observations are practical, it is unclear to what extent they provide a representative estimate of residual consciousness. As for now, much remains uncertain, which comes at the risk of reductionist approaches. As it is the essence of scientific endeavor to reduce subjectivity to the minimum, we need to avoid reductionism to the maximum and pursue a valid and reliable stratification of states of consciousness. However, the current need to make decisions on the clinical management of patients does not allow to merely wait until emerging technologies and scientific insights will strongly reduce the level of uncertainty. Pragmatically, it therefore makes sense to get the best out of the current state of the art, even when the current determination of states of consciousness is not infallible at all. Upon condition that the benefits (e.g. of a more accurate diagnosis and prognosis) outweigh the risks (e.g. false positive/negative observations of consciousness, and the related risk of devaluing human persons building on an uncertain or false observation), it might therefore be better to have something rather than nothing, and to keep considering the ever-richer information on the functioning of the brain in the definition and assessment of states of consciousness.

This brings us to the second challenge, which lies in the validity of measurements. Machines can quantify certain traits of patients, but since data do not speak up for themselves it is clinicians who have to interpret the significance of quantitative differences. But what if different measurements and/or interpretations do not match? It has been suggested that in the future classifiers should take into account the ensuing clinical cost when predicting one class over the other (Noirhomme et al. 2017). What level of validity is required for adequate patient management? It has been suggested that from the patient's perspective, a false positive observation of consciousness may not be as burdensome as being falsely identified as unresponsive (Jox et al. 2012; Peterson et al. 2015). Additionally, Shea and

Bayne (2010) argue in this respect that a new consciousness metric should not be checked for validity against a single gold standard, simply because there is still no consensus as to the nature of consciousness and how this phenomenon is best measured (Peterson 2016). Rather, accumulative information needs to come from multiple independent sources (Seth et al. 2008). Such mode of reasoning, i.e. reasoning by conciliation, necessitates to test the degree to which results of multiple tests conciliate or deviate with a patient's broader assessment (Peterson 2016). Naturally, reasoning by conciliation can be challenged by the presence of discordant results (which may result from practical problems during assessment or actually reflecting neurobiological differences), by the selection of what to include and exclude as sources of information and by disentangling actual preserved consciousness from mere statistical associations among different signals. This rationale has been integrated in a multidimensional model which tracks the recovery of aspects of consciousness after brain injury (Bayne et al. 2016). Recent experimental data further support this model, by showing that EEG evaluation covering different dimensions of cognitive processing increased diagnostic sensitivity: the presence of high-level effects distinguished between MCS and VSUWS, while the presence of low-level effects was similar in both groups. Such results show that a multidimensional evaluation not only probed patients' consciousness but also established a more general and nuanced profile of the residual cognitive capacities (Sergent et al. 2016). In a third-person perspective, false positives may instil false hopes about a positive clinical outcome in patients' caregivers, and also in socioeconomic terms false positive cases have associated costs, such as intensified efforts to exhaust available treatment options (Kahane and Savulescu 2009; Jox et al. 2012).

With consciousness being surrounded by ambiguity and its measurements being imperfect and (still) of uncertain significance, how can we consider the appropriate way forward? First, we lack convincing arguments against measuring, while the advantages of adding up to scientific insight and providing best available information on brain functioning for patients that show no clear clinical signs of consciousness are valuable. Second, notwithstanding these advantages, one must keep in mind that any level of consciousness is essentially a human construct that is open to bias and misinterpretation. Taking into account that decisions on patients are by definition being made by proxy decision-makers, prudence to not overestimate the value of measurements is essential. Therefore, continuous reflection, also on the normative implications, should be in place when building on measured levels of consciousness for clinical management or end-of-life decisions.

