
Volume 178, issue 2, 15 April 2009
ISSN 0165-0270
ELSEVIER
JOURNAL OF
NEUROSCIENCE
METHODS
Editors-in-Chief
Vincenzo Crunelli
Greg A. Gerhardt

Journal of Neuroscience Methods 178 (2009) 350–356

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

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

A noninvasive, fast and inexpensive tool for the detection of eye open/closed state in primates

Rea Mitelman[a,b,*,1], Mati Joshua[a,b,1], Avital Adler[a,b], Hagai Bergman[a,b,c]

[a] Department of Physiology, The Hebrew University – Hadassah Medical School, Jerusalem 91120, Israel
[b] The Interdisciplinary Center for Neural Computation, The Hebrew University, Jerusalem 91904, Israel
[c] Eric Roland Center for Neurodegenerative Diseases, The Hebrew University, Jerusalem 91904, Israel

ARTICLE INFO

Article history:
Received 5 October 2008
Received in revised form 4 December 2008
Accepted 4 December 2008

Keywords:
Electrophysiological recordings
Image processing
Eyeblink conditioning
Primates
Eyelid

ABSTRACT

Accurate detection of the eye state (i.e., open or closed) of animals during electrophysiological recordings is often crucial for analyzing physiological data. This requires a system which is reliable, and preferably noninvasive and inexpensive. Here we present such a tool incorporating a standard digital camera and a semi-automatic eye state detection (ESD) algorithm that can be used easily in typical primate electrophysiological setups.

The ESD algorithm is based on the high light absorbance of the iris and pupil relative to the eyelid and takes advantage of the unique conditions found in primate physiological recordings (minimal area of sclera and head fixation). The ESD algorithm is as accurate as a human observer, and is not vulnerable to variance inherent to human decisions that it requires (i.e., eye location setting, training set classification and threshold setting). The temporal resolution with standard interlaced digital cameras is 17–20 ms. This is sufficient for the detection of eye state changes during electrophysiological recordings including spontaneous blinking and eye blink conditioning, as demonstrated here. Furthermore, the ESD tool can be applied to other physiological areas of research in which changes in eye state are critical to analyzing neuronal activity.

1. Introduction

Vision is the main sense by which primates (both human and nonhuman) perceive the world. Unlike other senses, visual input can be completely blocked at the level of the sensory organ by the eyelid. Therefore, understanding any neuronal activity involving the visual system requires an accurate recording of the state of the eyelids, i.e., whether they are open or closed. Furthermore, detecting the state of the eyelid is crucial for monitoring motor output during eyeblink conditioning (Marquis and Hilgard, 1937). Finally, detection of eyeblink enables the study of the natural frequency of blinking, which is altered in different pathological states such as schizophrenia or Parkinson's disease (Ponder and Kennedy, 1927; Stevens, 1978; Karson, 1983).

Several methods have been suggested for detection of the eye state of primates. One useful technique is electromyography (EMG) of the orbicularis oculi, the main muscle that is involved in blinking movement, and detecting its activation during eye closure (Silverstein et al., 1978; Blazquez et al., 2002). Another more direct method attaches the eyelid by a wire to a microtorque potentiometer that can measure its movements (Pennypacker et al., 1966). These methods are somewhat invasive, and it is unclear how these devices influence the natural movements of the eyelid.

Less invasive ways include connecting an electromagnetic search coil to the eyelid (Robinson, 1963; Porter et al., 1993). Here, a wire coil is secured to the upper eyelid of the animal, and placed in a weak magnetic field. This generates a current in the coil that is proportional to the angular velocity of the eyelid, thus enabling detection of changes in the state of the eye. Another noninvasive method uses an infrared light-emitting diode (LED) and a photo sensor (Thompson et al., 1994; Clark and Zola, 1998). However, this method requires placing the detector at a distance of 4–5 mm from the animal's eye, which may block significant parts of its field of view. These methods may be irritating to the primates, and therefore could influence their behavior.

The least invasive method that has been used by researchers is direct detection by a human observer. This is usually done offline, after videotaping the animal's behavior (e.g., Nevet et al., 2004). However this method is very cumbersome and time consuming, and therefore is not feasible for processing large amounts of data. Furthermore, human observers are prone to mistakes when asked

* Corresponding author at: Department of Physiology, The Hebrew University – Hadassah Medical School, POB 12272, Jerusalem 91120, Israel. Tel.: +972 2 6757388; fax: +972 2 6439736.
E-mail address: ream@alice.nc.huji.ac.il (R. Mitelman).

[1] These authors contributed equally.

0165-0270/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.jneumeth.2008.12.007

to classify long video sequences and may be biased by their a priori expectations.

Several automatic visual analysis based methods have been suggested for eye state detection in other mammals. In humans there are several algorithms (Tian et al., 2000; Miyakawa et al., 2004; Benoit and Caplier, 2005; Tan and Zhang, 2006; Heishman and Duric, 2007), but they are rather complex, and do not take into account some of the differences between human and non-human primates (e.g. the difference in the sclera's relative size). Moreover, these algorithms are primarily designed for non-scientific goals such as driver fatigue detection, and are intended to achieve impressive stability under unsupervised circumstances. On the other hand, they do not take advantage of the typical primate physiological recording setting, and fall below the performance level of human observers. A system that was suggested for use in rabbits (Bracha et al., 2003) has the disadvantage of attaching markers on the upper and lower eyelid of the animal and therefore is less suitable for daily repeating recording sessions that are typical of physiological studies of awake behaving primates.

