



Readout of the intrinsic and extrinsic properties of a stimulus from un-experienced neuronal activities: Towards cognitive neuroprostheses

Saba Farbod Kia^{a,b}, Elaine Åstrand^a, Guilhem Ibos^a, Suliann Ben Hamed^{a,b,*}

^a Centre de Neuroscience Cognitive, CNRS-UMR 5229, Université Claude Bernard Lyon I, 67 Bd. Pinel, 69675 Bron Cedex, France

^b Bahá'í Institute for Higher Education, www.bihe.org, Iran

ARTICLE INFO

Keywords:

Attention
Decoding
Brain-machine interface (BMI)
Cognition
Vision
Non-human primate
Readout
Frontal eye field
FEF

ABSTRACT

While sensory and motor systems have attracted most of the research effort in the field neuroprosthetics, little attention has been devoted to higher order cortical processes. Here, we propose a first step in the direction of applying neural decoding to the study and manipulation of visuospatial attention, an endogenous process at the interface between sensory and motor functions. To this aim, we investigate whether the offline activity of a population of non-human primate frontal eye field neurons (FEF) in response to an endogenous cue can be readout *on a trial by trial basis* to provide a precise description of the cue's attributes, namely, its location and identity, but also the allocation of attention following its interpretation. Using a linear decoder, we reach up to 86% correct predictions for the different decoded variables, including the spatial allocation of endogenous attention. We show that the decoding performance drops on incorrect trials, indicating that cue encoding participates to the animal's behavioral performance. Last, we show that the temporal resolution of the decoding influences readout performance. These results are a strong indication of the feasibility of the readout of endogenous variables by standard decoding algorithms, on a suboptimal dataset. However, its validity remains to be proved in a real-time situation.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

The field of neural prosthetics is rapidly developing in the larger field of neuroscience (Lebedev and Nicolelis, 2006; Nicolelis and Lebedev, 2009; Green and Kalaska, 2010; Fetz, 2007; Hatsopoulos and Donoghue, 2009). Its general aim is to use preserved electrophysiological nervous activities in order to counter specific functional dysfunctions or deficits by driving external palliative devices. One of the most representative examples of its applications is the use of the neuronal activity of the motor cortex in a tetraplegic patient to control a computer driven environment, thus allowing him a certain degree of mobility and independence (Hochberg et al., 2006; Donoghue et al., 2007). Another important advance in the field is the demonstration that other regions than the motor cortex can be used to drive neural prostheses, such as parietal cortex (Musallam et al., 2004), or dorso-lateral prefrontal cortex (Vansteensel et al., 2010) hence providing a potential substitute for motor cortex activities when these are not available, following for example an acute injury of this region. This field holds the potential to improve the life of thousands of patients and the research effort is directed to alleviate a diversity of pathological

conditions including injuries of the motor system (e.g. lesions of the motor cortex, cortico-spinal pyramidal tract or spinal cord), certain types of neurodegenerative conditions (e.g. cerebral palsy, amyotrophic lateral sclerosis) or distal limb injuries (e.g. amputations).

However, while most of the research effort in neural prosthetics has concentrated on the use of motor signals to drive external devices, new directions in the field of neuroprostheses are also emerging. For example, a recent study has demonstrated that incorporating sensory feedback to a motor neural prosthesis improves its performance (Suminski et al., 2010). On another line, Musallam et al. (2004) have shown, in the context of a motor behavior, that cognitive signals such as the expected value of a reward, i.e. the subject's motivation, can be decoded from parietal neural activity. The work of Jerbi and collaborators also demonstrates that such signals as attention orientation signals and mental calculation signals can be used to drive a cognitive brain-machine interface (Jerbi et al., 2009).

Our scope in the present work is to test whether other types of cognitive variables in relation with visuospatial processing can be decoded from neuronal populations and thus serve to enhance the functions of neural prostheses. We will specifically focus on endogenous spatial attention which guides the active selection of sensory information based on cognitive priors (for example when looking for the phone, the likelihood of it being on the table is higher than

* Corresponding author at: Centre de Neuroscience Cognitive, CNRS-UMR 5229, 67 Bd. Pinel, 69675 Bron Cedex, France.

