NeuroImage xxx (2010) xxx–xxx

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/ynimg

"Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients

Christophe L. Phillips [a,b,*], Marie-Aurelie Bruno [a,c], Pierre Maquet [a,d], Mélanie Boly [a,c], Quentin Noirhomme [a,c], Caroline Schnakers [a,c], Audrey Vanhaudenhuyse [a,c], Maxime Bonjean [a], Roland Hustinx [a,e], Gustave Moonen [a,d], André Luxen [a], Steven Laureys [a,c,d]

[a] Cyclotron Research Centre, University of Liège, B30, 4000 Liège, Belgium
[b] Department of Electrical Engineering and Computer Science, University of Liège, B28, 4000 Liège, Belgium
[c] Coma Science Group, Department of Neurology, University Hospital, University of Liège, B35, 4000 Liège, Belgium
[d] Department of Neurology, University Hospital, University of Liège, B35, 4000 Liège, Belgium
[e] Department of Nuclear medecine, University Hospital, University of Liège, B35, 4000 Liège, Belgium

ARTICLE INFO

Article history:
Received 15 December 2009
Revised 4 May 2010
Accepted 28 May 2010
Available online xxxx

Keywords:
FDG-PET
Vegetative state
Locked-in syndrome
Consciousness
Classifier
Relevance vector machine

ABSTRACT

The vegetative state is a devastating condition where patients awaken from their coma (i.e., open their eyes) but fail to show any behavioural sign of conscious awareness. Locked-in syndrome patients also awaken from their coma and are unable to show any motor response to command (except for small eye movements or blinks) but recover full conscious awareness of self and environment. Bedside evaluation of residual cognitive function in coma survivors often is difficult because motor responses may be very limited or inconsistent. We here aimed to disentangle vegetative from "locked-in" patients by an automatic procedure based on machine learning using fluorodeoxyglucose PET data obtained in 37 healthy controls and in 13 patients in a vegetative state. Next, the trained machine was tested on brain scans obtained in 8 patients with locked-in syndrome. We used a sparse probabilistic Bayesian learning framework called "relevance vector machine" (RVM) to classify the scans. The trained RVM classifier, applied on an input scan, returns a probability value (p-value) of being in one class or the other, here being "conscious" or not. Training on the control and vegetative state groups was assessed with a leave-one-out cross-validation procedure, leading to 100% classification accuracy. When applied on the locked-in patients, all scans were classified as "conscious" with a mean p-value of .95 (min .85). In conclusion, even with this relatively limited data set, we could train a classifier distinguishing between normal consciousness (i.e., wakeful conscious awareness) and the vegetative state (i.e., wakeful unawareness). Cross-validation also indicated that the clinical classification and the one predicted by the automatic RVM classifier were in accordance. Moreover, when applied on a third group of "locked-in" consciously aware patients, they all had a strong probability of being similar to the normal controls, as expected. Therefore, RVM classification of cerebral metabolic images obtained in coma survivors could become a useful tool for the automated PET-based diagnosis of altered states of consciousness.

Introduction

Following traumatic or non-traumatic brain damages, some patients will fall into an irreversible coma, and possibly brain death. Still some others will "awaken" (i.e., recover sleep–wake cycles) from their coma but will remain fully unaware (i.e., will only show reflex movements without command following) in a "vegetative state" (VS). "Locked-in syndrome" (LIS) patients also awaken from their coma but recover full consciousness while being selectively de-efferented (i.e., they have no means of producing speech, limb or face movements in response to commands except for small eye movements or blinks) (Laureys, 2005; Laureys et al., 2005). Patients with LIS sometimes may remain comatose for days or weeks, needing artificial respiration and then gradually wake up, remaining paralyzed and voiceless, superficially resembling patients in VS. Distressingly, studies reported that the diagnosis of LIS on average takes over 2.5 months and in some cases took 4 to 6 years before aware and sensitive patients, locked in an immobile body, were recognized as being conscious (León-Carrión et al., 2002). Recognizing unambiguous signs of conscious perception of the environment and of the self in coma survivors can be very challenging. This difficulty is reflected in the frequent misdiagnoses (approximately 40%) of the vegetative state (Laureys et al., 2004; Schnakers et al., 2009b). An accurate and reliable evaluation of the level and content of consciousness in severely brain-damaged patients is thus of paramount importance for their appropriate management.

* Corresponding author. Cyclotron Research Centre, University of Liège, Sart Tilman, Bldg. B30, B-4000 Liège, Belgium. Fax: +32 4 366 2946.
E-mail address: c.phillips@ulg.ac.be (C.L. Phillips).

1053-8119/$ – see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2010.05.083

Table 1
Demographic, clinical, electroencephalographic (EEG), structural neuroimaging of patients in a vegetative state (VS) and locked-in syndrome (LIS).

Patient number	Gender /age (year)	Etiology	Interval (m)	Arousal[a]	Auditory function[a]	Visual function[a]	Motor function[a]	Oro-motor function[a]	Communication[a]	EEG	MRI/CT
VS_01	M/53	Anoxia	10	Eyes opening without stimulation	Auditory startle	Visual startle	Spastic quadriplegia, abnormal posturing	Oral reflexive movement	None	Generalized suppression	No focal lesions[b]
VS_02	F/65	Anoxia	9	Eyes opening with stimulation	Auditory startle	Visual startle	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	None	Low voltage delta activity more prominent on the right	Corticosubcortical atrophy and secondary hydrocephalus without focal lesion
VS_03	M/48	Anoxia	30	Eyes opening with stimulation	Auditory startle	Visual startle	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	None	Low-voltage delta, discontinuous generalized paroxistic activity	Diffuse corticosubcortical atrophy and secondary hydrocephalus without focal lesion
VS_04	M/69	Anoxia	1	Eyes opening with stimulation	None	None	Spastic quadriplegia, abnormal posturing	Oral reflexive movement	None	Generalized theta-delta activity	Lesions in bilateral paraventricular, insular, prefrontal and left parietal areas
VS_05	M/53	CVA	2	Eyes opening without stimulation	Auditory startle	Visual startle	Spastic quadriplegia, flexion withdrawal	None	None	Generalized theta-delta activity	Quadri-ventricular hemorrhage, multifocal cortical and brainstem lesions
VS_06	F/42	Anoxia	3	Eyes opening without stimulation	Auditory startle	None	Spastic quadriplegia, abnormal posturing	Vocalization/ oral movement	None	NA	NA
VS_07	F/36	Anoxia	30	Eyes opening without stimulation	Auditory startle	Visual startle	Spastic quadriplegia, flexion withdrawal	None	None	Generalized theta activity, sporadic delta dysrythmia	Corticosubcortical atrophy, bilateral subcortical and prefrontal lesions
VS_08	F/69	Trauma	1	Eyes opening with stimulation	Auditory startle	None	Spastic quadriplegia, no response to noxious stimulation	Oral reflexive movement	None	Generalized theta-delta activity, left frontal paroxistic activity	Multifocal meningeal hemorrhages, bilateral temporal contusions, left frontal subdural hematoma
VS_09	F/79	Sub-arachnoid hemorrhage	1	None	None	None	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	None	Generalized theta activity, left central paroxistic activity	Multifocal meningeal and intraventricular hemorrhages[b]
VS_10	M/62	Anoxia	9	Eyes opening without stimulation	Auditory startle	None	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	None	Theta-delta activity (parietal and occipital)	Corticosubcortical atrophy, periventricular, basal ganglia, cerebellar and brainstem lesions
VS_11	F/70	Encephalitis	1	Eyes opening with stimulation	Localization to sound	Visual startle	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	None	Generalized non-reactive delta activity	Bilateral prefrontal, temporal and periventrical lesions