In this manuscript we suggest a simple, noninvasive and inexpensive video-based method to detect the state of the eye of primates under head immobilization conditions. The system takes advantage of the typical setting of primate physiological experiments, and operates on the basis of minimal changes in the position of the eyes during a recording session. The video camera can be positioned at a distance from the monkey (depending on its zoom properties) and therefore does not obscure the visual field and does not modify natural blinking behavior. The method is also highly accurate, with a performance level equivalent to that of a human observer (a mean normalized error of 0.15%). Furthermore, since this method works with infrared videotaping, the eye state can be detected in a dark environment.

2. Materials and methods

The tool we describe in this paper includes standard hardware and simple custom-made software. We present the hardware we used in the experiments, and the way we chose to implement the algorithm, although any equivalent hardware and software implementation can be employed.

2.1. Physical setup and data acquisition

All experimental protocols were performed in accordance with the National Institutes of Health Guide for the Care and Use of Laboratory Animals and with Hebrew University guidelines for the use and care of laboratory animals in research, supervised by the institutional animal care and use committee. Briefly, monkeys went through an operation during which a head holder and a recording chamber were attached to their head. During recording sessions, the monkeys' heads were immobilized and microelectrodes were advanced into different targets in the basal ganglia.

A standard infrared digital surveillance camera was used to digitally record the monkey's facial movements (AVer-s 2.54, AverMedia Systems, Taipei, Taiwan). The recording was done in an interlaced mode, with sampling rate of 25 frames per second (PAL mode). In the interlaced mode, each frame is composed of two separately sampled fields: one occupying the even rows and the other the odd ones, without smoothing them. Movies were saved in AVI format in 640×480 pixel resolution, with a grayscale color depth of 8 bits (i.e., 256 levels of infrared brightness).

All data analysis was done in Matlab (Version 7.5, R2007b, The MathWorks). Movies or single frames were easily imported to Matlab, such that each frame is a single brightness matrix and an entire movie is a hypermatrix. To improve performance, importing was done in blocks of a few dozen frames. Each frame was de-interlaced

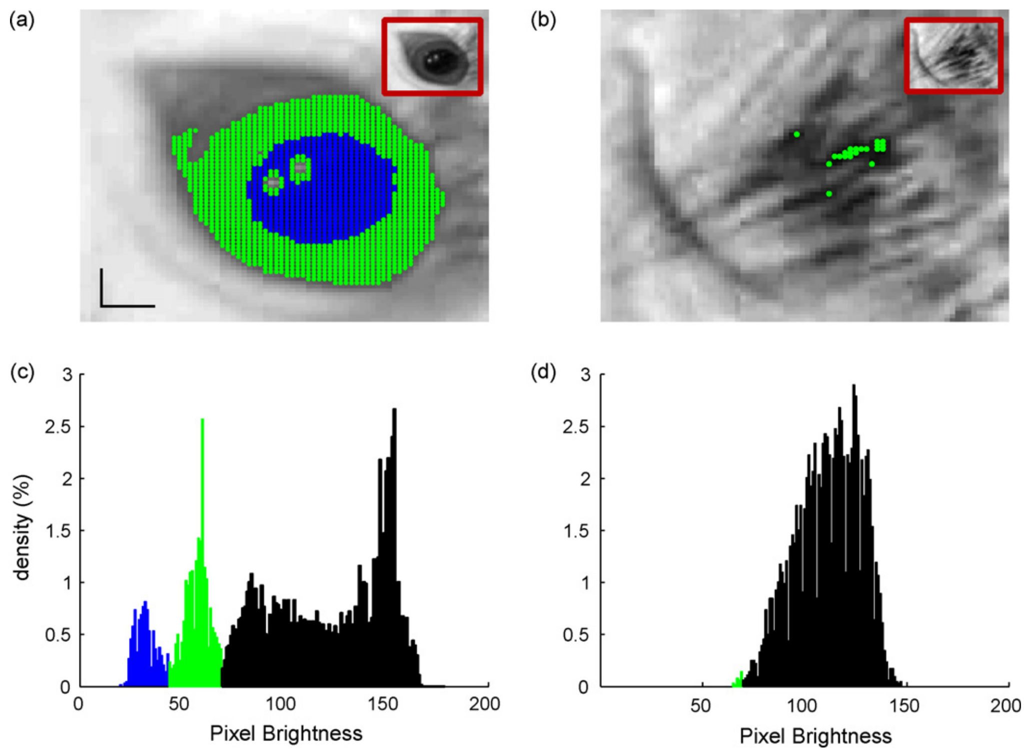

Fig. 1. Example of density histograms of the brightness of a closed and open eye. (a and b) Randomly chosen open (a) and closed (b) eye field. Pixels in two ranges of brightness are color marked, and the original eye fields are shown in the inset. The darker hue, marked in blue, is seen specifically in the pupil, and the intermediate hue, marked in green, is seen mostly in the iris. Scale indicates 10 pixels horizontally and vertically. (c and d) Brightness histograms of the corresponding pictures. The two peaks were manually marked, and the color codes are as in a and b. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

to its two fields, and missing lines were interpolated by averaging each adjacent pair of existing lines. This recreated a full-sized image and doubled the sampling rate to 50 images per second.