E-mail address: sbenha@isc.cnrs.fr (S. Ben Hamed).

on the ceiling, as a result, most visual resources will be focused on the table). This process, which is distinct from subsequent perceptual processes, is a key cortical function. Following parietal acute lesions, hemineglect patients develop the inability to attend and thus to perceive and interact with their contralesional environment (Husain and Nachev, 2007; Riddoch et al., 2010). In order to address this question, we use a database of neurons recorded from the frontal eye fields (FEF) while monkeys were engaged in a variant of a cued target detection which allows to dissociate in time the processes related to the orientation of attention from processes related to target detection (Ibos et al., 2009). The FEF is a prefrontal area that plays a crucial role in the control of eye movements and visual attention (Ibos et al., 2009; Schall et al., 2004; Bruce and Goldberg, 1985), as assessed both from its neuronal response profiles (Bruce et al., 1985) and the effect of its perturbation on behavior (Moore and Armstrong, 2003; Wardak et al., 2006; Bruce et al., 1985). Collecting neuronal data from non-human primates allows a precise targeting of the cortical area of interest over several recording sessions. On the opposite, the electrodes used for intracortical recordings in humans are primarily placed to identify epileptogenic foci and are therefore not optimal to address specific cognitive processes such as those described in the present paper. Here, we provide a quantification of how efficiently the allocation of spatial attention following cue presentation can be readout on a trial by trial basis. We contrast this readout performance with that achieved in reading out the physical attributes of the cue known to be represented in the FEF, such as its position. We also describe several parameters that affect the final prediction of the classifier.

2. Methods

2.1. Description of the neurophysiological database

2.1.1. Behavioral task

The activity of 142 frontal eye field (FEF) neurons was recorded from two macaque monkeys, while performing a cued target detection task (Ibos et al., 2009). The experimental design of this task allows to dissociate in time the processes related to the orientation of attention from those related to target detection. In particular, the cue is an abstract cue whose joint identity and position information need to be combined by the monkey in order to know towards which hemifield he should direct his attention. Briefly, the monkeys had to fixate a central point on the screen (Fig. 1a). Two streams of visual objects were presented, one in the visual receptive field of the neuron being recorded and the other in the contralateral side. One of the streams included a cue which instructed with a certain probability the position of the target. The cue could be green (resp. red), predicting that the target would appear in the same (resp. other) stream. Therefore, the monkey had to combine the information related to both physical attributes of the cue (its location and identity) to find out its instruction. This instruction is thus a cognitive cue eliciting a cognitive response in the brain. The monkey had to release a lever to report the presence of the target. In 67% of the trials, the target appeared in the instructed stream (valid trials), in 17% of the trials, it appeared on the opposite stream (invalid trials), and in 16% of the trials it did not appear at all (catch trials) to discourage systematic responses. The monkey was rewarded for releasing the lever 150–750 ms following target




Fig. 1. (a) Example of an RSVP trial for a valid shift cue condition. This trial being a valid trial, the target is presented in the stream opposite to the cue. On target detection, monkeys have to release the lever between 150 ms and 750 ms following target presentation. They are rewarded for successful detections. (b) Behavioral validity effects of both monkeys M and Z as measured on reaction times (RT) and detection rates (DR).

onset on valid and invalid trials and holding to it on catch trials. Invalid trials were used to check that the monkey used the predictive information provided by the cue in order to optimize his behavior. Sessions in which this was not the case were discarded from the analysis.

2.1.2. Cell categories

The neurons were recorded successively and categorized as a function of their response. Neurons were defined as cue-related neurons when their response following the presentation of at least one type of cue was significantly different from the pre-cue response (t -test, $p < 0.05$). Cue-related neurons were further subdivided into different groups (Ibos et al., 2009): neurons encoding the position of the cue (defined as those cells having a significant difference between the neuronal cue-related response to ipsilateral and contralateral cues on a t -test, $p < 0.05$, and no significant difference between the neuronal cue-related response to cues instructing attention ipsilateral and contralateral on a t -test, $p > 0.05$, $n = 18$), neurons encoding the instructed position of attention (defined as those cells having no significant difference between the neuronal cue-related response to ipsilateral and contralateral cues on a t -test, $p > 0.05$, and a significant difference between the neuronal cue-related response to cues instructing attention ipsilateral and contralateral on a t -test, $p < 0.05$, $n = 20$) and neurons specifically encoding the identity of the cue (defined as the remaining cue-related cells, $n = 43$). These cell categories were exclusive. In all, there were therefore 81 cue-related neurons. Decoding was also performed on the whole dataset including both the cue-related and cue unrelated neurons ($n = 142$).