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

VS_12	F/35	Anoxia	17	Eyes opening without stimulation	Auditory startle	Visual startle	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	None	Disorganized low-voltage theta-delta basic rhythm	Focal lesions in basal ganglia
VS_13	M/55	Anoxia	6	Eyes opening with stimulation	Auditory startle	None	Spastic quadriplegia, abnormal posturing	Oral reflexive movement	None	Generalized delta activity	NA
LIS_01	M/44	Anoxia	285	Eyes opening with stimulation	Reproducible movement to command	Visual pursuit	Spastic quadriplegia, flexion withdrawal	Vocalization/ oral movement	None	Alternating alpha-theta and sporadic delta activity	Bilateral frontal corticosubcortical, brainstem and cerebellar atrophy
LIS_02	F/21	Basilar artery thrombosis	0.5	None	None	None	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	None	Unstructured theta basic rhythm	Multifocal ischemic lesions in brainstem, bilateral cerebellum and thalamus
LIS_03	M/44	Basilar artery thrombosis	51	Attention	Consistent movement to command	Object recognition	Spastic quadriplegia, functional object use	None	Non-functional : intentional	Reactive alpha basic rhythm	Brainstem, bilateral cerebellum, thalamus, right occipital and mesiotemporal lesions
LIS_04	F/37	Basilar artery thrombosis	1	Eyes opening without stimulation	Reproducible movement to command	Object recognition	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	Non-functional : intentional	Reactive 6 Hz basic rhythm	Multifocal ischemic lesions in brainstem, right temporal and prefrontal areas[b]
LIS_05	M/53	Basilar artery thrombosis	3	None	Reproducible movement to command	Visual pursuit	Spastic quadriplegia, abnormal posturing	Vocalization/ oral movement	Non-functional : intentional	6 Hz basic rhythm	Ischemic lesion in brainstem, midbrain and left cerebellum
LIS_06	M/35	Trauma	285	Attention	Auditory startle	Visual startle	Spastic quadriplegia, flexion withdrawal	Vocalization/ oral movement	None	Alpha activity with posterior theta dysrythmia	Diffuse corticosubcortical atrophy and secondary hydrocephalus[b]
LIS_07	F/46	Basilar artery thrombosis	1	Attention	Consistent movement to command	Object recognition	Spastic quadriplegia, flexion withdrawal	Oral reflexive movement	Functional: accurate	Reactive alpha activity, intermittent delta-theta	Ischemic lesion in brainstem extending to midbrain
LIS_08	F/42	Brainstem hemorrhage	2	Eyes opening with stimulation	Reproducible movement to command	Visual pursuit	Spastic quadriplegia, flexion withdrawal	None	Non-functional : intentional	Diffuse theta activity	Hemorrhagic lesion in brainstem

Demographic data: patient number, gender and age (at scanning time), etiology, interval (in month) between the acute brain insult and scanning, Coma Recovery Scale—Revised subscores observations (arousal, auditory function, visual function, motor function, oro-motor function, communication), EEG characteristics and structural MRI/CT diagnostics. NA = not available.

[a] Based on Coma Recovery Scale—Revised assessment (CRS-R).

[b] X-ray CT data.

For the time being, consciousness cannot be measured objectively by any equipment and its bedside estimation requires the interpretation of multiple clinical signs. Behavioural assessment remains the "gold standard" for detecting signs of consciousness and, hence, for determining diagnosis. Clinically, we are limited to the appraisal of the patient's capacity to perceive the external world and to voluntarily interact with it (i.e., perceptual awareness). Bedside evaluation of residual brain function in severely brain-damaged patients is complicated by the presence of motor impairment, tracheotomy, fluctuating arousal level or ambiguous and rapidly habituating responses (Majerus et al., 2005).

We here aim to disentangle VS from LIS patients by teaching a "machine" to discriminate consciousness, using fluorodeoxyglucose PET (FDG-PET) scans obtained in healthy controls and patients in VS. The capacity of the machine at discriminating new (or unseen) images was directly assessed via a cross-validation procedure on data from the healthy controls and VS patients. Next, the trained machine was further tested on brain scans obtained in patients with LIS.

Materials and methods

PET data

The data consisted of FDG-PET scans from 37 healthy control subjects, further referred to as the "CO" group, and 45 patients. The 45 patients are characterized by different pathologies and split into three subgroups: 13 in a vegetative state ("VS"), 8 with locked-in syndrome ("LIS"), plus 24 patients with ambiguous consciousness signs.

FDG-PET reflects the uptake of glucose by the brain, i.e., its energy consumption. Areas of hypometabolism and relatively preserved metabolism can be associated to a decrease (preservation) of local brain activity (Laureys et al., 2008). Cerebral metabolism data were acquired after intravenous injection of five to ten mCi (185–370 MBq) FDG (as described in Laureys et al., 2000) on a Gemini PET scan (Philips Medical Systems) at Liège university hospital.

Subjects

The 37 control subjects were healthy volunteers (17 men; mean age 45 ± 17 years; range 18–80 years). All patients had standardized behavioural assessment using the Coma Recovery Scale—Revised (CRS-R) (Giacino et al., 2004; Kalmar and Giacino, 2005) performed the day of PET scanning by an experienced neuropsychologist (MAB, CS and AV). The CRS-R is a standardized and validated (Schnakers et al., 2008) neurobehavioural assessment scale to determine patients' level of consciousness. CRS-R has shown superior performance in detecting consciousness in severely brain-damaged patients when compared to other scales (Giacino et al., 2002) or unstructured neurological assessment (Schnakers et al., 2009b). It assesses auditory, visual, verbal and motor functions as well as communication and arousal level. The total score ranges between 0 (worst) and 23 (best).

All VS patients fulfilled the international criteria for VS: (1) spontaneous eye opening with intermittent wakefulness and behaviourally assessed sleep–wake cycles; (2) no evidence of language expression or comprehension; and (3) no evidence of reproducible voluntary behavioural responses to any stimuli, i.e., no evidence of awareness of the environment (ANACoEA, 1993; MSTFoPVS, 1994). All LIS patients fulfilled the international criteria for LIS: (1) spontaneous eye opening with intermittent wakefulness and behaviourally assessed sleep–wake cycles (bilateral ptosis was ruled out as a complicating factor); (2) no evidence of verbal language expression (e.g., aphonia or severe hypophonia); (3) quadriplegia or quadriparesis; and (4) primary mode of communication that uses vertical or lateral eye movement or blinking of the upper eyelid (ACoRM, 1995).

The 45 patients (26 men; mean age 46 ± 16 years; range 14–79 years) were scanned between 14 days and 24 years following the acute brain damage: 16 patients within 3 months and 29 in the chronic stage, i.e., more than 3 months after brain damage. The 24 scans from patients with ambiguous consciousness signs were only used at the spatial pre-processing stage in order to balance the number of patients and healthy CO subjects in the creation of the study-specific PET template (see next section). Table 1 summarizes the demographic and clinical characteristics (Giacino et al., 2004) of the 13 VS and 8 LIS patients (10 men; mean age 50 ± 15 years; range 21–79 years), which are the focus of this study. Informed consent was obtained for all control subjects and for LIS patients, and from the legal representative of VS patients and those with ambiguous consciousness signs. The study was approved by the Ethics Committee of the University and University Hospital of Liège.

Visual "gold standard"

To compare our classification procedure (described in the next 2 sections) with a "gold standard", four experts (1 neuropsychologist, MAB, and 3 neurologists, MB, RH and SL, used to examine patient PET scans) were required to visually classify anonymized raw PET images. The images from the 13 VS and 8 LIS patients and 13 healthy CO subjects (picked at random among the 37 images available) were randomly presented to each examiner separately. Using only his expertise and visual inspection, he had then to decide if the images came from a conscious or unconscious subject. Since the differentiation of LIS patients from VS patients is the most challenging and relevant clinical question, we considered the classification error rate of only the VS and LIS patients. Then the experts' specificity and sensitivity of "consciousness detection" were estimated. Finally the mean and standard deviation of the error rate, specificity and sensitivity were calculated.

Spatial pre-processing and data extraction

Before applying any analysis tool, either "classic" statistical mapping or classification method, images from the different subjects and patients must be brought into a common (standard) space. This should ensure that homologous brain regions from the different subjects/patients are aligned, and that their activity can be compared. The data were spatially pre-processed using the "Statistical Parametric Mapping" Matlab toolbox SPM8 (Wellcome Trust Centre for Neuroimaging, University College London, UK).

The first task is to bring all the images, from patients and controls, into a standard stereotactic space, here the Montreal Neurological Institute (MNI) space (Evans et al., 1993). This image normalization is complicated by 2 factors: first, the PET template provided with SPM was built using $H_2^{15}O$ cerebral blood flow PET (i.e., not FDG cerebral metabolism) images from healthy subjects, and second, we are here dealing with a mix of images from normal healthy controls and patients, some of which with pathological brain deformations. Indeed some of the patient images present abnormal brain structures, mainly overinflated ventricles (i.e., ex-vacuo hydrocephalus secondary to cortical and subcortical lesions and atrophy). The normalization therefore proceeded in three successive steps:

- normalize all 82 images, 45 patients and 37 controls, individually using the SPM-PET template along with the a priori "gray matter" (GM) and "white matter" (WM) density images, also available with SPM. The combination of those 3 images (PET, GM and WM templates) into a weighted template should better account for differences in brain tissue intensity and distribution.
- average together the 82 normalized images and smooth with an 8-mm FWHM kernel. This smoothed averaged image can then be used as the study-specific template, see Fig. 1 (bottom left).
- redo the normalization of the 37 controls, 13 VS and 8 LIS images but using the study-specific template instead of the SPM-PET template. The images are finally resampled into a $40 \times 48 \times 34$ voxels image, cubic $4 \times 4 \times 4$ mm^3 voxels.