2.2. Algorithm description

The eye state detection (ESD) algorithm is based on the difference in light absorbance between the eyelid and the eye itself – the pupil, as well as the iris. Unlike humans, non-human primates have a relatively small sclera, so the pupil and the iris occupy most of the eye opening space. Visible light, as well as infrared radiation, is absorbed by the pupil and the iris considerably more than it is from the eyelid (Durkin et al., 1990; Thompson et al., 1994). As a result, in an open eye image, a certain number of pixels are dramatically darker than all other pixels, whereas in a closed eye image there are hardly any such dark pixels. This can be seen in Fig. 1 which plots the brightness histogram of the area of the eye. The brightness histogram of a typical open eye has two peaks that do not appear in the unimodal brightness histogram of the closed eye. The darkest peak originates from the high light absorbance of the pupil, and the second dim peak from the iris. As outlined below, proper detection of these peaks is enough for correct automatic classification of the state of the eye.

The ESD algorithm is semi-automatic and requires three quick human decisions. First, the user is asked to indicate the location of the monkey's eye (termed "eye field") in an arbitrarily chosen frame, by marking two opposite corners of a rectangle (Fig. 2a). Since the head was immobilized during the experiment reported here, this rectangle only needed to be marked once per experimental day.

The next step is training the algorithm. Eye fields from the video are chosen randomly by the algorithm, and are presented to the user. The user classifies the state of the eye in each eye field as open, closed, or inconclusive (Fig. 2b). This step is completed when the user determines that enough eye fields of both conclusive states have been classified (usually about 20 fields in total). Most videos contain more open eye fields than closed ones and a similar ratio is therefore found in the training set.

The last step is to set the thresholds for the classification: brightness threshold and eye state threshold. This is done by pooling the eye field matrices for each conclusive state. This yields two brightness histograms, one for the open and one for the closed eye (Fig. 2c). The brightness histogram of the open eye fields consistently includes two peaks of darker pixels that fail to appear in the brightness histogram of the closed eye. The user is asked to set a brightness threshold that includes the maximum area of these peaks and the minimum area of the closed eye histogram (the dashed line in Fig. 2c). All pixels darker than this threshold are considered "black" for the following stage of the algorithm.

The ESD algorithm calculates the number of black pixels in each eye field in the training set. The closed eye field with the maximal number of black pixels and the open eye field with the minimal number of black pixels are defined as 'anchors'. The average number of black pixels in the two anchors is set as the eye state threshold. Taking the midpoint of the anchors as a threshold yields optimal separation in the training set, in the aspect of minimizing the generalization error. Such a threshold is conceptually similar to the one-dimensional case in the support vector machine (SVM) classification algorithm (Cortes and Vapnik, 1995). The calculation of the eye state threshold completes the training stage, and the algorithm now has all the necessary data for classification of the entire day of the experiment. The entire training stage takes the user about 30 s, and contains all the human-based decision input to the process.

The eye state classification is obtained by calculating the number of black pixels for each eye field of the entire video sequence, based on the user's chosen brightness threshold. Each eye field is then classified according to the eye state threshold: eye fields with more black points than the eye state threshold are classified "open", and ones with fewer black points than the eye state threshold are classified "closed". This classification is 'hard' in the sense that a decision is forced. Therefore, eye fields that could be perceived by a human observer as inconclusive are also classified according to the number of black pixels. Using our hardware (2.8 GHz Pentium with 2 GB of RAM), the classification of a 2-h video took about 20 min (Matlab m files can be found at http://basalganglia.huji.ac.il/assets/ESD.zip).

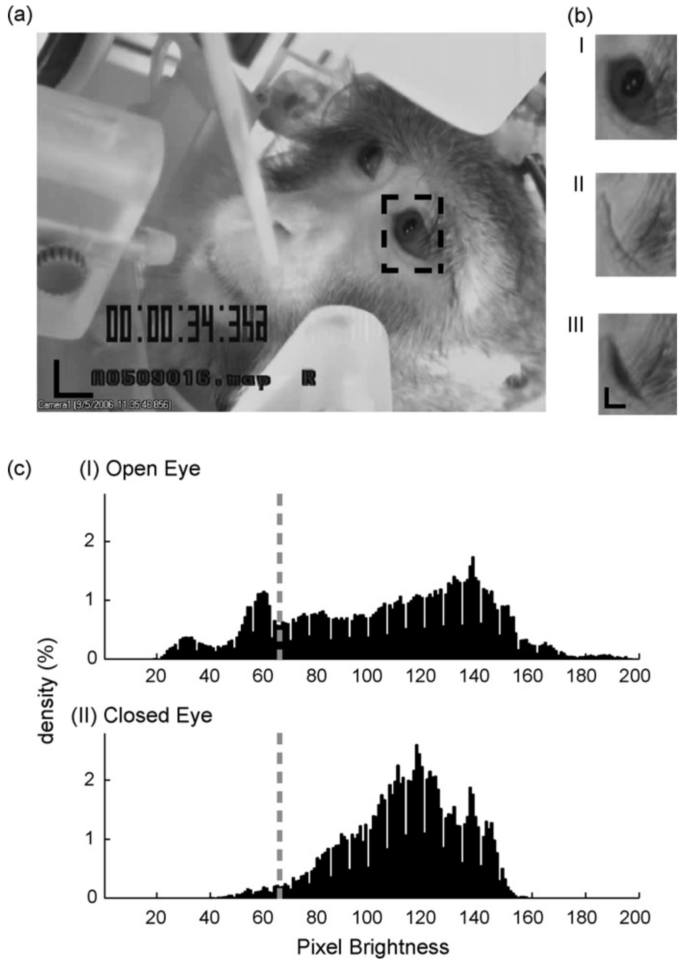

Fig. 2. The training stage of the eye state detection algorithm (ESD). (a) Arbitrarily chosen frame presented to the user, and the marked location of the eye. Scale indicates 50 pixels horizontally and vertically. (b) Three randomly chosen eye fields that were categorized by the user as open (I), closed (II) and inconclusive (III). Scale indicates 20 pixels horizontally and vertically. (c) Brightness histogram of the open (top) and closed (bottom) eye fields of the training set. The vertical dashed line is the brightness value that was chosen by the user as the threshold for the categorization of the entire video.