2.2. Classifier

We used a one-layer perceptron based on the Levenberg–Marquardt backpropagation. Fifty percent of the trials are used for training, 20% for validation and 30% for testing. This network is formally equivalent to a linear decoding approach. More specifically, it is constructed as follows (Fig. 2). Throughout this study, the input layer of the network has one unit per cell recorded in FEF during the whole length of the study, and the total number of input cells depending on the category of interest. Cells are thus artificially considered as one neuronal population. The output layer is a single binary unit corresponding to the state of one possible attribute of the cue: physical position (left/right), instructed position of attention (left/right) or cue identity (stay/shift). Each state can hence take two different values, corresponding to the two different decoded classes. A linear activation function is used to correctly classify the neural data. Then, the outputs of the network are clipped using a symmetrical hard limit function to one of two discrete values (-1 and 1) corresponding to the attributes of the cue. This clipping procedure thus applies a classification threshold at zero, corresponding to the mean value between the two theoretical outputs (1 and -1). This is legitimate given the fact that the two decoded classes are homogeneously distributed within the data sample. **Training:** The input data is the cell activity averaged over 100 ms time bin centered around the time point where responses of interest have been shown to be maximal (200 ms following cue). For each cell, the network is trained on 50% of randomly drawn trials among those available. Training is performed using a Levenberg–Marquardt backpropagation and hyperbolic tangent sigmoid transfer functions. The weights of the activation function are therefore updated for every run to optimize and minimize




Fig. 2. Network architecture. The input layer encodes the activity of N FEF neurons following cue presentation ($N = 18, 20, 43, 81$, depending on the functional category of interest). The output layer contains 1 unit. The network is trained to encode the predicted attribute of the cue amongst two possible states, given current input activities. Training involves optimizing weights using a Levenberg–Marquardt back propagation algorithm and a tangsig transfer function. A symmetric hard limit function is applied to the 1 output unit in order to clip the readout to two discrete values, -1 and 1 .

the square distance between the estimate of the cue attribute and its true value. **Stopping:** Training is stopped using a regular stopping technique based on the reversal of the performance trend of the network in response to a second subset of trials, called the validation set (20% of the remaining trials drawn randomly). **Testing:** The performance of the network is then tested for this time point (200 ms post-cue) on the remaining trials, never experienced by the network, as well as on all trials for activities averaged over successive 100 ms ranging from 200 ms before the training time reference to 600 ms following it, with a 1 ms resolution. The readout performance of the network is defined as the percentage of match between the outputs of the network and the actual physical state of the cue on the corresponding trials. This performance thus reflects whether the network can predict the state of the cue when it is experiencing neuronal inputs that were collected when the monkey needed to interpret it for a successful behavior. The major assumption here is that monkeys had to use the cue to orient their attention to the stream of stimuli in which the target has the highest probability to appear in order to succeed in the trial and be rewarded. The behavioral performance of the animals on invalidly cued trials with respect to validly cued trials indicates that they are indeed using the predictive information held by the cue

(Fig. 1b). The training testing procedure is repeated over 15 runs, each run varying in the initial weight matrix and in the training/validation/testing trial sets (i.e. this is a 15-folds cross-validation, with random initialization of the weights for each classifier). Each run yields a decoding performance defined as the percentage of correct output predictions as compared to the actual output state, when the network is observing the successive population activities of the test dataset. In other words, this decoding performance corresponds to the probability that, on a given trial, the network provides a correct guess. The main measure that will be discussed in the following is the average decoding performance over the 15 successive runs.

3. Results

Incoming sensory stimuli can be described by their physical attributes such as their position, color or shape. These physical exogenous attributes are independent of the subject himself. However, incoming sensory stimuli can also have more complex attributes that are related to how the subject interprets them. For example, is the stimulus rewarding or not? Is it predictive of some change in the environment? Is it instructing a behavior? This type




Fig. 3. Readout of (a) cue instruction, (b) cue position and (c) cue identity from different FEF neuronal subpopulations. Left panels: average activity of cells selective of the instructed position of attention (a), cue position (b) or cue identity (c) in response to the preferred (solid line) or non-preferred condition (dashed line) and significant differences (gray segment, $p < 0.05$). Cue onset is denoted by the vertical line. Middle panels: readout performance (mean $\pm$ s.d.) as a function of time around cue onset (vertical line), compared to chance level (dashed line). Right panels: average readout performance over the two different time windows shaded in the middle panel. Att: cells selective for position of attention; Pos: for cue position; ID: for cue identity. Error bars indicate $\pm$ s.d. The gray horizontal shaded area corresponds to the average readout performance of the position of the first stream of visual stimuli ($\pm$ s.d.).