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

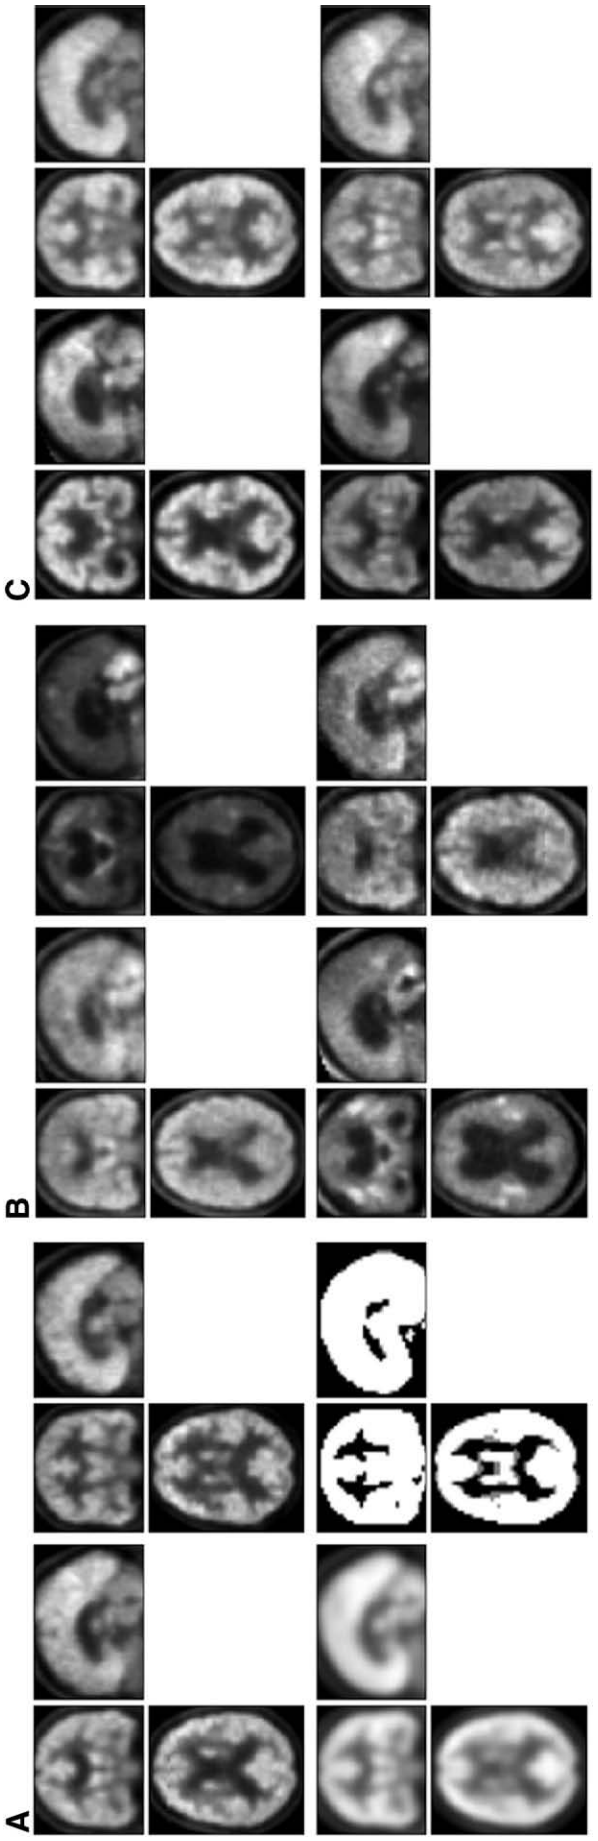

Fig. 1. Examples of normalized PET images of (B) VS patient number VS_04, VS_07, VS_10, VS_13, (C) LIS patient number LIS_01, LIS_02, LIS_03, LIS_04 and (A top row) 2 healthy controls. These images show the raw PET signal after spatial normalization normalization to the MNI space before being rescaled by the image global mean. The bottom row of panel A shows the study-specific template obtained after the first normalization step (left) and the gray matter mask used to select voxels (right).

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

At each of the normalization steps, the regularisation imposed on the nonlinear warping of the normalization procedure was also increased by one order, compared to the standard setting, to prevent unrealistic warping, as suggested in (Crinion et al., 2007).

After the second normalization step, all images are thus aligned within the MNI standard space, see Figs. 1A (top row), B and C. Since the a priori gray matter density map is also in the MNI space, in each image we sampled the 24603 voxels (out of 64280 voxels from the whole image), which had a probability higher than 50% to be gray matter, see Fig. 1 (bottom right).

Finally, to account for differences of global signal between the images, each set of 24603 retained voxels was scaled by its global mean signal. The brain metabolism of each healthy control or patient was thus "summarized" by a single 24603×1 data vector extracted from their FDG-PET scan after normalization, resampling and scaling.

Relevance vector machine classification

Our objective was first to train a "pattern classification" machine (Bishop, 2006a) to discriminate two groups of subjects, then to apply the trained classifier on the data from a third group. We employed a linear "relevance vector machine" (RVM) classifier (Tipping, 2001; Bishop, 2006b) to discriminate the CO and VS groups, using only the data vector of each subject and the associated "CO" or "VS" label. Then the trained RVM classifier was applied on the data from the LIS group, the output indicating if each LIS patient brain metabolism is "more like" the conscious CO or the unconscious VS.

RVM is a kernel method, linear in the parameters and similar to the "support vector machine" (SVM) (Müller et al., 2001). However, it offers several advantages over SVM, mainly probabilistic predictions and automatic estimation of hyper-parameters. See the Appendix section for a detailed mathematical description of the RVM approach. In few words, each "data element" consisted in a 24603×1 vector with the sampled metabolic values of each subject's PET image and a label, 0 or 1 for the VS and CO, respectively. The data from the 37 CO and 13 VS, i.e., 50 data vectors and their label, are used to train the RVM classifier. After training, the RVM classifier is applied on the LIS data vector and returns a posterior probability value (p-value) of being in one class or the other: here being more like the CO ($p > .5$) or more like the VS ($p < .5$). An SVM-like "hard decision" could be easily obtained by rounding the p-value to the closest integer value, i.e., 0 for VS or 1 for CO, but throws away any probabilistic information provided by the p-values.

RVM is by definition a multivariate approach and all the features (here metabolic voxel values) are used simultaneously to obtain a prediction for an input data vector. Still not all the features from all the