3. Results

3.1. Algorithm performance and stability

The eye state detection algorithm is based on two thresholds. As described above, the first is the brightness threshold, which determines how dark (on a scale of 0–255) a pixel needs to be so as to be considered black and is set by the user during the training stage. The second is the eye state threshold, which determines the state of the eye according to the number of black pixels, and is calculated automatically (based on the user's classification during the training stage). Note that these two thresholds could cause performance instability in the algorithm, since they are influenced by the features of the training set and the user's decisions. Therefore, we calculated

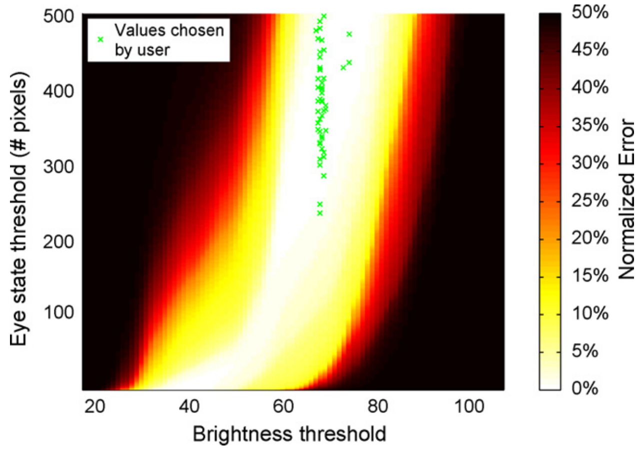

Fig. 3. ESD algorithm error as a function of the two thresholds. Normalized error is shown (color coded) as a function of the two types of classification thresholds used in the ESD algorithm (brightness threshold and eye state threshold). Normalized error is the probability of a classification error, assuming equal probability of open and closed eye. Brightness threshold is the threshold which defines which pixels are dark enough and is determined by the user. The eye state threshold defines how many black pixels (as defined by the first threshold) are enough to determine that the eye is open, and is determined by the algorithm according to the human decisions on the training set. The green Xs denote brightness thresholds chosen by the user, and the corresponding number of eye state thresholds for different training sets. This was done by repeating the algorithm 50 times with a training set of 20 images chosen randomly. Although the eye state threshold has a relatively large variance, all repetitions produced a normalized error of 0–1.5%. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

the algorithm error as a function of these two thresholds. We chose a 20,000 field (400 s) video, and one of the experimenters (RM) classified its fields manually into open eye fields (88.0%), closed eye fields (8.9%) and inconclusive eye fields (3.1%). The inconclusive eye fields were discarded in this analysis, since the human observer performance was considered the gold standard.

Two types of error are possible: detecting a closed eye when it is open, and vice versa. We were interested in performance independent of the eye state statistics in the chosen video. Therefore we used a normalized error (E_N), which is identical to the probability of an error in the case of equiprobability of open and closed eyes:

$$
\begin{aligned}
E_N &= \frac{1}{2}\frac{P(\text{classify open, closed eye})}{P(\text{closed eye})} + \frac{1}{2}\frac{P(\text{classify closed, open eye})}{P(\text{open eye})} \\
&= \frac{1}{2}\left[\frac{P(\text{classify open, closed eye})}{P(\text{closed eye})} + \frac{P(\text{classify closed, open eye})}{P(\text{open eye})}\right]
\end{aligned}
$$

By the definition of conditional probability, this is also the average of the conditioned probability of an error:

$$E_N = \frac{1}{2}[P(\text{classify open}|\text{closed eye}) + P(\text{classify closed}|\text{open eye})]$$

Each run of the algorithm produces two values, as described above: a brightness threshold selected by the user, and an eye-state threshold calculated by the algorithm according to the user's open/close classification in the training set. To show how the normalized error depends on these two values, Fig. 3 plots the normalized error as a function of them. The figure shows that a large range of these two thresholds results in a relatively low error; hence the algorithm's performance is stable over a large range of thresholds.

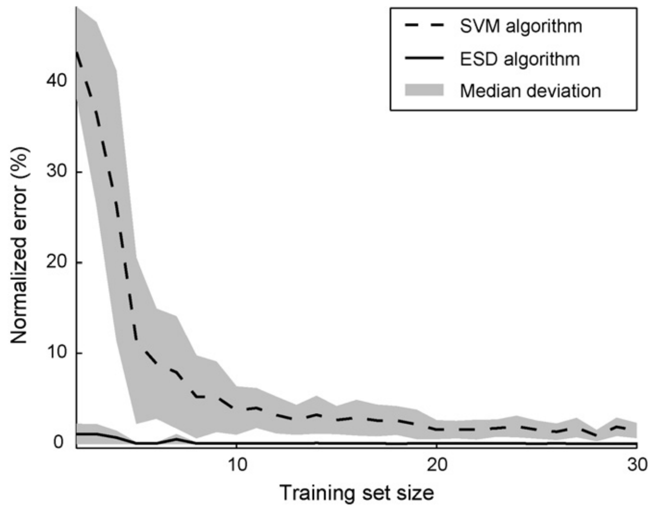

Fig. 4. ESD algorithm performs more accurately than the SVM algorithm and needs a much smaller training set. The median of the normalized error of the ESD algorithm (continuous), and the SVM algorithm (dashed) ± the median absolute deviation (gray shadow), calculated following 10 repeats per training set size. The training sets were chosen randomly, while keeping a constant ratio of closed and open eye fields. SVM fails to match the performance of the ESD algorithm for median performance, variability and dependence on training size.