of endogenous attributes arises from the subject himself. In this context, the cue has two exogenous attributes, its physical position and its identity/color. The cue also has an endogenous attribute, resulting from cue interpretation, namely, the desired final position of attention allocation. This endogenous attribute is not confounded with either the position or the identity of the cue and requires that the monkey processes both of these attributes to correctly direct his attention in space: both a red cue on the left and a green cue on the right instruct attention to the right while both a red cue on the right and a green cue on the left instruct attention to the left. It is important to note here that while in humans, a verbal report of how the cue was interpreted by the subjects can be asked for (did the cue instruct to hold attention or to shift it elsewhere?), this cannot be done in monkey experiments. However, the design of the task allows us to indirectly infer the subject's interpretation of each key element of the task. Indeed, the target detection difficulty is adjusted so that, on average, a behavioral benefit (shorter reaction times and higher detection rates on valid trials than on invalid trials) is drawn from the correct interpretation of the cue (Fig. 1b). In the following, we will focus on the cue related signals and we will compare the readout performance of the endogenous attributes of this key element of the environment to the readout performance of its exogenous attributes.

3.1. Readout of cue attributes

Based on a statistical analysis of their response characteristics to the cue, the cue-related cells were classified into three distinct groups: one group selectively encoding the location of the cue (*Position* cells, Fig. 3a, left hand plot), another group selectively encoding the identity of the cue (*Identity* cells, Fig. 3b, left hand plot) and a last group selectively encoding the instructed position of attention (*Attention* cells, Fig. 3c, left hand plot). In the following, we explore how these different cells contribute to the readout performance of each of the cue attributes. For comparison, the average readout performance achieved on decoding the position of the first stream of visual stimuli from all FEF cells is plotted on all performance data plots (Fig. 3, right panels). This performance is indicative of the performance of the decoding algorithm at reading out the position of a simple visual stimulus.

3.1.1. Cue instruction

A first network was trained to predict the instructed position of attention by the cue on a single trial basis, as described in the methods, either from the specific subpopulation of *Attention* cells (Fig. 3a, middle plot) or from different subpopulations of cue cells (Fig. 3a, right hand plot). A maximum readout performance of $82.5 \pm 2\%$ is achieved in the 150–250 ms time window following cue onset. Because *Identity* cells also hold partial information about the instructed position of attention, the majority of them representing either the contralateral *shift* cue (i.e. instructing attention ipsilateral to the cue) or the contralateral *stay* cue (i.e. instructing attention at the location of the cue), readout performance of cue instruction is also carried out on the joint *Attention + Identity* cue populations. The obtained performance is significantly higher from the previous one ($85 \pm 2\%$, $p < 0.0001$). Decoding of the instructed position of attention from all cue population is slightly lower than that achieved with the joint *Attention + Identity* cue populations but significantly higher than that achieved with just the *Attention* cue cells ($84 \pm 2\%$, $p < 0.0001$). However, decoding from the whole FEF population introduces noise and slightly deteriorates readout performance with respect to decoding from just *Attention* cue cells ($80 \pm 7\%$, $p < 0.0001$).

Interestingly, when measured further away from cue onset, in the 450–550 ms time window, the readout performance for cue instruction remains high when decoded from the *Attention*

cue population ($73.5 \pm 2\%$), the *Attention + Identity* cue population ($73 \pm 3\%$, $p = 0.0001$) or all the cue population ($72.5 \pm 2\%$, $p = 0.0001$). However, this performance drops when decoded from all the FEF population ($68 \pm 3\%$, $p = 0.0001$), suggesting that the cue instruction information is held over time and that this is essentially achieved by the subpopulation of cells selected on the basis of its cue-related responsiveness. In order to compare the readout performance achieved when decoding the cognitive instruction held by the cue to that achieved when decoding lower order cue properties, we explore in the following how well the identity of the cue or its position can be readout from the different cue-related populations.

3.1.2. Cue position

The maximum readout performance achieved when decoding the position of the cue from *Position* cells in the 150–250 ms time window following cue onset is equal to $76 \pm 2\%$ and it is significantly different from that achieved on decoding the instructed position of attention from *Attention* cells ($p < 0.0001$). Here again, *Identity* cells also hold partial information about the position of the cue, as the majority of them represents a contralateral cue. Readout performance of cue position was thus carried out on the joint *Position + Identity* cue populations. The obtained performance is significantly improved with respect to the previous one ($81 \pm 4\%$, $p = 0.0001$). It remains intermediate between these two values when decoding is achieved on all cue population ($80 \pm 4\%$, $p = 0.0001$) and drops again when achieved on all FEF population ($76.5 \pm 4\%$, $p = 0.2127$). This contrasts with what is seen for the readout of cue instruction, as the total FEF population seems to hold as much information about cue position than the *Position* cells themselves. In the 450–550 ms time window, the readout performance for cue position drops around 60% (60% for *Position* cells, 61% for joint *Position + Identity* cells, 58% for all cue cells and 56% for all FEF cells), suggesting that the encoding of cue position in the context of the ongoing behavior is not sustained over time. Overall, it is surprising that the readout of cue position is significantly worse than the readout of the instructed position of attention whatever the subpopulation considered ($p < 0.0001$ for all comparisons). This is a strong indication that endogenous information can be decoded as good if not better from FEF neuronal responses than exogenous information.