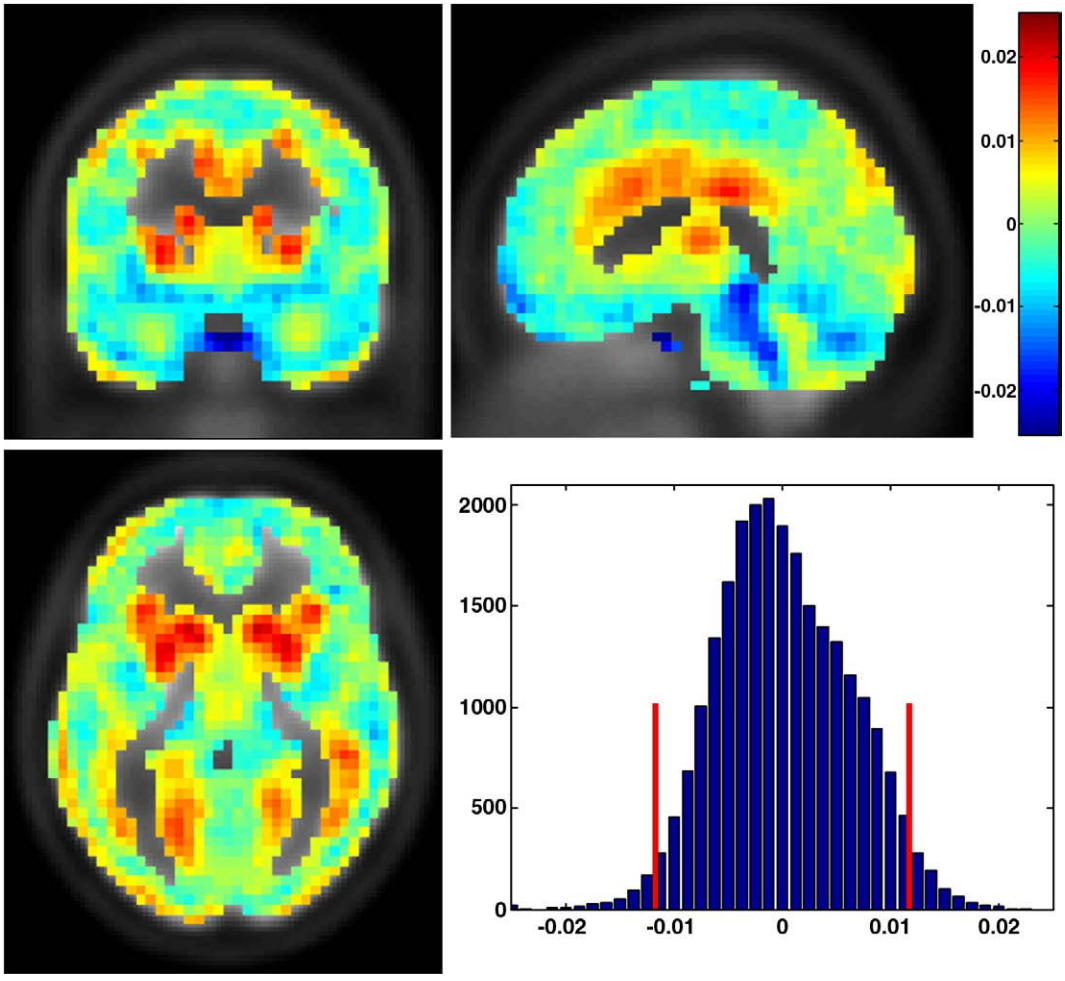

Fig. 2. Untresholded relevance map, displayed on top of the study-specific FDG-PET template, and histogram of the voxel relevance. The relevance map shows the relevance of each of the voxels as a tri-dimensional image (top left, coronal view; top right, sagittal view; bottom left, transverse view) and is built when training the RVM classifier on the whole VS-versus-CO data set. A positive value (voxel displayed in yellow-red) indicates that a relatively larger metabolic value at those voxels drives the classification of the scan towards the CO class. Conversely, negative values (voxels displayed in blue-purple) "pushes" to a classification in the VS class if the metabolic values are relatively large in those areas. Voxel with low relevance, displayed in green, simply do not contribute much to the classification of data. The histogram of the voxel relevance summarizes the distribution of the relevance values, indicating that a majority of voxels have relevance close to zero. The two red lines on the histogram show the percentile 95 threshold (± 0.0117) applied to select a subset of voxels.

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

images are equally important for the posterior *p*-value estimation. With a linear RVM classifier as we used, the relevance of each voxel (see the last section of the Appendix) can be produced as a weighted linear combination of the data elements used for training: "voxel relevance" simply indicates how each individual voxel in an image contributes to the classification of this image into one class or the other. When a new image, i.e., an image not used for training the RVM classifier, has to be classified, the values at its voxels are weighted by the corresponding voxel relevance and summed together. This value is either positive or negative, i.e., in one category or the other, providing a simple binary classification. Finally the weighted sum of voxel values is turned into a posterior probability value, i.e., a value between 0 and 1, through a sigmoid function, which gives insight on how certain the classification is.

Since there is one relevance value per voxel, this can be displayed as an image or "relevance map", see Fig. 2. Note that relevance is *not* a statistical value per se but is just a weighting factor attributed to each voxel. A whole range of voxel relevance is obtained, which can be arbitrarily thresholded to keep only the most relevant voxels (or any other a priori criteria could be used). We here fix a percentile 95 threshold and keep only the 1230 voxels, the "relevant" set, which have the 5% highest (in absolute value) relevance, see Figs. 2 and 3. The RVM classifier is then trained again only on those 1230 most relevant voxels, providing a different classifier based only on a small subset of the original features. The goal is to assess if those 1230 voxels contain on their own enough information to classify CO and VS data. This second classifier is applied on the LIS data vector limited to the same 1230 voxels, providing a second classification estimate for these data.

Crucially to validate the RVM training and ensure its robustness to new or unseen data vectors, the RVM classifier, with the full (24603 voxels) or "relevant" (1230 voxels) data set, is cross-validated using a "leave-one-out" (LOO) technique: Practically the classifier is trained on the data set minus one "left-out" data element, i.e., 49 data vectors and their respective labels, and applied on the left-out data vector to estimate its posterior *p*-value. These individual *p*-values are looked at to check the certainty of the classification. Then the estimated label, i.e., rounded *p*-value, is compared with the known label and any discrepancy is counted as an error, leading to the classifier error rate. The training and testing for the LOO validation is thus performed 50 times, once for each CO and VS data element.

Results

When the full data set is used to train the classifier, cross-validation leads to 0% error rate: the label of each data vector, CO or VS, was correctly retrieved by a classifier trained on the other 49 data vectors and labels. The CO subjects obtained a mean *p*-value of .99 (minimum .89) and the VS patients, .06 (maximum .38), see Table 2.A (third column). When the trained RVM classifier was applied on the LIS data vectors, all 8 LIS patients had a good probability (mean *p*-value of .91 and minimum .61) of being similar to the CO, conscious healthy subjects, as was expected (see Table 2.B, third column).

The relevance map, see Fig. 2, shows the distribution of voxel relevance throughout the brain volume. Most voxels, as seen on the relevance histogram, have a relevance close to zero (greenish color on Fig. 2) and therefore do not contribute much to the *p*-value estimate and classification result. The few ones with large (in absolute value) relevance are almost equally distributed between those pushing the classifier towards the CO or VS class, in the two tails of this distribution.

By thresholding the relevance map at percentile 95 and selecting only the voxels with a relevance (in absolute value) larger than .0117, one ends up with the 1230 voxels (out of 24603) that have the largest influence on the RVM. This thresholded map includes 73 clusters of voxels, counting from 1 to 251 voxels, and provides information about which area ensembles in the brain are important to discriminate VS and CO, see Fig. 3 and Table 3. This thresholding effectively works as a straightforward feature selection procedure (Guyon and Elisseeff, 2003).

The classifier using only the "relevant" subset of 1230 voxels (95% percentile) had an overall 2% cross-validation error rate, with mean and maximum *p*-values for the VS of .05 and .54 (see Table 2.A, fourth column), and for the CO a mean and minimum *p*-value of .99 and .93.