Next, in order to verify that the thresholds generated during the training stage of the algorithm actually resulted in a low normalized error, we ran the semi-automatic algorithm 50 times with different randomly chosen training sets of 20 conclusive eye fields (i.e., open or closed) while applying the same open and closed eye statistics as in the entire video (i.e. 18 open eye fields and 2 closed). The resulting threshold pair of each run is plotted as an '*x*' in Fig. 3. The brightness thresholds have a relatively low range of values, whereas the eye-state thresholds have a higher range of values. This is due to the low variance in the brightness of the iris and the pupil, in comparison to a higher variance in the number of pupil and iris pixels. Nevertheless, the number of classification errors for this range of thresholds falls within the span of a relatively small error. The normalized classification error had a mean of 0.15%, and a maximal value of 1.4%.

Next, to assess the semi-automatic algorithm's dependence on the size of the training set, we trained it on different training set sizes, ranging from two to thirty eye fields. Again, this was done by forcing the statistics of the entire video on the training sets (while keeping at least one eye field of each conclusive type). This was repeated 10 times for each training set size. The median ± absolute median deviation of the classification's normalized error is plotted in Fig. 4 (continuous line). This shows an impressively low normalized error median even on a training set with a single eye field for each category, and a negligible error with training set as small as 8 eye fields.

To demonstrate the robustness of the ESD algorithm, we compared its performance to the SVM algorithm, which is an optimal linear classification algorithm, in the sense of minimizing generalization error. Briefly, this algorithm finds the linear classifier of two clusters (in an *n*-dimensional hyperplane) with maximal margins (Cortes and Vapnik, 1995). We reshaped the eye field matrices to vectors, and used randomly chosen training sets of these vectors to train the SVM algorithm. We used the linear classifier to classify the entire video, and calculated the normalized error (identically to the normalized error of the ESD algorithm). This was repeated with different training set sizes, 10 times per size, and we obtained the median ± absolute median deviation (Fig. 4, dashed line). This shows that the ESD algorithm performs considerably better than the SVM algorithm. Although the difference decreases with the increase

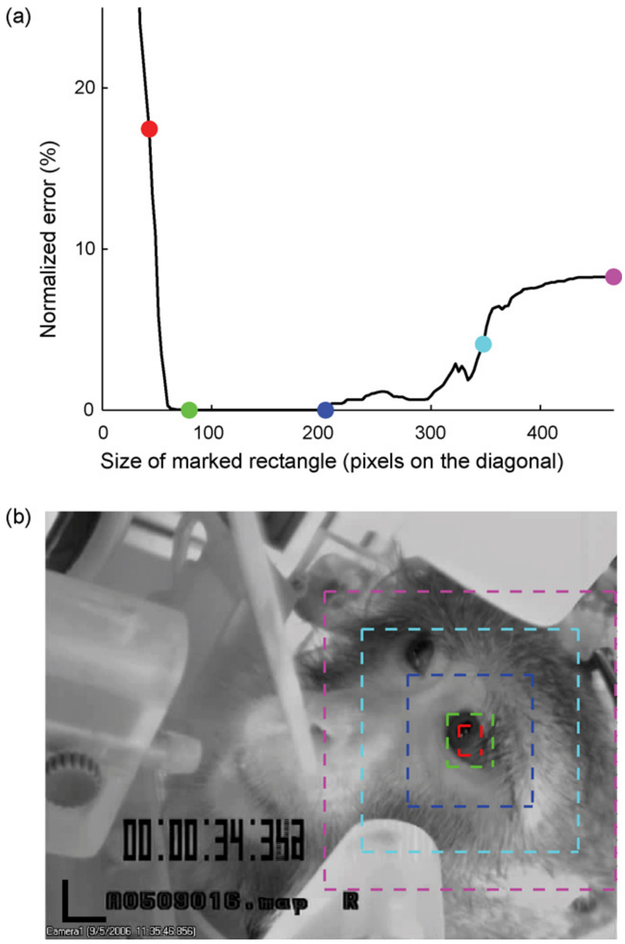

Fig. 5. ESD algorithm error as a function of the size of the rectangle which marked the eye location. (a) Normalized error as a function of the size of the eye field (in pixels on the diagonal). The algorithm was trained with the same training set, but with different sizes of eye rectangles. The normalized error of each eye rectangle size is shown as a single point on the curve. (b) An arbitrary field, with different sizes of marked rectangles. The rectangles correspond to the same color points as in (a). There is a broad area between the green and blue frames in which the error is zero, and a broader area between the green point and the cyan in which the error is less than 5%. The error only becomes considerable beyond this range. Scale indicates 50 pixels horizontally and vertically. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

in training set size, it remains noticeable even in larger ($n = 100$) training sets (data not shown). Furthermore, the ESD algorithm emerges as more reliable, since the deviation around the median error is smaller than the deviation obtained in SVM (Fig. 4, gray shadow).