3.1.3. Cue identity

The majority of *Identity* cells represents contralateral *shift* and *stay* cues and does not encode any of the ipsilateral cues. Therefore, on decoding cue identity, we actually measure the performance of the network in predicting the identity of a contralateral cue. Accordingly, the maximum readout performance achieved from *Identity* cells in the 150–250 ms time window following cue onset is equal to $86 \pm 4\%$ and is significantly different from that achieved on decoding the instructed position of attention from *Attention* cells ($p < 0.0001$). The readout of cue identity thus outperforms that of cue instruction and cue position, corroborating the fact that the FEF receives information both from the dorsal “where” and the ventral “what” pathways. This performance is significantly deteriorated when other cells are included in the decoded population (Fig. 3c, left hand plot). However, as in the case of cue position, the readout performance drastically drops in the 450–550 ms time window (below 50%, Fig. 3c, left hand plot), suggesting that cue identity information is not maintained in the FEF in the context of the present task.

As a result, it appears that FEF holds information both about the physical attributes of the cue (identity and position) as well as of its endogenous representation in relation to the ongoing behavior (the instructed position of attention). Interestingly, these three types of information can be decoded with very high performance

rates, the instructed position of attention and cue identity being readout best, with a performance above 80% correct.

3.2. Effect of the animal's behavioral performance on the readout of the physical attributes of the cue

The task was designed so that the monkey needed to correctly interpret the cue in order to optimize its behavioral performance. Invalid cue trials in which the cue does not correctly predict the position of the target allow us to check that this is the case. Indeed, on these trials, the animals' reaction times are significantly slower and their performance significantly deteriorated with respect to valid trials, thus reflecting the behavioral cost of cue invalidity (Fig. 1). This proves that, on average, the cue is used by the monkeys to improve their performances in the task. The question we are asking here is whether the encoding of the cue, as reflected by the population response is different on incorrect trials than on correct trials. We therefore presented the networks trained to decode the different cue parameters (position, identity and instruction) with the population response on incorrect trials and compared the performance to that obtained for correct trials (Fig. 4, dark gray, correct trials, light gray, incorrect trials). The readout performance for the position of the cue ($80 \pm 2\%$ vs. $61.4 \pm 6.1\%$ for correct vs. incorrect trials respectively, $p < 0.0001$), the identity of the cue ($86.1 \pm 2.3\%$ vs. $65.2 \pm 7.1\%$, $p < 0.0001$) and the position of attention ($78.5 \pm 2.8\%$ vs. $62.4 \pm 6\%$, $p < 0.0001$) are significantly deteriorated on incorrect trials with respect to correct trials, though they remain significantly different from chance ($p < 0.0001$). This suggests that a miss-encoding of the cue has a partial causal relationship with the animal's behavioral performance. Obviously, the encoding of the target is also expected to have a causal effect on behavioral performance, error trials arising from a conjunction of miss-encoding of both the cue and the target.

3.3. Effect of population size on readout performance

The results presented up to now were carried out on subpopulations of cells selected on the basis of their response pattern to the cue or to the target. An important question to investigate in this context is the relationship between readout performance and the size of the neuronal population. Here we focus on the readout per-

formance of the instructed position of attention and we look at how this performance evolves as cells get randomly dropped out from the decoding process (Fig. 5). As expected, readout performance decreases as the number of cells decreases, both when considering the *Attention* cue cells only or all the FEF cells, though it is worth noting that with as few as 3 cells, readout performance is already significantly above the 50% chance level. Not surprisingly, the decoding performance starts higher, improves and tends towards a saturation asymptote faster when specifically adding *Attention* cue cells to the decoded population than when adding FEF cells independently of their functional response.