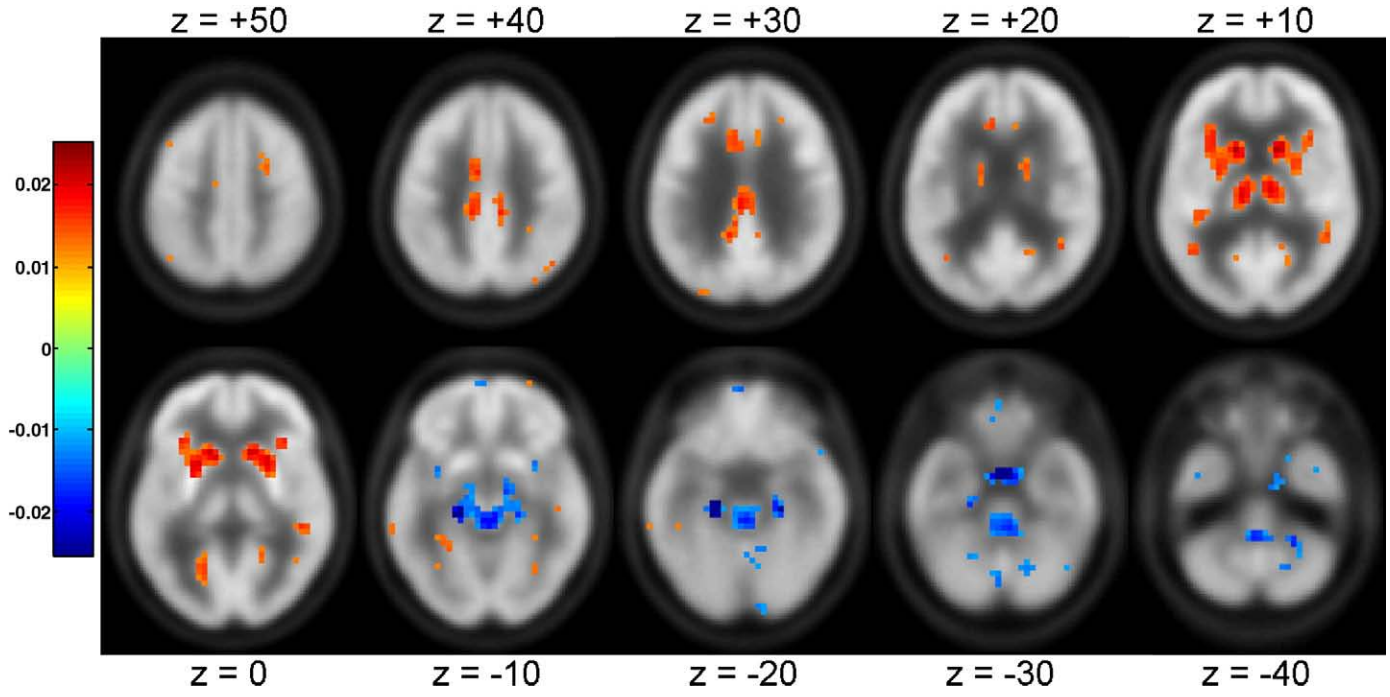

Fig. 3. Relevance map thresholded at percentile 95: only the 5% most relevant (in absolute value) voxels with relevance larger than .0117 are displayed, over the study-specific FDG-PET template. The same color scale as in Fig. 2 is used, and the interpretation of the voxel relevance is similar. Transverse slices are in MNI space at $z = 50, 40, 30, 20, 10, 0, -10, -20, -30, -40$ mm.

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

Table 2
Results of the RVM classifier: (A) patients in a vegetative state (VS) and (B) locked-in syndrome (LIS).

Patient number	Gender/age (year)	*p*-value (full set)	*p*-value (relevant set)	*p*-value ("no-edge" set)
(A)				
VS_01	M/53	0	0	0
VS_02	F/65	.0003	.0001	.0001
VS_03	M/48	.0014	.0052	.0021
VS_04	M/69	.0099	.0102	.0098
VS_05	M/53	.3768	.5383	.7384
VS_06	F/42	.0003	.0036	.0020
VS_07	F/36	0	0	0
VS_08	F/69	.1688	.0042	.0022
VS_09	F/79	.0111	.0085	.0144
VS_10	M/62	0	0	0
VS_11	F/70	0	0	0
VS_12	F/35	.2408	.0268	.0109
VS_13	M/55	.0108	.0058	.0016
(B)				
LIS_01	M/44	.6094	.4835	.2821
LIS_02	F/21	.9982	.9992	.9996
LIS_03	M/44	.9997	.9907	.9822
LIS_04	F/37	.9956	.9893	.9936
LIS_05	M/53	.9887	.8236	.8485
LIS_06	M/35	.7094	.9869	.9866
LIS_07	F/46	.9994	.9986	.9980
LIS_08	F/42	.9900	.9900	.9889

RVM classifier results: *p*-values indicate the probability of being conscious, using the full set, the "relevant" set (95 percentile most relevant voxels) or "no-edge" set (same as the "relevant" set without the small clusters of voxels on the edge of the brain) of voxels.
Note that the *p*-values for the VS patients are those of the leave-one-out cross-validation, while those for the LIS patients are obtained after training of the full CO and VS data sets and thus predict their consciousness state.

The only cross-validation error comes from the VS patient which already had the highest *p*-value (.38) when employing all the voxels. When the "relevant" set classifier was applied on the LIS data, 7 out of 8 LIS patients were classified as CO. For the whole LIS group, the mean and minimum *p*-value are .91 and .48 (see Table 2.B, fourth column).

Table 3
Description of the main clusters of most relevant voxels (percentile 95, threshold at ±.0117) in the brain volume.

Direction (CO/VS)	Cluster size (no. voxels)	Relevance in cluster, mean ± SD	Anatomical structure	Slices
VS	251	−0.01666 ± 0.00775	Brainstem & bilateral lower temporal lobe	$z = -10, -20, -30, -40$
VS	56	−0.01367 ± 0.0017	Right cerebellum	$z = -20, -30, -40$
VS	16	−0.01315 ± 0.00114	Left cerebellum	$z = -30, -40$
CO	211	0.01484 ± 0.0223	Left caudate & thalamus	$z = 0, 10, 20$
CO	141	0.01493 ± 0.0025	Right caudate	$z = 0, 10, 20$
CO	45	0.01451 ± 0.00219	Right thalamus	$z = 0, 10, 20$
CO	13	0.01242 ± 0.00073	Posterior cingulate cortex & precuneus	$z = 30, 40$
CO	69	0.01334 ± 0.00097	Frontal medial cortex	$z = 30, 40$
CO	5	0.01246 ± 0.00079	Left dorsolateral prefrontal cortex	$z = 30$
CO	8	0.01233 ± 0.00057	Right dorsolateral prefrontal cortex	$z = 50$

Direction of relevance (towards the unconscious vegetative, VS, or conscious control, CO, group), size of cluster (in voxels), mean ± standard deviation of relevance over the cluster, anatomical description of the cluster volume and transverse slice position (in MNI space).

Table 4
Results of the "gold standard" visual classification of consciousness by four experts, compared to those of the RVM classifier.

	Error rate	Sensitivity	Specificity
Expert 1	33%	50%	77%
Expert 2	19%	75%	85%
Expert 3	33%	75%	62%
Expert 4	19%	88%	77%
Mean ± SD of experts	26 ± 7%	72 ± 16%	75 ± 10%
RVM, full set	0%	100%	100%
RVM, relevant/no-edge set	10%	88%	92%

The error rate, sensitivity and specificity take into account only the 13 VS and 8 LIS patients, as this classification is the most relevant from a clinical point of view.

The misclassified LIS patient had a *p*-value only marginally in favour of the VS group (.48) and was the same atypical LIS patient who also had a low *p*-value when using all the voxels.

As can be seen on Fig. 3, among the 1230 voxels of the "relevant" set, a few of them (usually isolated) are located on the edge of the brain or ventricles. To avoid any spurious "edge effect", clusters of relevant voxels counting less than 5 voxels or clearly lying on an edge were manually removed from the "relevant" set. This new "no-edge" set of voxels counts 1056 voxels, distributed in 20 clusters (down from 73 with the "relevant" set) with 6 to 251 voxels each. A third classifier was thus trained and validated using this "no-edge" set of 1056 voxels. The results are very similar to those obtained with the "relevant" set with an overall 2% cross-validation error rate and the same misclassified LIS patient (see Tables 2.A and B, fifth column).

The results of the "gold standard" visual classification by the experts are summarized in Table 4. In this table the error rate, sensitivity and specificity take into account only the 13 VS and 8 LIS patients, as this classification is the most relevant from a clinical point of view. Moreover, only one expert did misclassify 2 healthy CO as unconscious. The performances of the RVM classifier, with either the full voxel set or the relevant and no-edge voxel sets, are estimated for the same 21 VS and LIS patients only.

Discussion

The very low error rate of the cross-validation, 0% and 2% for the full and "relevant" voxel sets respectively, shows the reliability and efficacy of the RVM classifier at discriminating CO and VS subjects, even though training is performed on only 49 images. By using only a subset of voxels, we were expecting a decrease in the performances of the classifier, even if the voxels selected were the most relevant ones in the complete set. Still the observed error rate remains relatively low,[1] indicating the robustness of the classifier and the relevance of the few voxels selected.

When the trained classifier is applied on the LIS data, 8 (resp. 7) out of the 8 LIS images are classified as conscious CO's when using the full (respectively "relevant") voxel set. This result is crucial for medical staff and patient carers as the LIS scans are not easily distinguishable from the VS ones by visual examination. Currently LIS patients are distinguishable from VS patients only through the interpretation of multiple clinical signs from bedside estimation and examination of various data recordings (PET, MRI/CT, EEG). The subject LIS_01, with the lowest probability, $p = .61$ (resp. $p = .48$) of being similar to the CO using the full (resp. "relevant") set of voxel, is in fact an atypical post-traumatic LIS patient with both brainstem and cortical lesions. This intermediate *p*-value is a sign that the PET image of that patient does not allow a reliable classification in the CO or VS class. The

[1] Note that, if the threshold for voxel selection is fixed at the 90% percentile, i.e., 2430 voxels are retained, the error rate for the cross-validation and LIS patient classification is zero.