Another user-selected parameter which could be an additional source of error is the location and size of the area of the eye (eye field). To test ESD algorithm's stability for different eye field sizes, we compared the algorithm's performance on a randomly chosen training set. Fig. 5 depicts the normalized error as a function of the eye field's rectangular size, measured by the length of diagonal (in pixels). The ESD algorithm maintained its level of performance for a large range of sizes, with a normalized error of less than 10% even for a large rectangle that occupied almost the entire monkey's face. Furthermore, a low level of error was maintained for a relatively small rectangle.

Finally, surveillance cameras are sensitive to visual light; hence changes in the luminance might affect the tool's performance. Luminance can go through high frequency modulations, e.g., because the camera's sampling and the refresh rate of the video screen (where visual stimuli were presented to the monkey) were not synchronized. Additionally, transient luminance changes may occur, e.g. due to opening of the recording chamber cover. To test the stability of the ESD algorithm as regards changes of luminance, we calculated the average brightness of the entire image (while omitting the area of the frame showing the time, see Fig. 2a bottom right) for each field. This is plotted in Fig. 6a, and reveals high frequency changes in luminance, as well as a transient change around 200 s. We manually split the video into brighter and darker periods, which are gray color coded in the figure. We ran the ESD algorithm 50 times with a training set of 20 fields, and calculated the average error per field. Fig. 6b shows the average normalized error in each of the video streams ± the standard deviation. The difference between the two error values was not significant (Student's *t*-test, $p = 0.19$), which shows nicely that the open/close classification error is negligibly affected by the general luminance.

3.2. *Possible applications of the ESD algorithm*

Eye state classification has potential for a wide variety of applications. We used the system on a delayed probabilistic classical conditioning task (Joshua et al., 2008). In this task, the monkey was repeatedly presented with one of a set of visual stimuli, each predicting an outcome with a different probability. Three stimuli predicted the administration of liquid food and three predicted the delivery of an airpuff with the same probabilities. Fig. 7a shows the percentage of the trials in which the eye is closed, in 20 ms bins, with respect to the administration of the outcome (food/airpuff). This indicates that the monkey closes its eyes to airpuffs but not to food. Furthermore, the monkey indeed learned to distinguish between these stimuli, as can be seen by the timing of the response which preceded the airpuff itself (Fig. 7a). The algorithm provides an elegant way of showing that the monkey's eye state has an increasingly higher probability of being closed as the time of the airpuff approaches.

We further used the algorithm to assess the blinking response to apomorphine (Apo) induced dyskinesias. Systemic injection of Apo, an ultra-fast dopamine agonist, induces orofacial dyskinesias which are known to include higher blinking rates (Blin et al., 1990; Kleven and Koek, 1996; Nevet et al., 2004). This is usually measured by human observers who count the number of blinks, a method which is prone to error and bias. Fig. 7b depicts the blinking rate of a monkey after intramuscular injection of 0.1 mg/kg Apo HCl 1%, as measured by the ESD algorithm. Closed eye events that lasted less than a second were defined as a blink. The blinks were counted in 1 min bins, and then smoothed with a Gaussian window with a standard deviation of 1 bin. As described in previous studies (Nevet et al., 2004, Fig. 1c), the blinking rate increased with the Apo administration, and remained so for at least 25 min.

4. Discussion

This manuscript describes a fast, simple, inexpensive and non-invasive tool for eye state detection during electrophysiological studies of primates. It is adapted to perform optimally in the typical setting of primate physiological studies; e.g., head fixation (Lemon, 1984). This type of tool is valuable for many primate studies, and can be easily adapted to most setups, since the only hardware it requires is a digital surveillance camera. The use of such camera, which is sensitive to the infrared wavelength and has its own source of infrared radiation, makes it possible to detect eye state under different illumination conditions, including darkness. Furthermore, using a digital camera rather than an analog device makes the eye state detection less vulnerable to electrical noise.

The temporal resolution of these cameras is usually 25–30 frames per second, which can be doubled by de-interlacing. This

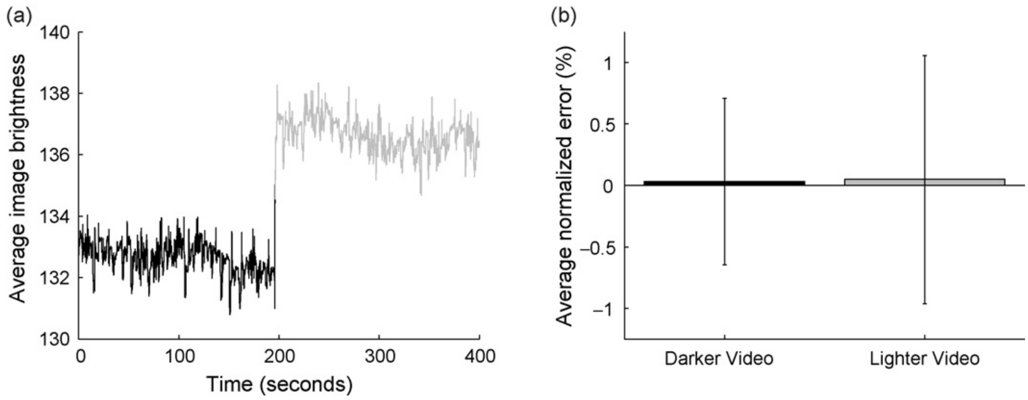

Fig. 6. ESD algorithm error as a function of general scene luminance. (a) Luminance in the entire scene as a function of time. Luminance was calculated by taking the average brightness value in each field during the 400-s video (omitting the area of the video image presenting the time). A striking increase in luminance can be seen around 200 s (due to opening of the recording chamber cover). This was manually marked and the video was split to a darker, early period (marked in black) and a lighter, late period (marked in gray). (b) The average normalized error of the darker and the lighter periods is presented (same color code as in a). The error was calculated by repeating the ESD algorithm 50 times with a training set of 20 eye fields. The error bars represent the standard deviation of the normalized error. Both periods present a very low error level, with no significant difference (Student's t-test, $p = 0.19$), indicating the low dependence of ESD performance on general scene luminance.