3.4. Effect of trial number on readout performance

Because our decoding is performed offline, the different networks were trained and tested on all available trials. However, in a real online decoding situation, as the subject starts working, the trials that can be used for training and testing increase progressively. As trial number increases, both additional information and additional noise are expected to affect decoding. In order to quantify this, we simulated an online decoding situation in which training takes place on the previously acquired population activities and testing is performed on the last recorded activities, as trials are added one by one to the training–testing database. Specifically, we measured the readout performance of the instructed position of attention from the *Attention* cue cells and *All* FEF cells as trial number increases (Fig. 6). The central result of this analysis is that readout performance steadily increases and stabilizes after an average of 60 trials.

3.5. Effect of temporal resolution on the readout of cue instruction

Although the decoding is performed on 1 ms time-steps, the activity is averaged over 100 ms time bins. The resulting effect is a reduction of the noise and a smoothing of the neuronal response variations due to noise. Here, we compare the position of attention readout performance obtained for activities averaged over different window lengths (Fig. 7). We show that decoding of cue instruction from all FEF population improves steadily as bin width increases although this improvement stabilizes for analysis time windows beyond 100 ms. While this difference in decoding




Fig. 4. Cue attribute readout performance as a function of the animal's behavioral performance. In dark gray, decoding performance on trials on which the monkey successfully performed the task. In light gray, decoding performance is shown on trials on which the monkey failed in correctly reporting the target's presence. *** $p < 0.0001$.




Fig. 5. Readout performance of the instructed position of attention as a function of the number of cells. Decoding is performed either on the population activity averaged over the 150–250 ms time window following cue onset (* symbols, gray: total population, black: attention population) or as a control on the 150–250 ms time window preceding cue onset (+ symbols) and is averaged, for each cell number tick, over 10 different random combination of cells.




Fig. 6. Readout performance of the instructed position of attention as a function of the number of trials. Decoding is performed either on the population activity averaged over the 150–250 ms time window following cue onset (* symbols, gray: total population, black: attention population) or as a control on the 150–250 ms time window preceding cue onset (+ symbols) and is averaged, for each trial number tick, over 10 different random combination of trials.




Fig. 7. Readout performance as a function of neuronal activities averaged over different time window lengths.

performance between small and larger bins can seem surprisingly important, it is worth noting that the information that is being decoded here is not a sensory information but rather a cognitive information involving an endogenous processing of the cue. Indeed, the cue is embedded into a visual stream of stimuli. We have shown elsewhere (Ibos et al., 2009) that the FEF does not encode the transient changes in the stream of visual stimuli and that all the observed response modulations (apart from the response to stream onset), both at the single cell and at the population levels are due to the cognitive processing of either the cue or the target. While sensory responses are expected to be locked to the stimulus onset, cognitive processes are expected to have a more variable temporal dynamics. Corroborating this point, the cue responses have an overall higher onset variability (s.d. = 18.75 spikes/s) than the visual response to the first stream onset (s.d. = 15.76, $p < 0.000001$) on an F -test for variance equality. This might explain in part why larger bins provide more information than small time analysis bins. For comparison, Musallam et al. achieve their decoding of the monkey's intention of movement from the 900 ms inter-

val of the delay period implemented in the behavioral task the monkey is performing (Musallam et al., 2004). While their decoding performance is comparable to the one described here, the temporal resolution is several orders higher than the one achieved in the present work. From a neuroprosthetics point of view, the trade-off between the improved decoding performance on larger time averaging bins and a decreased time resolution is an important factor to be considered.

4. Discussion and conclusion

This work shows that the cognitive variables related to processes underlying endogenous attention can be reliably decoded from the single trial FEF neurons activity with an accuracy ranging from about 76% to 86%, giving it a high reliability. Such results confirm and extend the work of Musallam et al. (2004) who for the first time show that cognitive intention of movement signals can be used for the purpose of a motor neuroprosthesis, thus paving the way toward new cognitive neural prostheses. We consider the present report as a starting point towards integrating perception-related variables into motor neuroprostheses. These results are based on the analysis of single neurons recorded independently. These recordings are therefore free of correlations which have been described to be detrimental to information content in specific situations. Indeed, theoretical studies suggest that the correlations in the activity of neurons in response to noise can either increase or decrease the information that can be read out from their activities depending on the spatial distribution of the recorded cells (Abbott and Dayan, 1999; Panzeri et al., 1999), which therefore, rules out a possibility to draw a generalized rule about their impact and necessitates to discuss the effects of correlations case by case (Averbeck et al., 2006). This taken into consideration, it turns out that in most examples of empirical data, correlations add to the information (Oram et al., 1998; Hung et al., 2005; Panzeri et al., 1999, 2003; Averbeck and Lee, 2006), suggesting that the performance we obtain here is a lower bound and can still be improved.