Conclusions

When predicting whether a patient belongs to a clinical class, it is possible to be confronted with the disagreement between the calculated thresholds and clinical labels. In that case, one of two sources is less reliable about the current status of the patient. It can be that the metric outperforms the physician by capturing subclinical features of eventual recovery of consciousness. For example, the EEG-based classification showed that while the machine and the physician matched in diagnosis of 50 patients, 25 clinically diagnosed VS/UWS patients were classified as in MCS by the EEG-based system. Interestingly, these patients showed a higher rate of recovery (Sitt et al. 2014). Similarly, for PCI, although not intended as a prognostic marker, the outcome was more favorable in the high-complexity

VS/UWS subgroup (six out of nine patients transitioned to a behavioral MCS within 6 months), whereas such transition was observed only in 5 of 21 low-complexity patients and in none of the no-response subgroup (Casarotto et al. 2016). Finally, for the resting state fMRI connectivity, the patient who was misclassified as in MCS had a profile of VS/UWS on the day of scan but evolved to MCS 38 days later; the other patient was misclassified as VS/UWS but had a clinical profile of MCS on the day of scanning based on the presence of localization to noxious stimulation but this behavior could not be elicited in any other evaluations (Demertzi et al. 2015). On the other hand, it can be that the biomarker predicted the class incorrectly. This is not infrequent as the metrics rely on data representing neural function instead of direct neural signals. Also, machines can quantify certain traits of patients. One way to gain confidence in what we measure is to utilize probabilistic predictions in place of binary distinctions. Like that, clinical reality may be more representative. Additionally, the use of accumulative evidence stemming from multiple nonoverlapping assessments with different modalities, which are sensitive and specific in detecting the capacity for conscious processing, are expected to boost our level of confidence in single-patient diagnostics. As regards the diagnostic procedure, to date we are in need of a framework which will determine the application order of each technological modality, balancing the availability of these technologies at each research/clinical site, each technology's sensitivity and specificity characteristics and the underlying clinical assumptions about a patient's conscious state. Such a framework is expected to ameliorate clinical management and further inform on the critical patterns of (un)consciousness.

Acknowledgements

This work was supported by the James S. McDonnell Foundation Scholar Award 2013, the Institut national de la santé et de la recherche médicale (INSERM), the EU 2020 Research and Innovation Programme under Grant Agreement No. 720270 (HBP SGA1), the Institut du Cerveau et de la Moelle Epinière (ICM, Paris) and the STIC-AmSud grant (Complex) awarded to JS.

References

- Arbabshirani MR, Plis S, Sui J et al. Single subject prediction of brain disorders in neuroimaging: promises and pitfalls. *Neuroimage* 2017;**145**:137–65.
- Baars BJ. Consciousness. *Scholarpedia* 2015;**10**:2207.
- Bayne T, Hohwy J, Owen AM. Are there levels of consciousness? *Trends Cogn Sci* 2016;**20**:405–13.
- Beckmann CF, DeLuca M, Devlin JT et al. Investigations into resting-state connectivity using independent component analysis. *Philos Trans R Soc Lond B Biol Sci* 2005;**360**:1001–13.
- Bekinschtein TA, Dehaene S, Rohaut B. Neural signature of the conscious processing of auditory regularities. *Proc Natl Acad Sci U S A* 2009;**106**:1672–77.
- Biswal BB, Van Kylen J, Hyde JS. Simultaneous assessment of flow and BOLD signals in resting-state functional connectivity maps. *NMR Biomed* 1997;**10**:165–70.
- Boly M, Garrido MI, Gosseries O et al. Preserved feedforward but impaired top-down processes in the vegetative state. *Science* 2011;**332**:858–62.
- Boly M, Tshibanda L, Vanhaudenhuyse A et al. Functional connectivity in the default network during resting state is preserved in a vegetative but not in a brain dead patient. *Hum Brain Mapp* 2009;**30**:2393–400.