makes the tool suitable for detection of non-human primates blinks (Baker et al., 2002), although cases of blinks with a fractional closure of the eyelid (Rambold et al., 2005) might be missed. For higher temporal resolution, a faster camera is needed, rendering the system more expensive, but requiring no changes in the algorithm. Correct timing of the electrophysiological recording and the eye state calls for accurate synchronization between the video and the electrophysiological recording. In our setup this was done by feeding a time signal from the recording system into the digital video, presenting it on the bottom left corner (e.g., Fig. 2a), and detecting it offline. Other synchronization signals, such as an auditory signal for the camera, can be used as well.

Although detection of the state of a single eye was sufficient for our purposes, the algorithm could be easily adapted to detect the states of both eyes by marking their location and possibly detecting the two thresholds separately. This would also require proper positioning of the camera, such that both eyes are clearly visible.

Both ESD and SVM algorithms are supervised learning binary classifiers, but ESD performs considerably better than the SVM algorithm in detecting the eye state, although SVM is considered a very robust linear classifier. This is probably because the ESD algorithm makes assumptions regarding the data, whereas SVM does not. The ESD assumes that the number of pixels in the eye field that are sufficiently dark is a strong enough rule for the detection of the state of the eye. Naturally, when this assumption does not hold, ESD will perform worse than SVM.

Indeed, ESD algorithm's errors occur mostly when there are fewer dark pixels, e.g. when the eye is turned away from the camera, or the eye-lid is partially closed. This also accounts for the increase in error with the decrease in eye field size (Fig. 5a). Nevertheless, the ESD algorithm maintains a low level of error for a relatively small rectangle. This suggests that the algorithm is stable to physiological changes in the angle of the eye in which smaller areas of the pupil and iris are visible. ESD algorithm is also stable to changes in general scene luminance. This is because the iris, and even more so the pupil, have a high light absorbance, and therefore there are enough dark pixels in an open eye image, even with greater light intensity.

This semi-automatic procedure was found to be adequate for our needs, because it was highly accurate and required very little training time. Therefore, we did not find it necessary to make it fully automatic. However, fully automatic algorithms that detect the location of the eye for a moving human face have been reported by other researchers (e.g., Craw et al., 1992), and could be adapted to our tool. This may enable using the tool for experiments that do not require a head restraint; e.g., with chronically implanted electrodes (Nordhausen et al., 1996; Jackson and Fetz, 2007).

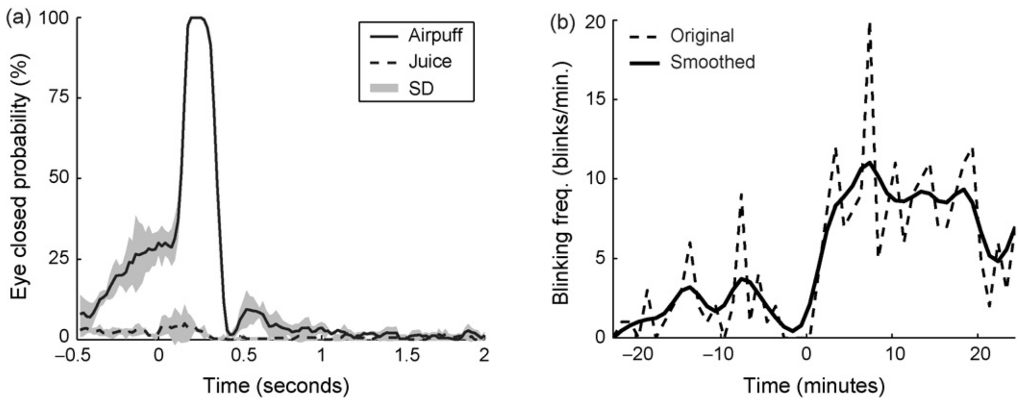

Fig. 7. Possible applications of the eye state detection tool. (a) Eye closure to airpuff. The monkey was presented with a reward (liquid food) or an aversive stimulus (airpuff) at $t = 0$, after a reward- or aversion-predicting visual stimulus. The percentage of times that the animal kept the eye closed in 20 ms bins is plotted ± variance (shaded). The monkey closed its eyes in anticipation and following the aversive stimulus, but not for the reward. (b) Apomorphine induced dyskinesias increases blinking frequency. Blinking rate was measured by counting eye closures lasting less than a second, before and after the Apomorphine injection (at $t = 0$). The dashed line is the raw blinking rate per minute and the continuous line is the blinking rate after smoothing with a Gaussian window of one bin (1 min).

Finally, there is continual interest in the effect of arousal levels on neural activity (e.g., Steriade and McCarley, 2005). Eye state detection, supported by analysis of eye-movements, EEG and EMG provide an accurate estimation of arousal state. Moreover, it provides a reliable estimation of blinking rate, which is affected by many physiological and pathological processes. Thus overall, our tool provides a reliable, noninvasive and inexpensive method for detection of eye open/closed states, and is therefore a recommended add-on for primate electrophysiological setups.