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

patient could simply have fallen asleep during the examination or temporarily been at a lower level of consciousness, which cannot be monitored. The *p*-values of the other LIS patients are much higher, usually above .95, and would provide a strong argument in favour of the LIS diagnosis. For example, the subject LIS_02 was clinically comatose when studied by PET 14 days post-brainstem stroke. Only ERP's were able to indicate presence of command following and hence of consciousness on day 49 (clinical details and cognitive ERP have been published in Schnakers et al., 2009a). This case is interesting as we believe the PET scan, showing no hypometabolism in any supra-tentorial area, is incompatible with the diagnosis of coma or VS. Indeed, we previously demonstrated that the latter conditions are characterized by a widespread frontoparietal cortical hypometabolism (Laureys et al., 1999b; Laureys et al., 2004). In our view, this patient classified as probably conscious by the RVM, indeed most likely was in pseudo-coma or total LIS (Laureys et al., 2005) during the PET study.

By comparison, the "gold standard" visual classification of PET scans by an expert is outperformed by the automatic RVM approach, as shown in Table 4. In all cases, the expert had a higher error rate of discriminating LIS from VS patients and their sensitivity and specificity was lower than that of the RVM classifier. This emphasises the need for a reproducible, systematic and user-independent interpretation of the PET data from patients with disorders of consciousness.

One key advantage of the RVM method over other similar approaches (like SVM) is clearly the estimated posterior probability associated with the classification: In a medical environment where a patient has to be diagnosed, one would consider more cautiously a classification result with a *p*-value of 51% to be in one class than with a *p*-value of 99%. Still this probability should not be regarded as a regressor estimate of the spectrum of consciousness level, since training is only performed on 2 classes of subjects/patients at both extremes, fully conscious or unconscious. Regressing out the consciousness level of patients would require a regression model and, for training, rely on data including patients at intermediate level of consciousness with their (independently) estimated consciousness level.

The validation of any "pattern classification" machine relies on the total independence between the training and the test data sets. Here this assumption is not totally met as LIS images were included in the building of the study-specific template. This could in theory bias the classification as the training data (CO/VS data) are not totally independent from the LIS images, through the study-specific template used to normalize all the images. This issue is actually the same, still in theory, for the LOO cross-validation procedure, as the "left-out" data are not entirely independent from the other "left-in" data used for the training. Nevertheless the bias, if any, should be minimal as the template is the smoothed average of 82 FDG-PET images (which by nature have a low spatial resolution): the inclusion/exclusion of one single image (for the CO-versus-VS cross-validation) or even 8 images (prediction for the LIS data) would hardly affect the study-specific template and the spatial normalization of the other images.

The relevance map, thresholded or not, effectively shows a network of brain areas whose relative level of activity taken together allows the discrimination between two classes of data, here conscious and unconscious subjects. Since the data were scaled by their mean, large negative (resp. positive) relevance at voxel *k* can be interpreted as this: a relatively larger metabolic value at voxel *k* pleads towards this data vector being from class VS (resp. CO). Still the relevance map is not like a classic statistical map and the values observed at each voxel have no statistical meaning on their own. Looking at the 5% most relevant voxels, as seen on Fig. 3 and summarized in Table 3, the highlighted areas broadly correspond to regions involved in consciousness and resting-state networks. Indeed VS patient cerebral dysfunction was not identified in one brain region but classic "group comparison" statistical analyses have shown a wide frontoparietal network encompassing the polymodal associative cortices (Laureys et al., 1999a, b, 2002): bilateral lateral frontal regions, parieto-temporal and posterior parietal areas, mesiofrontal, posterior cingulate and precuneal cortices and bilateral thalami and caudate nuclei (for a review, see Laureys et al., 2004). These regions are also known to be the most active "by default" in resting non-stimulated conditions (Raichle and Mintun, 2006) and to be important in various functions that are necessary for consciousness, such as attention, memory and language (Baars et al., 2003). Another hallmark of the vegetative state is the relative sparing of metabolism in the brainstem (encompassing the ponto-mesencephalic reticular formation, the hypothalamus and the basal forebrain; Laureys, 2004), allowing for the maintenance of vegetative functions in these patients such as sleep–wake cycles, autonomic and ventilatory control and cranial nerve reflexes.

There is a number of highlighted voxels on the edge of the brain. This is probably due to different normalization accuracy between the two groups: CO subjects had healthy brains and were certainly more easily normalized, while some VS patients, but certainly not all (see Fig. 1.B), had brain with abnormally large ventricles and reduced gray matter volume, making them much more difficult to normalize. These anatomical differences are more obvious at the "edge" of the brain or ventricles and are picked up by the classifier. Nevertheless, the LOO cross-validation proves that the classification is not entirely driven by such voxels from a single image, as not all VS images have abnormal anatomy neither the same deformations. The LIS scans are generally not easily distinguishable from those of VS patients and still RVM trained on the "relevant" set accurately and reliably estimated that they were more like the CO than VS in 7 out of 8 cases. Moreover, when those possibly spurious voxels are removed and the RVM is trained on the "no-edge" voxel set, the validation and classification results are similar to those obtained previously further indicating that RVM results are not strongly driven by those anatomical differences.

Conclusions

Bedside assessment in patients with disorders of consciousness following severe traumatic or non-traumatic brain damage continues to represent a major challenge. Despite the importance of diagnostic accuracy and advances in the past 15 years, Schnakers et al. (2009b) recently showed that the rate of misdiagnosis among those patients remains around 40%. While these figures cause concern, they at least emphasize that bedside diagnosis was possible—otherwise they would not have been identified as having been misdiagnosed. Hence, the here presented automated and objective "consciousness classifier" based on functional neuroimaging may have clinical relevance to confirm (or deny) the bedside diagnosis (Giacino et al., 2006; Laureys et al., 2006).

Consciousness in itself is not an on–off phenomenon but is part of a continuum (Baars et al., 2003), but here the groups of subjects used were sampled from both extremities of the consciousness spectrum: conscious LIS patients and controls versus unconscious VS patients. In this case, a two-class or binary classification does make sense. The extension of the RVM approach and its application to better diagnose intermediate states of consciousness, such as "minimally conscious state" (MCS) patients (Boly et al., 2008; Demertzi et al., 2009), remain to be explored. The proposed examiner-independent method for classification of the level of consciousness based on cerebral metabolic PET data sets in patients with severe brain damage is a much welcomed tool in clinical nuclear medicine and neurology. This not only has ethical consequences but is also crucial for the patient's rehabilitation and daily management (Laureys and Boly, 2007).

In this work, we demonstrated the ability of an RVM classifier machine to distinguish the brain metabolism of healthy controls and VS patients and we confirmed its utility in diagnosing LIS patients, sometimes clinically mistaken for VS (León-Carrión et al., 2002; Laureys et al., 2005). In conclusion, RVM classification of cerebral metabolic images obtained in coma survivors could become a useful

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

tool for the automated PET-based diagnosis of altered states of consciousness.

Acknowledgment

CP and SL are respectively Research Associate and Senior Research Associate at the Fonds National de la Recherche Scientifique de Belgique (FRS-FNRS). This research was supported by FRS-FNRS, Reine Elisabeth Medical Foundation, University and University Hospital of Liège, European Commission (DISCOS, DECODER, MINDBRIDGE, CATIA), James S. McDonnell Foundation, Mind Science Foundation, Concerted Research Action and Fondation Léon Frédéricq.

Appendix A. Relevance vector machine

In supervised learning, a training data set is used to teach a "machine". This data set comprises a set of N input vectors $\{x_n\}_{n=1}^N$ and associated target value $\{t_n\}_{n=1}^N$. This target value t_n can be a sample from a set of continuous values, i.e., a regressor, for a regression problem, or a binary value, i.e., class label, for a classification problem. From this training data set $\{x_n,t_n\}_{n=1}^N$, we wish to learn a mapping $f(x_n)$ predicting the target t_n, in order to predict the target t_* for any new input x_*.

A popular and flexible form for the function $f(x)$ is

$$f(x;w) = \sum_{i=1}^{M} w_i\phi_i(x) = w^t\phi(x) \tag{1}$$

and the output value is a linearly weighted, by the adjustable parameters $w=[w_1,w_2,...,w_M]^t$, sum of M basis function $\phi(x)=[\phi_1(x), \phi_2(x),...,\phi_M(x)]^t$. These basis functions are fixed and user-defined but of any form. The objective of the training is thus to estimate the "best" parameters w given a set of training data $\{x_n,t_n\}_{n=1}{}^N$ and fixed functions $\phi_i(x)$.