- Boveroux P, Vanhaudenhuyse A, Bruno MA et al. Breakdown of within- and between-network resting state functional magnetic resonance imaging connectivity during propofol-induced loss of consciousness. *Anesthesiology* 2010;113:1038–53.
- Bruno M, Vanhaudenhuyse A, Schnakers C et al. Visual fixation in the vegetative state: an observational case series PET study. *BMC Neurol* 2010;10. <https://www.ncbi.nlm.nih.gov/pubmed/23946194>
- Casali AG, Gosseries O, Rosanova M et al. A theoretically based index of consciousness independent of sensory processing and behavior. *Sci Transl Med* 2013;5:198ra105–198ra105.
- Casarotto S, Comanducci A, Rosanova M et al. Stratification of unresponsive patients by an independently validated index of brain complexity. *Ann Neurol* 2016;80:718–29.
- Clavagner S, Falchier A, Kennedy H. Long-distance feedback projections to area V1: implications for multisensory integration, spatial awareness, and visual consciousness. *Cogn Affect Behav Neurosci* 2004;4:117–26.
- Cruse D, Beukema S, Chennu S et al. The reliability of the N400 in single subjects: implications for patients with disorders of consciousness. *NeuroImage Clin* 2014;4:788–99.
- Demertzi A, Antonopoulos G, Heine L et al. Intrinsic functional connectivity differentiates minimally conscious from unresponsive patients. *Brain* 2015;138:2619–31.
- Demertzi A, Gomez F, Crone JS et al. Multiple fMRI system-level baseline connectivity is disrupted in patients with consciousness alterations. *Cortex* 2014;52:35–46.
- Demertzi A, Ledoux D, Bruno M-A et al. Attitudes towards end-of-life issues in disorders of consciousness: a European survey. *J Neurol* 2011;258:1058–65.
- Demertzi A, Liew C, Ledoux D et al. Dualism persists in the science of mind. *Ann N Y Acad Sci* 2009a;1157:1–9.
- Demertzi A, Racine E, Bruno A et al. Pain perception in disorders of consciousness: neuroscience, clinical care, and ethics in dialogue. *Neuroethics* 2013;6:37–50.
- Demertzi A, Schnakers C, Ledoux D et al. Different beliefs about pain perception in the vegetative and minimally conscious states: a European survey of medical and paramedical professionals. *Prog Brain Res* 2009b;177:329–38.
- Demertzi A, Whitfield-Gabrieli S. Intrinsic brain activity and consciousness. In: Laureys S, Gosseries O, Tononi G (eds), *The Neurology of Consciousness*, 2nd edn. Elsevier, 2016, 95–105.
- Eckert MA, Kamdar N V., Chang CE et al. A cross-modal system linking primary auditory and visual cortices: evidence from intrinsic fMRI connectivity analysis. *Hum Brain Mapp* 2008;29:848–57.
- Engel AK, Fries P, Singer W. Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 2001;2:704–16.
- Engemann D, Raimondo F, King J-R et al. Automated measurement and prediction of consciousness in vegetative and minimally conscious patients. *ICML Work Stat Mach Learn Neurosci* (Stammlins 2015).
- Engemann D, Raimondo F, King J-R et al. (n.d.) Robust quantification of consciousness from clinical EEG.
- Fridman EA, Schiff ND. Neuromodulation of the conscious state following severe brain injuries. *Curr Opin Neurobiol* 2014;29:172–77.
- Gantner IS, Bodart O, Laureys S et al. Our rapidly changing understanding of acute and chronic disorders of consciousness: challenges for neurologists. *Future Neurol* 2013;8:43–54.
- Giacino JT, Ashwal S, Childs N et al. The minimally conscious state: Definition and diagnostic criteria. *Neurology* 2002;58:349–53.