It is worth noting that the decoding was performed here using a linear decoding strategy which is compatible with the underlying neuronal properties (Ben Hamed et al., 2003, 2007; Pouget and Sejnowski, 1997). It is not clear whether other biologically plausible approaches such as Bayesian decoding (Oram et al., 1998) would outperform the present results. Given the underlying theoretical framework, it would be important to test this point although experimentalists might run into the major limitation of this approach, namely, the need for enough trials per specific state to provide a satisfactory description of the probability density response function of each neuron. This performance can potentially be improved by adopting more performance oriented algorithms such as a support vector machine approach (Morra et al., 2010; Hung et al., 2005). Other methods of decoding, involving for example, temporal decoding (Zacksenhouse and Ahissar, 2006) instead of or in addition to spike decoding, or extracting information both from spikes and local field potentials (Belitski et al., 2008), are also worth of interest for the present purpose.

While we precisely dissociate in this work the possibilities of reading out exogenous and endogenous attributes of stimuli, we only had two states to contrast and decode from every time. This can be considered as the most basic initial decoding step as it amounts to classifying population neuronal activities into two categories. Recent work has shown the feasibility of decoding several states of a given cognitive variable (Musallam et al., 2004) or continuous states of motor variables (Hochberg et al., 2006). It would be important to investigate whether this level of readout precision can be achieved for spatial endogenous variables.

Finally, when undergoing the present analyses, a crucial question was (1) whether decoding performance relied on the selection of the proper neuronal population specifically encoding the variable of interest, (2) whether it remained rather insensitive to the noise added by the non-responsive cells (3) or else, whether non-responsive cells contributed to total population information and thus to overall decoding performance. Overall, we show that the total population holds comparable or even more information available to the classifier about each variable than specific subpopulations. Although a more precise evaluation of this question is required, this is interesting from a neuroprosthetic point of view, because it suggests that recording from the largest population with no need for pre-selection of cells – which would induce a computational time cost – can be achieved without incurring the cost of losing information. This is also interesting from a physiological point of view, as it suggests that cells that cannot be shown to statistically encode a given variable still do contribute to the total population information, suggesting that single cell physiology might benefit from considering both single cell and population statistics when addressing the functional role of a given cortical area.

In conclusion, we believe that this work is a starting point to refine the decoding of motor intentions, improve neural prostheses and ultimately lead to novel rehabilitation techniques by introducing endogenous cognitive variables in relation with spatial processing and decision making.

Acknowledgments

This work has been supported by a CNRS NeuroInformatic Grant (2009). E. Åstrand is supported by a joint CNRS-DGA Grant. We thank the anonymous reviewers for their constructive comments on the present work.