Relevance vector machine for regression

"Relevance vector machine" (RVM) is a Bayesian framework for learning in general models described here above. RVM actually relies on a particular form of Eq. (1), similar to that used for "support vector machine" (SVM) (Burges, 1998; Vapnik, 1998; Schölkopf et al., 1999):

$$f(x;w) = \sum_{i=1}^{N} w_iK(x,x_i) + w_0 = w^t\phi(x) \tag{2}$$

with $\phi(x)=[1, K(x,x_1), K(x,x_2),...,K(x,x_n)]^t$ and $w=[w_0, w_1,..., w_N]^t$. Note that the constant term w_0 is introduced in the vector of unknown parameters w. The kernel function $K(x,x_i)$ so defines one basis function per "data point" x in the training set. RVM regression employs model (2) with an additive noise term to link the vectorial input x_n and scalar target variable t_n

$$t_n = f(x_n;\ w) + \epsilon_n \tag{3}$$

where ϵ_n is a zero-mean white noise process with variance σ^2, i.e., $p(\epsilon_n|\sigma^2)=\mathcal{N}(\epsilon_n|0,\sigma^2)$. Considering noise precision β instead of its variance σ^2, i.e., posing $\beta=\sigma^{-2}$, and assuming the independence of the samples t_n, the likelihood of the complete training data set is

$$p(t|X,w,\beta) = (2\pi\beta^{-1})^{-N/2}\exp\left(-\frac{1}{2}\beta||t-\Phi w||^2\right) \tag{4}$$

where $t=[t_1,...,t_N]^t$, $X=\{x_n\}_{n=1}^N$, and $\Phi=[\phi(x_1), \phi(x_2),...,\phi(x_n)]^t$ is a $N\times(N+1)$ design matrix. With more parameters $(N+1)$ than training data samples (N), direct maximum-likelihood estimation of w would lead to over-fitting. In the RVM Bayesian framework, zero-mean Gaussian shrinkage priors are imposed on every w_i and, assuming the independence of the parameters, we have:

$$p(w_i|\alpha_i) = N(w_i|0,\alpha_i^{-1}) \Rightarrow p(w|\alpha) = \prod_{i=0}^{N} N(w_i|0,\alpha_i^{-1}) \tag{5}$$

with $\alpha=[\alpha_0, \alpha_1,...,\alpha_N]^t$, a $N+1$ vector of hyper-parameters representing the precision on the parameters. Finally uniform hyper-priors are assumed for all the precision hyper-parameters, α and β. An interesting property of these hyper-priors is that when the evidence of the model is maximized with respect to the hyper-parameters, a few of them go to infinity which effectively constraints the corresponding parameters to be zero. This is a type of "automatic relevance determination" (MacKay, 1994; Neal, 1996) leading to a sparse set of parameters w. Using Bayes rule and the properties of Gaussian functions, the posterior distribution of the weight can also be described by a Gaussian:

$$p(w|X,t,\alpha,\beta) = N(w|m,\ \Sigma) \tag{6}$$

where the mean m and covariance Σ are given by

$$\begin{aligned} m &= \beta\Sigma\Phi^t t \\ \Sigma &= (A+\beta\Phi^t\Phi)^{-1} \end{aligned} \tag{7}$$

with $A=\mathrm{diag}(\alpha_0,...,\alpha_N)$ a diagonal matrix of precisions.

In practice, the values of α and β are estimated by maximizing the marginal likelihood $p(t|X,\alpha,\beta)$, i.e., using a type-II maximum-likelihood method (Berger, 1985). Only the most probable values are thus calculated, an approximation to estimating and using their full distribution. With this simplification, the marginal likelihood can be obtained by integrating out the weight parameters

$$p(t|X,\alpha,\beta) = \int p(t|X,w,\beta)p(w,\alpha)\mathrm{d}w = N(t|0,\beta^{-1}I+\Phi A^{-1}\Phi^t) \tag{8}$$

Values of α and β that maximizes (the log of) (8) can then be obtained iteratively, using the following update rules:

$$\begin{aligned} \alpha_i^{\mathrm{new}} &= \frac{1-\alpha_i\Sigma_{ii}}{m_i^2} \\ (\beta^{\mathrm{new}})^{-1} &= \frac{||t-\Phi m||^2}{N-\sum_{i=1}^{N}(1-\alpha_i\Sigma_{ii})} \end{aligned} \tag{9}$$

where m_i is the ith element of the estimated posterior weight w and Σ_{ii} the ith diagonal element of the posterior covariance matrix Σ from Eq. (7).

Once the iterative procedure has converged to the "most probable" values α_{MP} and β_{MP}, the distribution of target value t_* for a new data point x_* is also Gaussian and estimated through

$$\begin{aligned} p(t_*|X,t,\alpha_{\mathrm{MP}},\beta_{\mathrm{MP}}) &= \int p(t_*|X,w,\beta_{\mathrm{MP}})p(w|X,t,\alpha_{\mathrm{MP}},\beta_{\mathrm{MP}})\mathrm{d}w \\ &= \mathcal{N}(t_*|m^t\phi(x_*),\sigma_*^2) \end{aligned} \tag{10}$$

with the variance estimated as

$$\sigma_*^2 = \beta_{\mathrm{MP}}^{-1} + \phi(x_*)^t\Sigma\phi(x_*) \tag{11}$$

where Σ is given by Eq. (7) with α and β set at their optimal value.

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

Relevance vector machine for classification

RVM classification follows the same framework as described in the previous section but with a modified likelihood function. Two-class problems call for a binary target variable $t_n \in \{0,1\}$ and we want to predict the posterior probability of belonging to one of the two classes, given the input data x_n. The linear Eq. (2) is therefore generalized by applying a logistic sigmoid function $\sigma(a) = \frac{1}{1 + \exp(-a)}$ such that

$$f(x; w) = \sigma(w^t \phi(x)) = \frac{1}{1 + \exp(-w^t \phi(x))} \tag{12}$$

Note that there is no noise variance here. Then, using the Bernouilli distribution, the likelihood of the training data set is defined as

$$p(t|X, w) = \prod_{n=1}^{N} \sigma(w^t \phi(x_n))^{t_n} (1 - \sigma(w^t \phi(x_n)))^{1-t_n} \tag{13}$$

Unlike the regression case, it is now impossible to integrate out the weight parameters w to directly obtain the weight posterior like in Eq. (6) or the marginal likelihood (Eq. (8)). Using the Laplace approximation, and for a fixed value of α, the mode of the posterior distribution over w is obtained by maximizing:

$$\begin{aligned} \log p(w|X, t, \alpha) &= \log\left(p(t|X, w) p(w|\alpha)\right) - \log p(t|X, \alpha) \\ &= \sum_{n=1}^{N} \left(t_n \log f(x_n; w) + (1 - t_n) \log(1 - f(x_n; w))\right) \\ &\quad - \frac{1}{2} w^t A w + \text{const} \end{aligned} \tag{14}$$

The mode and variance of the Laplace approximation for w are

$$\begin{aligned} w_{\text{MP}} &= \Sigma_{\text{MP}} \Phi^t B t \\ \Sigma_{\text{MP}} &= (\Phi^t B \Phi + A)^{-1} \end{aligned} \tag{15}$$

where B is an $N \times N$ diagonal matrix with $b_{nn} = f(x_n; w)(1 - f(x_n; w))$. Using this Laplace approximation, the marginal likelihood is expressed as

$$\begin{aligned} p(t|X, \alpha) &= \int p(t|X, w) p(w|\alpha) \mathrm{d}w \\ &= p(t|X, w_{\text{MP}}) p(w_{\text{MP}}|\alpha) (2\pi)^{M/2} |\Sigma_{\text{MP}}|^{1/2} \end{aligned} \tag{16}$$

When maximizing Eq. (16) with respect to each α_i, one eventually obtains an update rule identical to Eq. (9).

Linear kernel and relevant features

When a linear kernel is used, $K(x, x_i)$ is simply the scalar product between the 2 vectors, i.e., $K(x, x_i) = x \cdot x_i$. Then Eq. (2) can be rewritten as

$$f(x; w) = \sum_{i=1}^{N} w_i K(x; x_i) + w_0 = x \cdot \left(\sum_{i=1}^{N} w_i x_i\right) + w_0 = x \cdot \mathbf{r} + w_0 \tag{17}$$

The vector $\mathbf{r} = \Sigma_{i=1}{}^{N} w_i x_i$ is thus a weighted mean of the training data vectors and can be interpreted as the relevance of each feature: First, the larger the absolute value of an element r_k, the more important it is to discriminate the classes. Second the sign of r_k indicates the direction of this feature's influence, i.e., towards the class labelled 0 ($r_k < 0$) or 1 ($r_k > 0$).