Acknowledgement

This work was partly supported by a Hebrew University Netherlands Association grant entitled "Fighting against Parkinson", and the Harry and Sylvia Hoffman leadership and responsibility program. We would like to thank E. Singer for language editing.

References

Baker RS, Radmanesh SM, Abell KM. The effect of apomorphine on blink kinematics in subhuman primates with and without facial nerve palsy. Invest Ophthalmol Vis Sci 2002;43:2933–8.

Benoit A, Caplier A. Hypovigilence analysis: open or closed eye or mouth? Blinking or yawning frequency? In: IEEE conference on advanced video and signal based surveillance, AVSS; 2005. p. 207–12.

Blazquez PM, Fujii N, Kojima J, Graybiel AM. A network representation of response probability in the striatum. Neuron 2002;33(6):973–82.

Blin O, Masson G, Azulay JP, Fondarai J, Serratrice G. Apomorphine-induced blinking and yawning in healthy volunteers. Br J Clin Pharmacol 1990;30:769–73.

Bracha V, Nilaweera W, Zenitsky G, Irwin K. Video recording system for the measurement of eyelid movements during classical conditioning of the eyeblink response in the rabbit. J Neurosci Methods 2003;125:173–81.

Clark RE, Zola S. Trace eyeblink classical conditioning in the monkey: a nonsurgical method and behavioral analysis. Behav Neurosci 1998;112:1062–8.

Cortes C, Vapnik V. Support-vector networks. Mach Learn 1995;20:273–97.

Craw I, Tock D, Bennett A. Finding face features. In: Goos G, Hartmanis J, editors. Computer vision – ECCV'92. Berlin: Springer; 1992. p. 92–6.

Durkin M, Prescott L, Jonet CJ, Frank E, Niggel M, Powell DA. Photoresistive measurement of the Pavlovian conditioned eyelid response in human subjects. Psychophysiology 1990;27:599–603.

Heishman R, Duric Z. Using image flow to detect eye blinks in color videos. In: IEEE workshop on applications of computer vision, WACV; 2007. p. 52.

Jackson A, Fetz EE. Compact movable microwire array for long-term chronic unit recording in cerebral cortex of primates. J Neurophysiol 2007;98:3109–18.

Joshua M, Adler A, Mitelman R, Vaadia E, Bergman H. Midbrain dopaminergic neurons and striatal cholinergic interneurons encode the difference between reward and aversive events at different epochs of probabilistic classical conditioning trials. J Neurosci 2008;28:116.73–84.

Karson CN. Spontaneous eye-blink rates and dopaminergic systems. Brain 1983;106(Pt 3):643–53.

Kleven MS, Koek W. Differential effects of direct and indirect dopamine agonists on eye blink rate in cynomolgus monkeys. J Pharmacol Exp Ther 1996;279:1211–9.

Lemon RN. Methods for neuronal recording in conscious animals. In: IBRO handbook series: methods in neurosciences. London: Wiley; 1984.

Marquis DG, Hilgard ER. Conditioned responses to light in monkeys after removal of the occipital lobes. Brain 1937;60:1–12.

Miyakawa T, Takano H, Nakamura K. Development of non-contact real-time blink detection system for doze alarm. In: SICE Annual Conference; 2004. p. 1626–31.

Nevet A, Morris G, Saban G, Fainstein N, Bergman H. Discharge rate of substantia nigra pars reticulata neurons is reduced in non-parkinsonian monkeys with apomorphine-induced orofacial dyskinesia. J Neurophysiol 2004;92:1973–81.

Nordhausen CT, Maynard EM, Normann RA. Single unit recording capabilities of a 100 microelectrode array. Brain Res 1996;726:129–40.

Pennypacker HS, King FA, Achenbach KE, Roberts L. An apparatus and procedure for conditioning the eye-blink reflex in the squirrel monkey. J Exp Anal Behav 1966;9:601–4.

Ponder E, Kennedy WP. On the Act of Blinking. Q J Exp Physiol 1927;18:89–110.

Porter JD, Stava MW, Gaddie IB, Baker RS. Quantitative analysis of eyelid movement metrics reveals the highly stereotyped nature of monkey blinks. Brain Res 1993;609:159–66.

Rambold H, El Baz I, Helmchen C. Blink effects on ongoing smooth pursuit eye movements in humans. Exp Brain Res 2005;161:11–26.

Robinson DA. A method of measuring eye movement using a scleral search coil in a magnetic field. IEEE Trans Biomed Eng 1963;10:137–45.

Silverstein LD, Graham FK, Eyeblink EMG. a miniature eyelid electrode for recording from orbicularis oculi. Psychophysiology 1978;15:377–9.

Steriade M, McCarley R. Brain control of wakefulness and sleep. 2nd ed. New York: Springer; 2005.

Stevens JR. Eye blink and schizophrenia: psychosis or tardive dyskinesia? Am J Psychiatry 1978;135:223–6.

Tan H, Zhang YJ. Detecting eye blink states by tracking iris and eyelids. Pattern Recog Lett 2006;27:667–75.

Thompson LT, Moyer JR, Akase E, Disterhoft JF. A system for quantitative analysis of associative learning. Part 1. Hardware interfaces with cross-species applications. J Neurosci Methods 1994;54:109–17.

Tian Yl, Kanade T, Cohn J. Eye-state action unit detection by gabor wavelets. In: Advances in multimodal interfaces – ICMI 2000. Berlin: Springer; 2000. p. 143–150.