- Giacino JT, Fins JJ, Laureys S et al. Disorders of consciousness after acquired brain injury: the state of the science. *Nat Rev Neurol* 2014;10:99–114.
- Giacino JT, Kalmar K, Whyte J. The JFK Coma Recovery Scale – Revised: Measurement. 2004;85:2020–29.
- Greicius MD, Srivastava G, Reiss AL et al. Default-mode network activity distinguishes Alzheimer’s disease from healthy aging: evidence from functional MRI. *Proc Natl Acad Sci U S A* 2004;101:4637–42.
- Gusnard DA, Raichle ME. Searching for a baseline: Functional imaging and the resting human brain. *Nat Rev Neurosci* 2001;2:685–94.
- Harrison AH, Connolly JF. Finding a way in: a review and practical evaluation of fMRI and EEG for detection and assessment in disorders of consciousness. *Neurosci Biobehav Rev* 2013;37:1403–19.
- Heine L, Soddu A, Gómez F et al. Resting state networks and consciousness: Alterations of multiple resting state network connectivity in physiological, pharmacological, and pathological consciousness states. *Front Psychol* 2012;3:1–12.
- Ilmoniemi RJ, Virtanen J, Ruohonen J et al. Neuronal responses to magnetic stimulation reveal cortical reactivity and connectivity. *Neuroreport* 1997;8:3537–40.
- Jennett B, Plum F. Persistent vegetative state after brain damage: a syndrome in search of a name. *Lancet* 1972;299:734–37.
- Jox RJ, Bernat JL, Laureys S et al. Disorders of consciousness: responding to requests for novel diagnostic and therapeutic interventions. *Lancet Neurol* 2012;11:732–38.
- Kahane G, Savulescu J. Brain damage and the moral significance of consciousness. *J Med Philos* 2009;34:6–26.
- King JR, Faugeras F, Gramfort A et al. Single-trial decoding of auditory novelty responses facilitates the detection of residual consciousness. *Neuroimage* 2013a;83C:726–38.
- King JR, Sitt JD, Faugeras F et al. Information sharing in the brain indexes consciousness in noncommunicative patients. *Curr Biol* 2013b;23:1914–19.
- Kirsch M, Guldenmund P, Ali Bahri M et al. Sedation of patients with disorders of consciousness during neuroimaging: effects on resting state functional brain connectivity. *Anesth Analg* 2017;124:588–98.
- Laird AR, Fox PM, Eickhoff SB et al. Behavioral interpretations of intrinsic connectivity networks. *J Cogn Neurosci* 2011;23:4022–37.
- Laureys S, Boly M. What is it like to be vegetative or minimally conscious? *Curr Opin Neurol* 2007;20:609–13.
- Laureys S, Celesia GG, Cohadon F et al. Unresponsive wakefulness syndrome: a new name for the vegetative state or apallic syndrome. *BMC Med* 2010;8:68.
- Laureys S, Pellas F, Van Eeckhout P et al. The locked-in syndrome: what is it like to be conscious but paralyzed and voiceless? *Prog Brain Res* 2005;150:495–611.
- Massimini M, Boly M, Casali A et al. A perturbational approach for evaluating the brain’s capacity for consciousness. *Prog Brain Res* 2009;177:201–14.
- Morlet D, Fischer C. MMN and novelty P3 in coma and other altered states of consciousness: a review. *Brain Topogr* 2014;27:467–79.
- Naro A, Leo A, Buda A et al. Do you see me? The role of visual fixation in chronic disorders of consciousness differential diagnosis. *Brain Res* 2016;1653:59–66.
- Noirhomme Q, Brecheisen R, Lesenfants D et al. “Look at my classifier’s result”: Disentangling unresponsive from (minimally) conscious patients. *Neuroimage* 2017;145:288–303.