References

- Abbott, L.F., Dayan, P., 1999. The effect of correlated variability on the accuracy of a population code. *Neural Comput.* 11, 91–101.
- Averbeck, B.B., Lee, D., 2006. Effects of noise correlations on information encoding and decoding. *J. Neurophysiol.* 95, 3633–3644.
- Averbeck, B.B., Latham, P.E., Pouget, A., 2006. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* 7, 358–366.
- Belitski, A., Gretton, A., Magri, C., Murayama, Y., Montemurro, M.A., Logothetis, N.K., Panzeri, S., 2008. Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *J. Neurosci.* 28, 5696–5709.
- Ben Hamed, S., Page, W., Duffy, C., Pouget, A., 2003. MSTd neuronal basis functions for the population encoding of heading direction. *J. Neurophysiol.* 90, 549–558.
- Ben Hamed, S., Schieber, M.H., Pouget, A., 2007. Decoding M1 neurons during multiple finger movements. *J. Neurophysiol.* 98, 327–333.
- Bruce, C.J., Goldberg, M.E., 1985. Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* 53, 603–635.
- Bruce, C.J., Goldberg, M.E., Bushnell, M.C., Stanton, G.B., 1985. Primate frontal eye fields. II. Physiological and anatomical correlates of electrically evoked eye movements. *J. Neurophysiol.* 54, 714–734.
- Donoghue, J.P., Nurmikko, A., Black, M., Hochberg, L.R., 2007. Assistive technology and robotic control using motor cortex ensemble-based neural interface systems in humans with tetraplegia. *J. Physiol. (London)* 579, 603–611.
- Fetz, E.E., 2007. Volitional control of neural activity: implications for brain–computer interfaces. *J. Physiol. (London)* 579, 571–579.
- Green, A.M., Kalaska, J.F., 2010. Learning to move machines with the mind. *Trends Neurosci.* <<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC21176975/>>.
- Hatsopoulos, N.G., Donoghue, J.P., 2009. The science of neural interface systems. *Annu. Rev. Neurosci.* 32, 249–266.
- Hochberg, L.R., Serruya, M.D., Friehs, G.M., Mukand, J.A., Saleh, M., Caplan, A.H., Branner, A., Chen, D., Penn, R.D., Donoghue, J.P., 2006. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442, 164–171.
- Hung, C.P., Kreiman, G., Poggio, T., DiCarlo, J.J., 2005. Fast readout of object identity from macaque inferior temporal cortex. *Science* 310, 863–866.
- Husain, M., Nachev, P., 2007. Space and the parietal cortex. *Trends Cogn. Sci. (Regul. Ed.)* 11, 30–36.
- Ibosa, G., Duhamel, J.-R., Ben Hamed, S., 2009. Differential contributions of parietal and frontal cortices to perceptual decision and attention control. Society for Neuroscience Annual Meeting.
- Jerbi, K., Freyermuth, S., Minotti, L., Kahane, P., Berthoz, A., Lachaux, J.-P., 2009. Watching brain TV and playing brain ball exploring novel BCI strategies using real-time analysis of human intracranial data. *Int. Rev. Neurobiol.* 86, 159–168.
- Lebedev, M.A., Nicolelis, M.A.L., 2006. Brain–machine interfaces: past, present and future. *Trends Neurosci.* 29, 536–546.
- Moore, T., Armstrong, K.M., 2003. Selective gating of visual signals by microstimulation of frontal cortex. *Nature* 421, 370–373.
- Morra, J.H., Tu, Z., Apostolova, L.G., Green, A.E., Toga, A.W., Thompson, P.M., 2010. Comparison of AdaBoost and support vector machines for detecting Alzheimer's disease through automated hippocampal segmentation. *IEEE Trans. Med. Imaging* 29, 30–43.
- Musallam, S., Corneil, B.D., Greger, B., Scherberger, H., Andersen, R.A., 2004. Cognitive control signals for neural prosthetics. *Science* 305, 258–262.
- Nicolelis, M.A.L., Lebedev, M.A., 2009. Principles of neural ensemble physiology underlying the operation of brain–machine interfaces. *Nat. Rev. Neurosci.* 10, 530–540.
- Oram, M.W., Földiák, P., Perrett, D.I., Sengpiel, F., 1998. The “Ideal Homunculus”: decoding neural population signals. *Trends Neurosci.* 21, 259–265.
- Panzeri, S., Schultz, S.R., Treves, A., Rolls, E.T., 1999. Correlations and the encoding of information in the nervous system. *Proc. Biol. Sci.* 266, 1001–1012.
- Panzeri, S., Pola, G., Petersen, R.S., 2003. Coding of sensory signals by neuronal populations: the role of correlated activity. *Neuroscientist* 9, 175–180.
- Pouget, A., Sejnowski, T.J., 1997. A new view of hemineglect based on the response properties of parietal neurones. *Philos. Trans. R. Soc. B-Biol. Sci.* 352, 1449–1459.
- Riddoch, M.J., Chechlacz, M., Mevorach, C., Mavritsaki, E., Allen, H., Humphreys, G.W., 2010. The neural mechanisms of visual selection: the view from neuropsychology. *Ann. N. Y. Acad. Sci.* 1191, 156–181.
- Schall, J.D., 2004. On the role of frontal eye field in guiding attention and saccades. *Vision Res.* 44, 1453–1467.
- Suminski, A.J., Tkach, D.C., Fagg, A.H., Hatsopoulos, N.G., 2010. Incorporating feedback from multiple sensory modalities enhances brain–machine interface control. *J. Neurosci.* 30, 16777–16787.
- Vansteensel, M.J., Hermes, D., Aarnoutse, E.J., Bleichner, M.G., Schalk, G., van Rijen, P.C., Leijten, F.S.S., Ramsey, N.F., 2010. Brain–computer interfacing based on cognitive control. *Ann. Neurol.* 67, 809–816.
- Wardak, C., Ibosa, G., Duhamel, J.-R., Olivier, E., 2006. Contribution of the monkey frontal eye field to covert visual attention. *J. Neurosci.* 26, 4228–4235.
- Zacksenhouse, M., Ahissar, E., 2006. Temporal decoding by phase-locked loops: unique features of circuit-level implementations and their significance for vibrissal information processing. *Neural Comput.* 18, 1611–1636.