References

ACoRM, 1995. American Congress of Rehabilitation Medicine: recommendations for use of uniform nomenclature pertinent to patients with severe alterations in consciousness. Arch Phys Med Rehabil, vol. 76, pp. 205–209.

ANACoEA, 1993. Persistent vegetative state: report of the American Neurological Association Committee on Ethical Affairs. Annals of Neurology 33, 386–390.

Baars, B.J., Ramsøy, T.Z., Laureys, S., 2003. Brain, conscious experience and the observing self. Trends Neurosci. 26, 671–675.

Berger, J.O., 1985. Statistical decision theory and Bayesian analysis.

Bishop, C.M., 2006a. Pattern Recognition and Machine Learning.

Bishop, C.M., 2006b. Pattern Recognition and Machine Learning. In: Jordan, M., et al. (Ed.), , pp. 325–358.

Boly, M., Faymonville, M.E., Schnakers, C., Peigneux, P., Lambermont, B., Phillips, C., Lancellotti, P., Luxen, A., Lamy, M., Moonen, G., Maquet, P., Laureys, S., 2008. Perception of pain in the minimally conscious state with PET activation: an observational study. Lancet Neurol. 7, 1013–1020.

Burges, C.J.C., 1998. A tutorial on support vector machines for pattern recognition. Data Min. Knowl. Discov. 2, 121–167.

Crinion, J., Ashburner, J., Leff, A., Brett, M., Price, C., Friston, K., 2007. Spatial normalization of lesioned brains: performance evaluation and impact on fMRI analyses. Neuroimage 37, 866–875.

Demertzi, A., Laureys, S., Boly, M., 2009. Coma, Persistent Vegetative States, and Diminished Consciousness. In: Banks, W. (Ed.), Encyclopedia of Consciousness, vol. 1. Elsevier Inc, pp. 147–156.

Evans, A.C., Collins, D.L., Mills, S.R., Brown, E.D., Kelly, R.L., Peters, T.M., 1993. 3D statistical neuroanatomical models from 305 MRI volumes. Proc IEEE-Nucl. Sci. Symp. Med. Imaging Conf. 1813–1817.

Giacino, J.T., Ashwal, S., Childs, N., Cranford, R., Jennett, B., Katz, D.I., Kelly, J.P., Rosenberg, J.H., Whyte, J., Zafonte, R.D., Zasler, N.D., 2002. The minimally conscious state: definition and diagnostic criteria. Neurology 58, 349–353.

Giacino, J.T., Kalmar, K., Whyte, J., 2004. The JFK Coma Recovery Scale-Revised: measurement characteristics and diagnostic utility. Arch. Phys. Med. Rehabil. 85, 2020–2029.

Giacino, J.T., Hirsch, J., Schiff, N., Laureys, S., 2006. Functional neuroimaging applications for assessment and rehabilitation planning in patients with disorders of consciousness. Arch. Phys. Med. Rehabil. 87, S67–S76.

Guyon, I., Elisseeff, A., 2003. An introduction to variable and feature selection. J. Mach. Learn. Res. 3, 1157–1182.

Kalmar, K., Giacino, J.T., 2005. The JFK Coma Recovery Scale—Revised. Neuropsychol. Rehabil. 15, 454–460.

Laureys, S., 2004. Functional neuroimaging in the vegetative state. NeuroRehabilitation 19, 335–341.

Laureys, S., 2005. The neural correlate of (un)awareness: lessons from the vegetative state. Trends Cogn. Sci. 9, 556–559.

Laureys, S., Boly, M., 2007. What is it like to be vegetative or minimally conscious? Curr. Opin. Neurol. 20, 609–613.

Laureys, S., Goldman, S., Phillips, C., Van Bogaert, P., Aerts, J., Luxen, A., Franck, G., Maquet, P., 1999a. Impaired effective cortical connectivity in vegetative state: preliminary investigation using PET. Neuroimage 9, 377–382.

Laureys, S., Lemaire, C., Maquet, P., Phillips, C., Franck, G., 1999b. Cerebral metabolism during vegetative state and after recovery to consciousness. J. Neurol. Neurosurg. Psychiatry 67, 121.

Laureys, S., Faymonville, M.E., Degueldre, C., Fiore, G.D., Damas, P., Lambermont, B., Janssens, N., Aerts, J., Franck, G., Luxen, A., Moonen, G., Lamy, M., Maquet, P., 2000. Auditory processing in the vegetative state. Brain 123, 1589–1601.

Laureys, S., Antoine, S., Boly, M., Elincx, S., Faymonville, M.E., Berre, J., Sadzot, B., Ferring, M., De Tiege, X., van Bogaert, P., Hansen, I., Damas, P., Mavroudakis, N., Lambermont, B., Del Fiore, G., Aerts, J., Degueldre, C., Phillips, C., Franck, G., Vincent, J.L., Lamy, M., Luxen, A., Moonen, G., Goldman, S., Maquet, P., 2002. Brain function in the vegetative state. Acta Neurol. Belg. 102, 177–185.

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., Giacino, J.T., Schiff, N.D., Schabus, M., Owen, A.M., 2006. How should functional imaging of patients with disorders of consciousness contribute to their clinical rehabilitation needs? Curr. Opin. Neurol. 19, 520–527.

Laureys, S., Boly, M., Tononi, G., 2008. The neurology of consciousness: functional neuroimaging. In: Laureys, S., Tononi, G. (Eds.), , pp. 31–42.

Laureys, S., Pellas, F., Eeckhout, P.V., Ghorbel, S., Schnakers, C., Perrin, F., Berré, J., Faymonville, M.-E., Pantke, K.-H., Damas, F., Lamy, M., Moonen, G., Goldman, S., 2005. The locked-in syndrome : what is it like to be conscious but paralyzed and voiceless? Prog. Brain Res. 150, 495–511.

León-Carrión, J., Eeckhout, P.v., Domínguez-Morales, M.D.R., Pérez-Santamaría, F.J., 2002. The locked-in syndrome: a syndrome looking for a therapy. Brain Inj. 16, 571–582.

MacKay, D.J.C., 1994. Models of Neural Networks III: Bayesian Methods for Back-propagation Networks. In: Domany, E., et al. (Ed.), , pp. 211–254.

Majerus, S., Gill-Thwaites, H., Andrews, K., Laureys, S., 2005. Behavioral evaluation of consciousness in severe brain damage. Prog. Brain Res. 150, 397–413.

MSTFoPVS, 1994. Multi-Society Task Force on PVS: Medical aspects of the persistent vegetative state. New England journal of medicine 330, 1499–1508.

Müller, K.-R., Mika, S., Rätsch, G., Tsuda, K., Schölkopf, B., 2001. An introduction to kernel-based learning algorithms. IEEE Trans. Neural Netw. 12, 181–202.

Neal, R.M., 1996. Bayesian Learning for Neural Networks.

Raichle, M.E., Mintun, M.A., 2006. Brain work and brain imaging. Annu. Rev. Neurosci. 29, 449–476.

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,

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083

S., 2008. A French validation study of the Coma Recovery Scale—Revised (CRS-R). Brain Inj. 22, 786–792.
Schnakers, C., Perrin, F., Schabus, M., Hustinx, R., Majerus, S., Moonen, G., Boly, M., Vanhaudenhuyse, A., Bruno, M.A., Laureys, S., 2009a. Detecting consciousness in a total locked-in syndrome: an active event-related paradigm. Neurocase 1–7.
Schnakers, C., Vanhaudenhuyse, A., Giacino, J., Ventura, M., Boly, M., Majerus, S., Moonen, G., Laureys, S., 2009b. Diagnostic accuracy of the vegetative and minimally conscious state: clinical consensus versus standardized neurobehavioral assessment. BMC Neurol. 9, 35.
Schölkopf, B., Burges, C.J.C., Smola, A.J., 1999. Advances in Kernel Methods: Support Vector Learning.
Tipping, M.E., 2001. Sparse Bayesian learning and the relevance vector machine. J. Mach. Learn. Res. 1, 211–244.
Vapnik, V.N., 1998. Statistical Learning Theory.

Please cite this article as: Phillips, C.L., et al., "Relevance vector machine" consciousness classifier applied to cerebral metabolism of vegetative and locked-in patients, NeuroImage (2010), doi:10.1016/j.neuroimage.2010.05.083