- Noreika V, Jylhänkangas L, Móró L et al. Consciousness lost and found: subjective experiences in an unresponsive state. *Brain Cogn* 2011;**77**:327–34.
- Pedregosa F, Varoquaux G, Gramfort A et al. Scikit-learn: Machine Learning in Python. *Mach Learn* 2012;**12**:2825–30.
- Perrin F, Schnakers C, Schabus M et al. Brain response to one's own name in vegetative state, minimally conscious state, and locked-in syndrome. *Arch Neurol* 2006;**63**:562–69.
- Peterson A. Consilience, clinical validation, and global disorders of consciousness. *Neurosci Conscious* 2016;**2016**:1–9. <https://academic.oup.com/nc/article-lookup/doi/10.1093/nc/niw011> (17 March 2017, last accessed).
- Peterson A, Cruse D, Naci L et al. Risk, diagnostic error, and the clinical science of consciousness. *NeuroImage Clin* 2015;**7**:588–97.
- Posner JB, Saper CB, Schiff ND et al. *Plum and Posner's Diagnosis of Stupor and Coma*, 4th edn. New York: Oxford University Press, 2007.
- Raichle ME, MacLeod AM, Snyder AZ et al. A default mode of brain function. *Proc Natl Acad Sci U S A* 2001;**98**:676–82.
- Rohaut B, Faugeras F, Chausson N et al. Probing ERP correlates of verbal semantic processing in patients with impaired consciousness. *Neuropsychologia* 2015;**66**:279–92.
- Sanders RD, Raz A, Banks MI et al. Is consciousness fragile? *Br J Anaesth* 2016;**116**:1–3.
- Schnakers C, Vanhaudenhuyse A, Giacino J et al. Diagnostic accuracy of the vegetative and minimally conscious state: Clinical consensus versus standardized neurobehavioral assessment. *BMC Neurol* 2009;**5**:1–5.
- Seel R, Sherer M, Whyte J et al. A practice parameter of the American Congress of Assessment Scales for disorders of consciousness: evidence-based recommendations for clinical practice and research. *Arch Phys Med Rehabil* 2010;**91**:1795–813.
- Sergent C, Faugeras F, Rohaut B et al. Multidimensional cognitive evaluation of patients with disorders of consciousness using EEG: a proof of concept study. *NeuroImage Clin* 2016.
- Seth AK, Dienes Z, Cleeremans A et al. Measuring consciousness: relating behavioural and neurophysiological approaches. *Trends Cogn Sci* 2008;**12**:314–21.
- Shannon CE. A mathematical theory of communication. *Bell Syst Tech J* 1948;**27**:379–423.
- Shea N, Bayne T. The vegetative state and the science of consciousness. *Br J Philos Sci* 2010;**61**:459–84.
- Siclari F, LaRocque JJ, Postle BR et al. Assessing sleep consciousness within subjects using a serial awakening paradigm. *Front Psychol* 2013;**4**.
- Sitt JD, King JR, Karoui IE et al. Large scale screening of neural signatures of consciousness in patients in a vegetative or minimally conscious state. *Brain* 2014;**137**:2258–70.
- Smith SM, Fox PT, Miller KL et al. Correspondence of the brain's functional architecture during activation and rest. *Proc Natl Acad Sci U S A* 2009;**106**:13040–45.
- Soddu A, Vanhaudenhuyse A, Bahri MA et al. Identifying the default-mode component in spatial IC analyses of patients with disorders of consciousness. *Hum Brain Mapp* 2012;**33**:778–96.
- Stender J, Gosseries O, Bruno MA et al. Diagnostic precision of PET imaging and functional MRI in disorders of consciousness: a clinical validation study. *Lancet Neurol* 2014;**6736**:8–16.
- The Multi-society Task Force on PVS. Medical aspects of the persistent vegetative state — 1. *N Engl J Med* 1994;**330**:1499–508.
- Tononi G. An information integration theory of consciousness. *BMC Neurosci* 2004;**5**:42.
- Tononi G, Boly M, Massimini M et al. Integrated information theory: from consciousness to its physical substrate. *Nat Rev Neurosci* 2016;**17**:450–61.
- Vanhaudenhuyse A, Noirhomme Q, Tshibanda LJ-F et al. Default network connectivity reflects the level of consciousness in non-communicative brain-damaged patients. *Brain* 2010;**133**:161–71.
- Weijer C, Peterson A, Webster F et al. Ethics of neuroimaging after serious brain injury. *BMC Med Ethics* 2014;**15**:41.